Seeing the Past with Computers

Experiments with
Augmented Reality
and Computer Vision
for History

Kevin Kee and Timothy Compeau, Editors

digital humanities

Seeing the Past with Computers

DIGITAL HUMANITIES

Series Editors:

Julie Thompson Klein, Wayne State University
Tara McPherson, University of Southern California
Paul Conway, University of Michigan

Seeing the Past with Computers: Experiments with Augmented Reality and Computer Vision for History

Kevin Kee and Timothy Compeau, Editors

A World of Fiction: Digital Collections and the Future of Literary History

Katherine Bode

Stamping American Memory: Collectors, Citizens, and the Post

Sheila A. Brennan

Big Digital Humanities: Imagining a Meeting Place for the Humanities and the Digital

Patrik Svensson

Manifesto for the Humanities: Transforming Doctoral Education in Good Enough Times

Sidonie Smith

Ethical Programs: Hospitality and the Rhetorics of Software

James J. Brown Jr.

Digital Rhetoric: Theory, Method, Practice

Douglas Eyman

Web Writing: Why and How for Liberal Arts Teaching and Learning

Jack Dougherty and Tennyson O'Donnell, Editors

Interdisciplining Digital Humanities: Boundary Work in an Emerging Field

Julie Thompson Klein

Pastplay: Teaching and Learning History with Technology

Kevin Kee, Editor

Writing History in the Digital Age

Jack Dougherty and Kristen Nawrotzki, Editors

Hacking the Academy: New Approaches to Scholarship and Teaching from Digital Humanities

Daniel J. Cohen and Tom Scheinfeldt, Editors

DIGITALCULTUREBOOKS, an imprint of the University of Michigan Press, is dedicated to publishing work in new media studies and the emerging field of digital humanities.

Seeing the Past with Computers

*Experiments with Augmented Reality
and Computer Vision for History*

KEVIN KEE AND
TIMOTHY COMPEAU, EDITORS

UNIVERSITY OF MICHIGAN PRESS
ANN ARBOR

This work is licensed under a Creative Commons Attribution 4.0 International License. Note to users: A Creative Commons license is only valid when it is applied by the person or entity that holds rights to the licensed work. Works may contain components (e.g., photographs, illustrations, or quotations) to which the rightsholder in the work cannot apply the license. It is ultimately your responsibility to independently evaluate the copyright status of any work or component part of a work you use, in light of your intended use. To view a copy of this license, visit <https://creativecommons.org/licenses/>.

Library of Congress Cataloging-in-Publication Data

Names: Kee, Kevin B. (Kevin Bradley), 1969– editor. | Compeau, Timothy, 1981– editor.

Title: Seeing the past with computers : experiments with augmented reality and computer vision for history / Kevin Kee and Timothy Compeau, editors.

Description: Ann Arbor : University of Michigan Press, [2019] | Series: Digital humanities | Includes index. | Identifiers: LCCN 2018034886 (print) | LCCN 2018042822 (ebook) | ISBN 9780472124558 (E-book) | ISBN 9780472900879 (Open Access) | ISBN 9780472131112 (hardcover : alk. paper) |

Subjects: LCSH: Digital humanities. | Augmented reality. | Computer vision.

Classification: LCC AZ105 (ebook) | LCC AZ105 .S4 2019 (print) | DDC 001.30285—dc23

LC record available at <https://lcn.loc.gov/2018034886>

doi: <https://dx.doi.org/10.3998/mpub.9964786>

The editors gratefully acknowledge the generous financial support of the University of Ottawa, the Canada Research Chairs program, the Social Sciences and Humanities Research Council of Canada, and Brock University, which provided the time and funds to bring the *Seeing the Past with Computers* project to completion.

Contents

Introduction: Seeing the Past <i>Kevin Kee and Timothy Compeau</i>	I
ONE: The People Inside <i>Tim Sherratt and Kate Bagnall</i>	II
TWO: Bringing Trouvé to Light: Speculative Computer Vision and Media History <i>Jentry Sayers</i>	32
THREE: Seeing Swinburne: Toward a Mobile and Augmented-Reality Edition of <i>Poems and Ballads</i> , 1866 <i>Bethany Nowviskie and Wayne Graham</i>	50
FOUR: Mixed-Reality Design for Broken-World Thinking <i>Kari Kraus, Derek Hansen, Elizabeth Bonsignore, June Abn, Jes Koepfler, Kathryn Kaczmarek Frew, Anthony Pellicone, and Carlea Holl-Jensen</i>	69
FIVE: Faster than the Eye: Using Computer Vision to Explore Sources in the History of Stage Magic <i>Devon Elliot and William J. Turkel</i>	83
SIX: The Analog Archive: Image-Mining the History of Electronics <i>Edward Jones-Imbotep and William J. Turkel</i>	95

SEVEN: Learning to See the Past at Scale: Exploring Web Archives through Hundreds of Thousands of Images <i>Ian Milligan</i>	116
EIGHT: Building Augmented Reality Freedom Stories: A Critical Reflection <i>Andrew Roth and Caitlin Fisher</i>	137
NINE: Experiments in Alternative- and Augmented-Reality Game Design: Platforms and Collaborations <i>Geoffrey Rockwell and Sean Gouglas</i>	158
TEN: Tecumseh Returns: A History Game in Alternate Reality, Augmented Reality, and Reality <i>Timothy Compeau and Robert MacDougall</i>	176
ELEVEN: History All Around Us: Toward Best Practices for Augmented Reality for History <i>Kevin Kee, Eric Poitras, and Timothy Compeau</i>	207
TWELVE: Hearing the Past <i>Shawn Graham, Stuart Eve, Colleen Morgan, and Alexis Pantos</i>	224
<i>Contributors</i>	237
<i>Index</i>	245

Digital materials related to this title can be found
on the Fulcrum platform via the following URL:
<https://doi.org/10.3998/mpub.9964786>

Introduction

Seeing the Past

Kevin Kee and Timothy Compeau

We live in a world of seeing computers. From airport security to social media to games, machines that see are embedded in our lives in subtle ways that, ironically, go unseen. For better or worse, face-recognition software and algorithms help law enforcement and security agencies track individuals in crowds; toll road cameras snap photographs of passing license plates to automatically bill the owners; computer imaging helps medical professionals diagnose patients; and smartphone cameras and GPS help game players capture “virtual creatures” in the “real world.” From the mundane to the life-altering, seeing computers are actively helping people see the present. The same technology can also be used to see the past, and in this book we explore the ways in which we can turn their gaze toward the study and communication of history.

In late 2014, a group of digital humanists met in Niagara-on-the-Lake, Ontario, to discuss how seeing technologies might be harnessed for historical research and teaching. Historians and archaeologists have all wished, at some point, that they could see the past firsthand. The confusion and ambiguity created by fragmentary sources tend to equally excite our doubts and imagination—we worry about the veracity of our sources, and we wonder what might exist in the gaps in our evidence. While that immediate encounter with the past remains a subject for science fiction writers, the scholars in this volume have taken up the challenge of exploring how seeing technolo-

gies can add to, and change our understanding of, the practice of history. *Seeing the Past with Computers* demonstrates new and innovative ways of accessing, understanding, writing about, teaching, and sharing history.

We focus on two related forms of seeing technology that are changing how some humanists work, but remain untapped and confusing for most scholars and students: computer vision and augmented reality. Computer vision (CV) is a technology that can access, process, analyze, and understand visual information. Consider, for instance, optical character recognition (OCR), which allows computers to read text from digitized print sources. Whereas scholars used to read a few books deeply (“close reading”), OCR has facilitated what Franco Moretti called “distant reading,” helping us mine and analyze thousands of books across eras, genres, and subjects.¹ Such quantitative approaches to textual analysis have their critics, but they also hold many lessons for those interested in history. Yet history involves more than just the textual evidence historians have traditionally privileged; traces of the past are also embedded in the visual—photographs, paintings, sketches—and material culture. The proliferation of digitized visual sources presents historians with exciting new technical and theoretical problems and opportunities. The scholars in this collection offer ways of thinking about where we might look for source material, and how we might use CV to analyze those sources, in the context of our research or teaching, to ensure broader, deeper, and more representative understandings of the past.

Computer vision supports augmented reality (AR), which overlays digital content onto the real-world objects that a computer sees, thereby “augmenting” our view of our surroundings. AR is also helping us study, see, and share history in compelling new ways. Although historians have long been interested in augmenting the present with the past through plaques and monuments, digital technologies increase the possibilities. For example, AR at museums or historical sites can provide on-the-spot supplementary information, telling hidden stories or offering lesser-known facts to enhance the visitor’s knowledge of an item or space. Historic photographs, 3D reconstructions, and re-enactments of historic events can be overlaid onto present-day geography, demonstrating how an object or environment has changed over time. And GPS-linked AR can prompt location-specific historical data, guiding visitors at sites of interest, and increasing their sensory engagement with a place. AR applications can provide a seemingly magical way to tease out the history all around us, allowing nonspecialists to see the past in a manner similar to historians and archaeologists.²

Other related experimental technologies have begun to be explored by historians: haptic interfaces, which privilege touch, and scent-scapes, which privilege smell, are two intriguing areas of sensory research. Our book focuses primarily on sight; AR and computer vision employ similar principles of seeing technology, and have reached a place in their development where they can be hacked and harnessed by the nonspecialist. In the pages that follow we explore the ways in which they can be employed in the archive, in the museum, in the classroom, and anywhere else that the imagination leads.

Most developers of CV and AR would not have envisioned the uses described in this book, and most historians will not have conceived of applying CV and AR to historical research and presentation. Alexander Manu, a writer specializing in the corporate technology industry, describes the space “between current capability and future possibility” as “the imagination gap.” For the corporate world, this is a warning to innovate or die. The stakes are not quite so high for university-based researchers, but the rewards can be just as impressive. Whether in a corporation, university, or memory institution, the need for freedom to experiment, speculate, and play is vital.³ As Manu notes, breakthroughs require a “willing trickster and a curious audience.”⁴ The chapters and experiments described in this book are linked by our desire to bridge one imagination gap in history. The designers who developed facial recognition software to monitor people in airports did not imagine that the technology would be applied to archival collections. Nor did the developers of AR for advertising imagine that their platform could be adapted for students inquiring into explanations that challenge the narratives found in their textbooks. Those who are interested in understanding and communicating the past now have an opportunity to embrace these technologies and bridge the imagination gap through the kind of speculative experimentation and pure investigation described in this book.

Indeed, in the chapters that follow we go further, and contend that seeing technologies are becoming *essential* tools for historians. The explosion of digitized and born-digital sources requires the use of computer vision. Some of these sources are text; others are images. As Ian Milligan writes in “Learning to See the Past at Scale,” most historians are “unprepared to engage with the quantity of digital sources that will fundamentally transform our profession.” To provide just one example, Dan Cohen has demonstrated the challenge of this multiplication of sources by noting that a scholar who wants to write a history of the Lyndon Johnson White House

must read and analyze the 40,000 memos issued during the president's administration. This will take time, but it is possible. In contrast, a historian writing about the Clinton White House must address four million emails (in addition to conventional administration documents); it is impossible to read these in one lifetime. The archives of the Bush White House, for its part, contain 200 million emails, and those of the Obama White House number many more. Without seeing computers, researchers will be incapable of sifting through this material appropriately, and unable to effectively write a history of the Clinton, Bush, Obama, or Trump White Houses. If historians expand their investigations to include the memes, photographs, cartoons, and video produced and shared *about* the White House, the challenge becomes greater. For this reason, the future of historical research requires the pioneering efforts of scholars like those in this book.

As the possibilities for exploring historical photographs, drawings, printed images, and the seemingly infinite images available on the internet continue to grow,⁵ we have a unique opportunity to use this technology in ways that enhance the study of history. Developing seeing technologies in field-sensitive ways necessarily requires that historians become familiar with these tools. These emerging forms of analysis and communication can then be married to, or used alongside, the time-honed and familiar methods of research, writing, and teaching. The creation of the new does not require the destruction of the old, and the chapters in this volume highlight the continued need for trained historians to use time-honored techniques to interpret and share their evidence. The technologies and techniques explored in this book revisit old problems and provide new answers, but generate ever more questions.

Computer vision, for example, prompts us to consider questions we may have never thought to ask: Are there cultural patterns lurking unseen in the infinite archive of the internet that can help us understand our world? Can we better grasp the mentalities of historical actors by dissecting their visual creations, when we view these creations in aggregate (by the thousands or millions)? Can we teach our students to think historically with these new methods? Seeing technologies also encourage us to experiment with methodologies in and beyond our field. In wrestling with such questions, and by embracing speculative investigation, the scholars in this book push the boundaries of what it means to “do history.” Our discipline includes history as generally understood, but also archaeology, science and technology studies, and Big History (from the Big Bang to the present). The objects of our study may be paper documents, but also Victorian jew-

ellery or circuit-board schematics. Our archives may include traditional folia, but also the GeoCities internet archive or vast collections of digitized photographs. And the scholarship produced may take the form of a book or an article, but also a reorganized archive, a game, or digital recreation. And while these early experiments are necessarily limited in scope, and the results sometimes make clearer our methodological shortcomings as much as our advances, our work makes one thing evident: new ways of seeing are also new ways of thinking.

Seeing the Past with Computers is in many ways a sequel to *PastPlay: Teaching and Learning with Technology* (2014), and we return to some of the ideas explored in that volume.⁶ Above all, however, this book is a testament to the power of playful experimentation with technology and techniques in our discipline, and in other domains of inquiry, simply to see what happens. None of the scholars in this collection set out to argue against more conventional ways of doing or presenting history; neither do they aim to fix deficiencies, nor lament the failings in their disciplines. Instead, we seek to explore our sources in new ways and to show how emerging technologies can enhance our understanding. Not every experiment has yielded earth-shattering results, but individually and collectively, the chapters in this collection leave us more convinced that breakthroughs emerge from a culture that values experimentation for its own sake.

Our opening chapter serves as a vignette to illustrate the spirit that runs through each of the essays that follow. “The People Inside” is the fruit of experimentation, collaboration, and a do-it-yourself attitude. Motivated by a desire to “see archive records differently,” but “without funding, without research partners, without timelines,” Tim Sheratt and Kate Bagnall recognized that it was “impossible to create and sustain a new research project. Until we just *did* it.” Though historians have studied “White Australia” policies, Sheratt and Bagnall use photographs as their entry point for examining the human cost of a racist and exclusionary immigration system. Employing computer vision and face-detection techniques to explore and present thousands of photographs attached to individual immigration files in the Australian National Archives, Sheratt and Bagnall reveal those affected by exclusionary policies in an immediate and touching way. Considering that the stiff institutional style of the mountains of files pertaining to government policies has a way of distancing and depersonalizing records, Sheratt and Bagnall’s work shows that human stories need not be lost when exploring vast data sources with computer vision; in fact, the authors demonstrate how technology can actually shine a light on individual lives.

Sheratt and Bagnall's work is proudly speculative, like virtually every project described in this book. In "Bringing Trouvé to Light," Jentery Sayers articulates and promotes the concept of "speculative computer vision" as a way of "multiplying how scholars see the past." His chapter describes how technology is helping to explore the little-known world of Victorian-era electronic jewelry, examining a mysterious skull pin that was said to move under power from tiny Victorian batteries. Because no one alive has ever seen the pin move and because the extant remains are inoperable, some historians question whether the jewelery ever truly worked. Sayers's research thus reveals how computer vision can serve not as an instrument of confirmation, but of experimentation and speculation. Using 3D scanning and modelling, Sayers demonstrates how these aspects of computer vision can be the key to uncovering how lost or broken machines once functioned, bringing us closer to uncovering the secrets of past electronic wizardry.

A speculative computing thread also runs through Bethany Nowviskie and Wayne Graham's exploration of a controversial book of Victorian poetry, but in this case AR is used to reveal how small changes between editions can expose profound ethical and cultural issues in the nineteenth-century publishing industry. Addressing the benefits and perils inherent in digitizing print media of the past, Nowviskie and Graham's work is a fascinating examination of how AR can be used to preserve, and even enhance, the way books as physical artifacts can preserve their own material and cultural history.

Kari Kraus and her team of researchers challenge the "completionist" approach historians and archaeologists often take when using AR, and highlight the potential of speculative research in the context of game development. Rather than use AR applications to reconstruct historical objects or landscapes, Kraus and her team focus on what happens when the technology is used to take things apart; in other words, what knowledge is gained from broken things? Situating their research in the 2015 AR game *DUST*, a collaborative effort between researchers at the University of Maryland and NASA to explore Big History concepts with young students, Kraus and her team show how the principles of "broken world thinking" can be used to explore the relationship and structure of parts within a whole.

Although broken things may confound and perplex humans, computer vision can convert degraded, altered, and decontextualized images into pattern-recognition algorithms that allow for new methods of seeing and understanding visual evidence. The next three chapters offer case studies

in seeing the past with computer vision, with a focus on the late nineteenth century, the early Cold War, and the heady days of the young internet. Each of the experiments by Elliott, Turkel, Jones-Imhotep, and Milligan generates patterns to better understand their respective material. Devon Elliott and William J. Turkel, for example, use computer vision to mine and “liberate” tens of thousands of pictures from digitized printed material to explore how magicians shared (or concealed) their secrets through images. Using digital techniques that are themselves “faster than eye,” Elliott and Turkel reveal intriguing patterns that tell us how networks of professional and amateur slight-of-hand artists shared and honed their craft during the golden age of Vaudeville.

Edward Jones-Imhotep and William Turkel’s work uses a similar technique, but with Mathematica, a powerful computational program originally designed for science and engineering applications, to explore another period of rapid technological change. By mining huge visual data sets of schematics and plans drawn by the pioneers of advanced electronics, Jones-Imhotep and Turkel reveal the confusion, anxieties, and struggles the profession faced as its practitioners argued over how to articulate and conceptualize the circuits that became the foundation of modern life. In doing so, Jones-Imhotep and Turkel provide new insights into this important period of technological evolution, while confirming the illuminating potential of applying seeing and learning machines to large sets of visual data.

Ian Milligan also uses Mathematica in his research, this time to make sense of the seemingly infinite digital images produced in the early years of the internet. The World Wide Web presents a vast new frontier for humanists, but the internet’s sheer size, complexity, and ephemeral nature pose significant hurdles. Milligan transforms these obstacles into opportunities by applying computer vision to scrape and assemble collections of thousands of images, revealing the unique patterns and commonalities formed in online communities in ways a single pair of eyes never could.

Computers help historians see the past, but historians can also help their students and public audiences see the past by using computers. Much like film and radio a century or more ago, AR presents a new exciting medium for teaching and learning, for play, and for telling stories.⁷ The chapters in this collection showcase applications that can help our students and the interested public immerse themselves in historical environments and ideas.

The success of these endeavours relies on creating the right partnerships between humanists and digital designers. In addition to showing how

historians experiment and think about seeing technologies, designers Caitlin Fisher and Andrew Roth share their experiences working with historians to produce AR learning opportunities for the classroom. Fisher and Roth's chapter describes the process of working with a team of historians to create AR exercises that explore America's Underground Railroad. Their team produced a prototype of a viable AR platform and workflow that others could use and adapt to create their own educational experiences, but Fisher and Roth's work also serves as a clarion call: historians must participate in the design process as much as they engage in the research informing the project. These are sentiments shared by Geoffrey Rockwell and Sean Gouglas, who explore their efforts to balance education, fun, and the experience of discovery within alternate and augmented reality place-based games. As they discuss in their chapter, experiments did not always meet with success, but were crucial steps to learning how to work with this emerging medium in the early years of the smart phone revolution.

Timothy Compeau and Robert MacDougall share their attempts to design games that teach not only history, but historical thinking skills. Employing often simple AR and alternate-reality techniques to create two versions of their game *Tecumseh Lies Here*, Compeau and MacDougall examine how AR can be used to encourage students to uncover and interpret the historical evidence for themselves, and to come to their own conclusions—in essence, to do the work of real historians.

Taking stock of the present state of AR applications found in museums and heritage sites, Kevin Kee, Eric Poitras, and Timothy Compeau consider some best practices for developing place-based AR in history. Extending their findings to classroom possibilities, the authors offer a plan to clinically test and evaluate the effect of AR applications on our students' learning and cognition, and to empirically determine the effectiveness of AR on the learning process.

Rounding out the chapters that explore seeing computers and seeing with computers, we conclude with an essay by Shawn Graham, Stuart Eve, Colleen Morgan, and Alexis Pantos, who remind us that AR is not limited to vision and that, perhaps, our fixation on the idea of *seeing* the past is little more than present-day occularcentrism: a vestigial relic of the Enlightenment hierarchy of the senses. In their work, this group of archaeologists suggests ways that aural AR could expand the sensory experience of historical sites. Experimenting with soundscapes of lost worlds to evoke emotions that sight cannot, the authors make a compelling case for their theory that, if we cannot see the past, perhaps we can hear it.

Just as seeing computers are altering how our world functions, computer vision and AR are broadening both the subject matter and the methods of historical research in ways that few predicted a generation ago. If the last half of the twentieth century saw historians develop new theoretical approaches, in which scholars, working alone, researched and treated history as text, the first half of the twenty-first century has witnessed teams of researchers working together to study and express history using a host of digital tools and media.

Framing the expansion of the discipline in this way gives the impression that one paradigm has replaced another. We suggest, however, that the use of computer vision and AR do not supplant previous methodologies and tools, so much as add to them. In addition to the cerebral and, sometimes, detached text-based treatment of history, seeing computers are contributing an immediate and occasionally affecting experience in which we not only see the past, but encounter history with multiple senses. The tools and media we use are becoming an additional form of scholarly publishing, adding to how historians and humanists tell stories and share knowledge with academic and public audiences.

Historians avoid making predictions about the future, but the work presented in this book represents a snapshot of the possibilities emerging in the midst of an unprecedented technological revolution. If this book has one overarching theory, it is that discovery comes from speculation and a playful approach to historical questions and problems. Imaginative experimentation with emerging technologies can generate conclusions that sometimes challenge the divisions between tried and tested theories, and in our case seeing technologies can help historians understand immense data sets at a distance, but also zoom in for a level of depth and engagement that was previously impossible. Both AR and computer vision show equal potential for near-simultaneous deep and distant seeing.

Recent decades have witnessed the “postmodern turn” in history, and more recently the “digital turn.” *Seeing the Past with Computers* shares some first efforts in the “visual turn” of digital history, and suggests pathways for future exploration. It also invites scholars and students to experiment with the methods described by a small group of colleagues, and to engage with us in the practice of critical reflection that has powered the humanities in the past and will propel them into the future. We look forward to new ways of seeing the past, so that together we can make sense of our world, and see across the gap that separates present and future history research and teaching.

Notes

1. Franco Moretti, *Distant Reading* (London: Verso, 2013). Also see Moretti's work at the Stanford Literary Lab, <http://litlab.stanford.edu/>
2. Alan B. Craig, *Understanding Augmented Reality: Concepts and Applications* (Waltham, MA: Elsevier, 2013); Robert MacDougal, "AR out of the Box," OARN 2012 Conference, Toronto, ON. <http://www.oarn.net/events/conference/videos/>, accessed July 28, 2016.
3. See Kevin Kee, ed., *Pastplay: Teaching and Learning History with Technology* (Ann Arbor, MI: University of Michigan Press, 2014) for explorations of play in history research and teaching.
4. Alexander Manu, *The Imagination Challenge: Strategic Foresight and Innovation in the Global Economy* (Berkeley, CA: New Rider, 2007), 15–17, 30.
5. For examples of the progress being made in AR development, see <http://www.augmentedworldexpo.com/>, accessed July 28, 2016.
6. Kee, ed., *Pastplay*.
7. thecreatorsproject.vice.com/blog/celebrating-georges-m%C3%A9li%C3%A8s-patron-saint-of-augmented-reality, accessed July 28, 2016.

CHAPTER I

The People Inside

Tim Sherratt and Kate Bagnall

Our collection begins with an example of computer vision that cuts through time and bureaucratic opacity to help us meet real people from the past. Buried in thousands of files in the National Archives of Australia is evidence of the exclusionary “White Australia” policies of the nineteenth and twentieth centuries, which were intended to limit and discourage immigration by non-Europeans. Tim Sherratt and Kate Bagnall decided to see what would happen if they used a form of face-detection software made ubiquitous by modern surveillance systems and applied it to a security system of a century ago. What we get is a new way to see the government documents, not as a source of statistics but, Sherratt and Bagnall argue, as powerful evidence of the people affected by racism.

In October 1911, the *Sydney Morning Herald* reported a local businessman’s complaints about his treatment by the Australian Customs Department. Charles Yee Wing, “a merchant of some standing, held in high esteem by Europeans and Chinese alike,” was planning a short trip to China.¹ He had applied to the department for a certificate that would allow him to re-enter Australia on his return but was annoyed when officials insisted that he be photographed “in various positions” to document his identity. A naturalized British subject, respectable family man, and long-term resident of Sydney, Charles Yee Wing objected to being treated “just like a criminal.”

Today we are accustomed to being identified by our image. Passports, driver’s licenses, student cards—we readily submit to being photographed for a variety of purposes, and we carry the images with us as proof that we

are who we say we are. The propagation and use of these likenesses has changed with the development of computer vision technology. Individual images can be discovered, analyzed, and compared across populations. The primary instrument of control has moved from document to database.

We are historians interested in bureaucratic systems for identification and control, and the impact of digital access on our understanding of how they worked. Kate's research explores the social and familial worlds of Chinese Australians, particularly those of mixed race, in the late nineteenth and early twentieth centuries. Tim is a hacker who uses digital technologies to open cultural collections to new forms of analysis and exploration. Together we have been focused on the vast collection of records generated by Australia's efforts to restrict non-European migration in the first half of the twentieth century. Among these records, preserved in the National Archives of Australia, are photographs and archival fragments documenting the life of Charles Yee Wing and thousands of others.

Computer vision can easily be used to find and recognize faces. Such technologies are often associated with the needs of law enforcement and national security, with the continued extension of systems for the identification and control of individuals. The latest facial recognition algorithms share a lineage with the thousands of immigration documents held by the National Archives. But can we use new technologies of identification to reveal the old? This chapter discusses an attempt to use facial detection technology to see archival records differently. What happens when instead of files and documents, systems and procedures, we see the people inside?

White Australia

Charles Yee Wing had a point in complaining about his treatment by the Customs Department. A century ago using portrait photography and fingerprinting to identify individuals was still fairly new, and until the early twentieth century, the most common official use of these technologies was to identify and manage criminals. Similar to law enforcement agencies in England, Europe, and the United States, the police in New South Wales, where Charles Yee Wing had lived since 1877, kept photographic gaol description books from around 1870.² The gaol description books placed "mug shots" of convicted criminals alongside biographic information and a physical description to identify and keep track of convicted criminals.³

By the 1890s Australia's colonial governments extended the use of these identification technologies to monitor and control the movement of peo-

ple across their boundaries, and these practices were continued on a national scale after Federation under the Immigration Restriction Act 1901. But not all travelers were treated equally under this new law. Passports in the modern sense were not introduced until later, during World War I, and this earlier regime targeted certain groups whose presence was seen to be at odds with white Australians' vision for their young nation.⁴ Charles Yee Wing's photographs identified his race as well as his face.

The Immigration Restriction Act remained in force, with amendments and a slight change of name, until 1958. The Act was the legislative backbone of what became known as the White Australia policy—a discriminatory system founded on the conviction that a strong and self-reliant Australia must, of necessity, be “white.”⁵ Yet the Act itself said nothing about “color” or “race.” It was, by design, a fairly inoffensive piece of bureaucratic machinery that empowered the Commonwealth to reject certain classes of immigrant, including convicted criminals, the physically or mentally ill, or those who were deemed morally unfit. The history of colonial cooperation and the movement to Federation told the real story, however, and debates surrounding the passage of the Act, both in Parliament and in the press, made the context explicit—“color” was crucial. In the words of Attorney-General and future Prime Minister Alfred Deakin in 1901, “The unity of Australia is nothing, if that does not imply a united race.”⁶

The practices of discrimination and exclusion at the heart of the Immigration Restriction Act were elaborated gradually through regulations, reviews, precedents, notes, and guidelines. Between 1902 and 1911, the head of the Department of External Affairs issued more than 400 circulars about immigration restriction to Customs staff,⁷ and while the Act may have fudged its racial dimensions, such advice to government officials did not. For example, one memorandum from 1936 plainly stated: “In pursuance of the ‘White Australia’ policy, the general practice is not to permit Asiatics or other coloured immigrants to enter Australia for the purpose of settling here permanently.”⁸

The principal instrument of exclusion under the Immigration Restriction Act was the innocuous-sounding Dictation Test. This test required an arriving passenger to write down a passage that was read to them in a European (later, any) language; failing the test meant deportation. To remove any misunderstanding of those administering it, the test's role was explained in a confidential note to Customs officials: “It is intended that the Dictation Test shall be an absolute bar to admission.”⁹ While the Act itself was silent on the details, officers were informed that all “persons of

coloured races” who were not otherwise exempted from the provisions of the Act would be subjected to the Dictation Test—and they would fail. The deterrent effect of the Dictation Test was striking. In 1902, 651 arriving travelers were tested but only 33 passed. In 1905, 107 were tested and just 3 passed. In 1913, 71 were tested and all failed.¹⁰ Ultimately people just stopped trying to come.

The Dictation Test was clearly effective at preventing new arrivals, but the system also had to accommodate the thousands of “colored” Australians who, like Charles Yee Wing, needed to be able to return home to Australia after traveling overseas. The first national census held in 1911 counted over 40,000 people of “non-European race” (not including Indigenous Australians) in the country, around 25,000 of whom were Chinese.¹¹ On their return they needed to prove their right to land by convincing Customs officials of their identity and of the validity of their claim to Australian domicile. Neither long-time residents, naturalized British subjects, nor the Australian-born could take for granted their right to re-entry if they looked “Asiatic” or “coloured.” They needed a piece of paper to prove it (see figure 1.1).

Some relied on naturalization papers or Australian birth certificates as proof, but most traveled after having applied and paid for an official certificate that would exempt them from the operations of the Dictation Test. The form of these certificates changed over time. In the earliest years of the Act, nonwhite residents could be granted a Certificate of Domicile. In 1905 this was replaced by the Certificate Exempting from Dictation Test (CEDT). Starting in 1903 these documents included photographs and handprints (later thumbprints), as well as a physical description, biographical information, and travel details.¹²

Nonwhite residents had to obtain a new CEDT for every journey. Two copies of the certificates were made—the traveler carried one, while a duplicate was retained by the Customs Department at the port of departure. On return the two copies were compared, the identity of the bearer was scrutinized, and officials decided if the traveler could stay or if the traveler would be deported as a “prohibited immigrant.” Many thousands of these certificates have been preserved. A growing number have been digitized and are available online. With portrait photographs and inky black handprints, these certificates are visually compelling documents.

Form No. 21. **COMMONWEALTH OF AUSTRALIA.** No. *8/301*
Immigration Restriction Acts 1901-1905 and Regulations.

CERTIFICATE EXEMPTING FROM DICTATION TEST.

I, *John Baxter* the Collector of Customs
for the State of *New South Wales* in the said Commonwealth,
hereby certify that *Yee Wing*
hereinafter described, who is leaving the Commonwealth temporarily, will be excepted from
the provisions of paragraph (a) of Section 3 of the Act if he returns to the Com-
monwealth within a period of *Six months* from this date.

Date *May 1908* *John Baxter* Collector of Customs.

DESCRIPTION.

Nationality *Chinese* Birthplace *Canton*
Age *46 years* Complexion *Dark*
Height *Medium* Hair *Dark*
Build *Medium* Eyes *Brown*

Particular marks, _____

(For impression of hand see back of this document.)

By Authority: ROBT. S. INNES, Government Printer, Melbourne.

Yee Wing

Date of departure *May 1908* Destination *Fiji*
Ship *Sura*
Date of return *June 30. 08* Ship *Sura*
Port *Sydney* *Wm. Cobb*
Customs Officer

Figure 1.1. Charles Yee Wing's Certificate Exempting from Dictation Test from 1908, when he traveled from Sydney to Fiji. NAA: ST84/1, 1908/301-310.

Records

The bureaucratic record-keeping system that underpinned the Immigration Restriction Act is preserved within the National Archives of Australia. As well as the exemption certificates, there are policy documents, departmental correspondence, case files, naturalization and birth records, reference letters, application forms, police reports, registers, indexes, and more. It is somewhat ironic that the records left by the bureaucracy of White Australia, an ideology that sought to marginalize and even deny the existence of “nonwhite” Australians, in fact document their lives in considerable detail and provide tangible evidence of Australia’s multiracial past.

Around the world there is a growing number of examples where records of control, surveillance, or oppression are being used to recover information about marginalized individuals or groups. “Records,” argues archivist Eric Ketelaar, “may be instruments of power, but, paradoxically, the same records can also become instruments of empowerment and liberation, salvation and freedom.”¹³ Ketelaar points to the use of Nazi documents in delivering compensation for assets seized during the Holocaust. In Australia, official records have been important in revealing the shocking history of forced removal of Aboriginal children from their families from the nineteenth century to as recently as the 1960s—they are known as the “Stolen Generations,” so named after groundbreaking work by historian Peter Read.¹⁴

Beyond supporting claims for social justice, such records can be embraced as sources of family or community heritage. For instance, historian Ricardo Punzalan describes how records of a US-administered leprosarium in the Philippines have been reclaimed as a symbol of community pride.¹⁵ Mark Aarons has written a history of his politically engaged family through detailed surveillance files accumulated by the Australian Security Intelligence Organisation.¹⁶ And records created under the White Australia policy, as well as similar systems in Canada, the United States, and New Zealand, are actively being used by family and community historians of Chinese, Japanese, Afghan, Indian, and Syrian descent to explore this aspect of their heritage for the first time. But these records are often preserved as evidence of systems rather than people. As archival theorists such as Terry Cook, Verne Harris, and Wendy Duff have argued, archival description is itself full of politics. Archivists “cannot describe records in an unbiased, neutral, or objective way,” note Duff and Harris. “Descriptions

inevitably privilege some views and diminish others.”¹⁷ What if, however, technology could open descriptive systems to new perspectives?

Historian Tim Hitchcock has written about how digitization and keyword searching has “freed us from the habit of mind implied by the structure of the archives.” We can see people as well as institutions, lives as well as bureaucratic processes. “What changes,” Hitchcock asks, “when we examine the world through the collected fragments of knowledge that we can recover about a single person, reorganized as a biographical narrative, rather than as part of an archival system?”¹⁸

Hacking the Archives

As historians, we have experienced many moments of excitement and inspiration in the collections of the National Archives of Australia. We are deeply in love with the records and the stories they reveal. We cannot say the same about the National Archives’ collection database, RecordSearch. Among its frustrations, RecordSearch’s authentication system makes sharing and citing links difficult. Until recently, its digitized file viewer lacked basic functionality and important contextual information. Despite some improvements over the years, it is a system that reflects the management practices of archives rather than the access needs of researchers. This, of course, is not unusual in the world of cultural heritage collections.

While working for the National Archives of Australia in 2007–08, we realized that it was possible to hack around some of these limitations. Creating a Zotero translator to extract structured data from RecordSearch pages revealed the power of screen-scraping—we did not have to live only with what was rendered in the browser. Around the same time we were involved in a project, Mapping Our Anzacs, to create a map-based interface to 375,000 World War I service records.¹⁹ This involved manipulating existing descriptive data to create new modes of access. Learning Python with the help of The Programming Historian finally pushed us over the edge and we created the first in a series of Python-based screen-scrappers to harvest data directly from RecordSearch. We were hooked.²⁰

Throughout this journey of exploration and enlightenment, the records of the White Australia policy remained close to our hearts. One of our hacks was a userscript that upgraded RecordSearch’s digitized file viewer. Userscripts are Javascript programs that run in the browser to rewrite the form and functionality of selected web pages. Our script added options

for navigation and printing, but it also made use of the Cooliris browser plugin to display images on a floating 3D wall. It offered a completely different way of seeing the archives.²¹ We excitedly pointed our new viewer at digitized files full of CEDTs. The wall of documents we saw, of faces and handprints of men, women, and children marginalized by White Australia, affirmed our belief that these records were not only historically significant, they were visually compelling. They had to be seen.

Our plotting continued after we left our jobs at the National Archives, but without an institutional home it seemed impossible to create and sustain a new research project. Until we just *did* it. Inspired by Hacking the Academy, we wondered what would happen if we just started talking about what we wanted to do—without funding, without research partners, without timelines.²² And so we launched Invisible Australians: Living Under the White Australia Policy, a web-based project designed to pull together the biographical threads embedded in the archives.²³ Our aim was to be “modular and opportunistic”—to be able to grow when resources allowed, to bolt on related projects, to absorb existing tools and technologies.²⁴ The history locked in these records was too important not to try.

How We Found the Faces

In the meantime other coders and hackers we knew were doing interesting things with cultural collections. Mitchell Whitelaw, for instance, started exploring the Visible Archive, developing techniques to see beyond a single file or document to the complete holdings of the National Archives of Australia.²⁵ Paul Hagon wondered whether facial detection technology could be used as a means of discovery within the photographic collections of the National Library of Australia.²⁶ Whitelaw’s challenge to show it all and Hagon’s idea to create new access points by extracting features from images inspired us to reconsider how we might see the records of the White Australia policy. Instead of a wall of documents, what could we learn from a wall of faces?

It would be nice to portray our process in hindsight as something careful and rigorous. But it was much more a case of playing with possibilities. Originally we had imagined that identifying and extracting photographs from CEDTs would be a semimanual, crowdsourced process, with volunteers marking up the coordinates of each individual photograph. But one weekend we just googled “facial detection python” and found a Python

script that used OpenCV to identify faces within images. OpenCV, we learned, was the go-to package for computer vision hackers.

People are really good at recognizing the characteristics of a face. We do it without thinking. Indeed, we are so good at it that we tend to “see” faces in all sorts of odd contexts—a phenomenon known as *pareidolia*. Computers have no such ability. They can be trained to detect a face, but generally this involves breaking the task down into many small, simple calculations.²⁷ Training computers to identify objects in images can be a complex and time-consuming business, but fortunately OpenCV ships with a number of pretrained “classifiers” that enable you to detect faces, eyes, and even cats.

With all the hard work of training done for us, finding faces was simply a matter of opening images using Python and feeding through them through the OpenCV classifier. We pointed the script we found at some CEDTs and, after a bit of tweaking, it worked! OpenCV fed us back the coordinates of any faces it found, and with some basic image manipulation in Python we could crop those areas and save them as new files. It was surprisingly easy to extract portrait photographs from archival documents.

We knew little about the technology of facial detection when we ran our first experiments. However, once we saw that it worked we started to think about what came next. Could we apply this extraction technique to the many thousands of documents held by the National Archives? First we needed easy access to all those image files. It is tempting to skim over the process of assembling our collection of images—downloading files is not as exciting as extracting faces—but, in reality, we have spent much more time wrestling with the frustrations of RecordSearch than with OpenCV. Cultural heritage institutions are starting to make metadata and images available in forms that encourage digital research. But research, by its nature, tests the boundaries of meaning, evidence, and access—we cannot be satisfied with what we are given. In the case of RecordSearch we had no option but to extract what we needed from the web interface through the process of screen-scraping. Reverse engineering an ASP.NET website with session-based authentication and a seemingly endless maze of redirects is challenging. Screen-scraping RecordSearch was only possible using libraries such as Mechanize and RoboBrowser, which mimic the behavior of web browsers.²⁸

Fortunately, when we began our facial detection experiments we had already built a working screen-scraping. It has since undergone several bouts of breakage and revision as changes to RecordSearch wrought new mys-

teries and complications. But at the time we simply pointed our screen-scraper at the archival series ST84/1, which we knew contained a large number of CEDTs, and waited. Eventually we had a collection of 12,502 images and we could start looking for faces.

Although looping through the images and applying the facial detection script was straightforward, there was much trial and error as we sought to improve the script's accuracy while minimizing the number of false positives.²⁹ Lacking a detailed understanding of how the facial detection algorithm actually worked, we simply plugged a variety of values into the "scale factor" and "minimum neighbors" parameters of the object detection module and observed the results. Eventually we settled on an appropriate balance and weeded out a few more false positives by applying an extra check to each cropped face. But this work forced us to ask, "What is a face?"

The facial detection algorithm simply returned a list of coordinates—a box for every face it thought it had found in the image, tightly focused on eyes, nose, and mouth. They were portraits reframed according to an algorithm's own assumptions of significance—*these* are the features that define a face. As we viewed the initial output of our script we made the small but important decision to expand the boxes. Adding an extra fifty pixels to each side kept the focus, but revealed more of the person. It seemed to make a difference.

With configuration complete we unleashed the script on all 12,502 images and watched with alarm as the CPU temperature of our three-year-old laptop soared. It is a computationally intensive process but possible even with modest technological resources. The laptop survived and, after several hours, we had a folder containing 11,170 cropped images. Despite our best efforts, many false positives remained. We simply weeded these out manually, leaving us with 7,247 faces (see figure 1.2).

The resulting images offered a powerful commentary about White Australia, and we wanted to display them in a way that was both simple and direct. It was the faces that mattered. Using the web application framework Django to manage the metadata and deliver the content, we created an interface using the Javascript libraries Isotope and Infinite Scroll. Although the project built on our history of RecordSearch hacking, it was a quick experiment that took little more than a weekend to harvest, process, and build. The result was a wall of faces—continuous, compelling, and unsettling. The more you scroll, the more faces appear. Faces of the people who destabilized Australia's claim of being a white nation—thousands of men, and a smattering of women and children, of Chinese, Japanese, Indian, and

Figure 1.2. The Real Face of White Australia, <http://invisibleaustralians.org/faces/>

Syrian heritage, to name a few. Faces of the people who lived their lives within a system of surveillance and control. Faces of the people who built homes, families, and businesses in a country that sought to deny their existence. This was the Real Face of White Australia.

Seeing

In 2012 we received an email from Mayu Kanamori, an artist researching the life of an early Japanese Australian photographer, Yasukichi Murakami, who had arrived and settled in Western Australia in 1897. Kanamori had come across the Invisible Australians site in her research and felt moved to thank us for what we were doing, writing: “When I scrolled down the Faces section of your website, browsing through the faces, tears welled up, and I couldn’t stop crying as if some sort of flood gates had been removed.”³⁰

We knew that the records, the photographs, the handprints, all carried emotive weight—it was the very reason we sought to expose them. What we did not quite realize was the effect of scale. Bringing all those photos together, without interpretation or intermediation, created a different type of experience. As Peter Binkley commented: “[The Real Face of White Australia] zooms you from the macro level of political criticism of the racist

policy down to the micro level of individual stories, and back again through the sheer accumulation of cases.”³¹

Our simple wall of faces showed that this was not just an archive, not just a policy.

The level of interest in the project from the international digital humanities (DH) community was also unexpected. In a September 2011 blog post we described our experiment, the technical details, and the context of the records.³² A few months later Tim explored the broader significance of the project in a presentation at the National Digital Forum (NDF) in New Zealand. One of our posts was picked up and reposted in South Africa.³³ Tim’s NDF talk was then published in the inaugural edition of the *Journal of Digital Humanities*.³⁴ Before long our weekend experiment was studied in digital history courses, discussed at museum conferences, and cited in research on a variety of topics including visualization, serendipity, and race.³⁵

Perhaps most surprising to us was the way the Real Face of White Australia was seen to illustrate key aspects of digital humanities practice. James Smithies described it as “one of the signal DH publications of 2011,” while Ted Underwood pointed to Tim’s NDF presentation as a “pep talk” for those uncertain where to start in the field.³⁶ Our experiment with computer vision offered an example of DH’s hacker ethos—of what becomes possible when you dig into code. At the same time it was also cited as a DH project that critically engaged with questions of race and power.³⁷

This was an experiment without an institutional home, built over a weekend on an aging laptop in our study. Building on our experience of getting data out of RecordSearch, and taking advantage of sophisticated open-source libraries such as OpenCV, we were able to create a new way of seeing and using the records. But the wall of faces was more than just an interface. The responses it garnered seemed to justify our decision to launch Invisible Australians as a research project without structure or support. The wall was a piece of opportunistic hacking that transformed our promises into something more tangible—it communicated our intentions more effectively than any manifesto or research plan.

Despite the project’s overall success, there were difficulties beyond the technical and logistical challenges. Concerns about representation and responsibility arose numerous times as we grappled with the technology and the records, and criticisms of the project have tended to focus on questions of context and selection. One post thought it was “ethically dubious” to present the photos without consent, separated from the original documents.³⁸ Others misunderstood the process and thought we were identify-

ing people *by* race. But rather than separating people by race, we sought to reveal the way in which the bureaucracy was creating the categories of “white” and “nonwhite” through the operations of the White Australia policy. As Kate’s work on Anglo-Chinese Australians has shown, officials administering the Immigration Restriction Act focused on identities such as “Chinese,” “Asiatic,” “half-caste,” or “colored” over “Australian,” “British,” “European,” or “white.”³⁹ The individuals whose photographs appeared on our wall were ordinary people living ordinary lives, subjected to a system of discrimination and control primarily because of their appearance—their images appeared on the wall because of the racist machinery of the Australian government.

It could be easy to think of Invisible Australians as some sort of rescue mission, liberating people from the archives of oppression. But archivists Wendy Duff and Verne Harris have noted dangers in taking it upon ourselves to restore missing voices to the historical record. “How can we avoid the danger of speaking for these voices?” they ask. “How can we avoid reinforcing marginalization by naming ‘the marginalized’ as marginal?”⁴⁰ Part of the task of our wider project in Invisible Australians is to provide space for people to be people. To have lives that surprise and confound us. To act in ways that challenge our categories. To resist us, to refuse to be aggregated, tallied, or visualized.

As we embarked on creating the wall of faces, historian Sophie Couchman’s work on early official identification photography and the Chinese in Australia was present in our minds. Couchman has written about a popular traveling exhibition, *Forgotten Faces: Chinese and the Law*, curated by the Public Record Office of Victoria and the Golden Dragon Museum in Bendigo in 2005. The exhibition presented large reproductions of gaol photographs of Chinese men imprisoned in Victoria between the 1870s and 1900, accompanied by brief biographical sketches drawn mostly from court and prison records. Couchman was critical of the exhibition for “deliberately pulling photographs of Chinese prisoners from the wider prison archive,” thereby presenting the Chinese in Victoria as criminals and powerless victims of government bureaucracy.⁴¹ She further noted that this process of selection obscured the fact that Victoria’s system of gaol photography treated Chinese criminals in the same way as white criminals.

As our script cleverly selected and cropped out face after face from the CEDTs, we thought about whether the same sorts of criticisms could be leveled at what we were doing. Was the Real Face of White Australia just another type of rogues’ gallery? Were we representing our subjects as more

than just passive victims of a racist bureaucracy? Were we using their images respectfully and decently? Could the images be understood by a contemporary audience? How could the resistance and agency of people like Charles Yee Wing be acknowledged?

Other work by Sophie Couchman looks closely at a series of photographs of Chinese entering Victoria that were used for immigration control purposes—one small part of the massive archive of the White Australia policy, like the CEDTs. In her reading of the 269 photographs, which date from 1899 to 1901, Couchman noted that these were “not so mug mug-shots” in that the Chinese subjects had a deal of autonomy in the way they represented themselves—in their choice of clothing and accessories such as umbrellas and hats (and even a bicycle), and in their poses and facial expressions.⁴² In light of this work, we realized that our wall of faces needed to be able to reflect the idiosyncrasies of the photographs, to acknowledge the self-representation within them (particularly early ones used before the administrative processes became more standardized), and to avoid assembling a gallery of mug shots. Therefore, we decided to leave the images at the different sizes they were, rather than resizing them for consistency. This, together with widening out the crop, allowed more of the person’s clothing, hairstyle, and background to be seen.

While the Real Face of White Australia is far from perfect, finding ways of representing agency has been important, particularly as massed groupings of portrait photographs are often associated with memorials as well. A “wall of faces” in the 9/11 Memorial Museum in New York displays photos of the nearly 3,000 people who died in the attacks “to try and communicate the scale of human loss.”⁴³ The United States Vietnam Veterans Memorial Fund has created an online “wall of faces” linked to profiles of individual service people killed in the war.⁴⁴ And the walls of one building at the Tuol Sleng Museum of Genocide in Cambodia are covered with photographs of victims, as is the ceiling of the Hall of Names in the Yad Veshem Holocaust History Museum in Jerusalem.⁴⁵ Yet the Real Face of White Australia is not a memorial. The people in the photos suffered oppression under the White Australia policy, but casting them as victims ignores their efforts to negotiate the system, to fight against its restrictions, to simply live their lives. This is a challenge we continue to grapple with, but perhaps part of the answer lies within the photos themselves.

Jenny Edkins suggests that despite conscious attempts to read meanings into portrait photographs, there are other, more visceral responses: “We are not merely passive spectators, but intimately involved, not separate be-

ings, but inevitably interconnected.”⁴⁶ A face in a photograph, she suggests, can reach us in ways that challenge systems of authority and power that bear a sense of connection and obligation. Perhaps the faces on our wall can speak for themselves?

Access Against the Grain

In our initial blog posts about the Real Face of White Australia, we described it as a finding aid. Despite some people’s concerns about context, all the photos are linked both to an uncropped image of the full exemption certificate and to further file details in RecordSearch, allowing users to navigate records in a different way. As Barbara Fister noted in a post about Tim’s NDF presentation:

In a sense he’s reverse-engineering the bureaucracy that once determined who was a proper Australian and is using the record-keeping used to control and oppress people to restore their history. He’s also taking what cultural institutions do—preserve, sort, interpret, and present culture—and reorganizing it using different rules.⁴⁷

Archival descriptive systems tend to be arranged in a hierarchy—from collections to parts. While keyword searching allows discovery across the hierarchy, items remain fixed in a matrix of significance, context, and containment. Mitchell Whitelaw’s Series Browser, for example, brings the properties of the containers to the surface, allowing users to see relationships across the whole collection.⁴⁸ Technologies that detect features in images or text, that aggregate and analyze existing metadata, allow us to turn descriptive hierarchies inside out. Within the National Archives our faces were locked away in photographs, themselves parts of larger digital images representing documents, contained in files, and organized in series. The Real Face of White Australia brought these buried features to the surface while retaining their archival context.

We could have manually cropped images from an assortment of files to create an exhibition of faces, but machine processing added the power of scale and the possibility of serendipity. As reactions to the wall have highlighted, the sheer number of faces, arranged in a seemingly endless array, carried both political critique and emotional engagement. Even Kate, who knows the records well, could observe new things through the machine’s computational gaze and contemplate new research methodologies

for questions we did not know to ask before. The very lack of curation of the original documents prompted new questions. For example, some viewers wondered about the preponderance of men and the absence of women. Where are the women? This is just one of many new questions revealed by the project.⁴⁹

In 2014 Tim updated our scripts to use the latest version of OpenCV and applied them to a very different set of images—photographs from Trove’s massive collection of digitized Australian newspapers.⁵⁰ The quality of these images is often poor and many contained no people, but the object detection module again worked its magic. This time the script looked not only for faces, but for eyes within those faces using another pretrained classifier. From a sample set of 12,000 photos Tim extracted around 800 faces and 1,000 eyes.

The interface for Eyes on the Past, built using the Python microframework Flask and MongoDB, presents a random selection of eyes, slowly blinking on and off. Clicking on an eye reveals the full face and the source of the image. Clicking again on the caption opens the full newspaper article in Trove. Where the Real Face of White Australia overwhelms with scale and meaning, Eyes on the Past is minimal and mysterious. It emphasizes absence, and the fragility of our connection with the past, even while it provides a new way of exploring the digitized newspapers. Some have found it beautiful; others just think it is creepy.

There is something glorious and exciting in the fact that the same technology can result in such different resources. Object detection cracks open images, treating them as assemblages to be queried and manipulated. New questions emerge and new experiences are possible. But these very technologies are also deeply embedded in modern systems of surveillance. While we explore the creative possibilities of facial detection, we should not ignore the historical threads that connect our own tools to the workings of discriminatory regimes like the White Australia policy.⁵¹ The ability to identify, to label, and to separate offers power to those who would control us. Under the computer’s gaze we can, like Charles Yee Wing, all be treated just like criminals.

Charles Yee Wing’s trip to China in 1911 was neither his first nor last dealing with the bureaucracy of White Australia. With a transnational business empire and political interests that stretched across Australia, New Zealand, Fiji, the United States, Hong Kong, and China, Yee Wing made a dozen or so journeys from Australia between the 1890s and 1920s.⁵² The paperwork kept by the Australian government on Yee Wing’s many over-

seas trips documents the tightening laws, policies, and administrative procedures over this period, but it also shows how he and his family pushed back against the system as regulations grew stricter.

Yee Wing was often accompanied on his travels by his white Australian wife, Susan Beck, or some of their many children. Eldest child Mabel first went overseas with her parents in 1903 at age three, and twenty-five years later displayed the same indignant spirit as her father when officials questioned her identity on returning home from New Zealand. Because she looked “Chinese,” Mabel was not permitted to leave the ship with the other passengers and was detained for some time while a Customs inspector assessed her right to enter Australia. In fact Mabel held an Australian passport, and had done so since a trip to China almost a decade earlier, but she did not initially present it as identification. As her solicitor wrote in a letter of complaint to the Collector of Customs, “she made the usual Declarations as to her place of birth, from which it should clearly have appeared that she was a natural born British Subject and was entitled to enter Australia without question.” Instead she was submitted to “indignities” and “humiliations.”⁵³

While we cannot deny the politics of the technologies we use, like Charles and Mabel Yee Wing we can find opportunities for resistance, subversion, and play. The Real Face of White Australia displays photos extracted from the existing record-keeping system, but what if we turned this around? On a whim we created another RecordSearch hack—a userscript that queries our database of faces and inserts them back into RecordSearch results.⁵⁴ The faces appear just alongside the archival metadata as if they’re bubbling up from records below. It is a hack that offers no improvements to the functionality of RecordSearch, but by seeing the faces of those who confronted discrimination, it adds a level of understanding because it can make us *feel* differently. Maybe this is what happens when instead of just files and documents, we can see the people inside.

Notes

1. “An Indignity,” *Sydney Morning Herald*, October 26, 1911, 9, <http://nla.gov.au/nla.news-article15283925>, accessed May 25, 2015.
2. See Jens Jäger, “Photography: A Means of Surveillance? Judicial Photography, 1850 to 1900,” *Crime, History & Societies* 5, no. 1 (2001): 27–51.
3. “Index to the Gaol Photographs,” State Records NSW, <http://www.records.nsw.gov.au/state-archives/indexes-online/gaol-records>, accessed June 9, 2015.
4. Passports were first introduced in Australia through the War Precautions

Act 1914–15, and made permanent in 1920 with the introduction of the Passports Act. See Jane Doulman and David Lee, *Every Assistance and Protection: A History of the Australian Passport* (Leichhardt, NSW: Federation Press, 2008).

5. On the history of the White Australia policy, see Myra Willard, *History of the White Australia Policy to 1920* (Melbourne: Melbourne University Press, 1923); A. T. Yarwood, *Asian Migration to Australia: The Background to Exclusion, 1896–1923* (Melbourne: Melbourne University Press, 1964); Anthony Palfreman, *The Administration of the White Australia Policy* (Melbourne: Melbourne University Press, 1967); Gwenda Tavan, *The Long Slow Death of White Australia* (Melbourne: Scribe, 2005).

6. “The ‘White Australia’ Programme,” *Gympie Times and Mary River Mining Gazette*, September 21, 1901, 3, <http://nla.gov.au/nla.news-article171004149>, accessed June 9, 2015.

7. Paul Jones, “Alien Acts: The White Australia Policy, 1901 to 1939” (PhD diss., University of Melbourne, 1998), 55.

8. National Archives of Australia (NAA): A4144, 860/1943.

9. NAA: A1, 1917/16266.

10. Barry York, *Immigration Restriction 1901–1957: Annual Returns as Required under the Australian Immigration Act between 1901 and 1957 on Persons Refused Admission, Persons Who Passed the Dictation Test and Departures of Coloured Persons from Australia* (Canberra: Centre for Immigration & Multicultural Studies, Australian National University, 1994).

11. “Part VIII: Non-European Races,” *Census of the Commonwealth of Australia, 1911* (Melbourne: Commonwealth of Australia, 1911), <http://www.abs.gov.au/AUSSTATS/abs@.nsf/mf/2112.0>, accessed June 9, 2015.

12. Originally the Certificate of Domicile form was printed without space for photographs or a handprint. From mid-1903 photographs and handprints were placed on the blank backs of the certificates until the form was redesigned in early 1904 with space for photographs on the front and a handprint on the back. See NAA: ST84/1, 1903/161–170 and ST84/1, 1904/71–80 for examples from New South Wales.

13. Eric Ketelaar, “Archival Temples, Archival Prisons: Modes of Power and Protection,” *Archival Science* 2 (September 2002): 221–38, doi:10.1007/BF02435623.

14. See Peter Read, *A Rape of the Soul So Profound* (St Leonards, NSW: Allen & Unwin, 1999).

15. Ricardo L. Punzalan, “Archives and the Segregation of People with Leprosy in the Philippines” (ICHORA-4, University of Western Australia, 2008), 93–102.

16. Mark Aarons, *The Family File* (Melbourne: Black Inc., 2010).

17. Wendy M. Duff and Verne Harris, “Stories and Names: Archival Description as Narrating Records and Constructing Meanings,” *Archival Science* 2 (September 2002): 263–85, doi:10.1007/BF02435625.

18. Tim Hitchcock, “Digital Searching and the Re-Formulation of Historical Knowledge,” in *The Virtual Representation of the Past*, eds. Mark Greengrass and Lorna Hughes (Farnham, UK: Ashgate, 2008), 81–90.

19. Tim Sherratt, “Bringing Life to Records: ‘Mapping Our Anzacs’ at the National Archives of Australia,” in *A Different Kind of Web: New Connections between Ar-*

chives and Our Users with Web 2.0, ed. Kate Theimer (Chicago: Society of American Archivists, 2011), 128–38.

20. William J. Turkel et al., eds., *The Programming Historian*, 2nd ed. 2016. <http://programminghistorian.org/>

21. Tim Sherratt, “Archives in 3D,” *Discontents*, December 17, 2008, <http://discontents.com.au/archives-in-3d/>

22. Dan Cohen and Tom Scheinfeldt, eds., *Hacking the Academy: The Edited Volume*, <http://hackingtheacademy.org/>

23. <http://invisibleaustralians.org/>

24. Tim Sherratt, “Hacking a Research Project,” *Discontents*, June 17, 2010, <http://discontents.com.au/hacking-a-research-project/>

25. Mitchell Whitelaw, “Visualising Archival Collections: The Visible Archive Project,” *Archives and Manuscripts* 2 (2009): 22–40.

26. Paul Hagon, “Everything I Learned about Cataloguing I Learned from Watching James Bond” (paper presented at VALA2010 Connections Content Conversations, Melbourne, 2010), <http://www.vala.org.au/vala2010-proceedings/vala2010-session-7-hagon>

27. Adam Harvey, ‘OpenCV Face Detection: Visualized’, Vimeo, 2010, <https://vimeo.com/12774628>, accessed September 26, 2017.

28. Tim Sherratt, “RecordSearch-Tools,” *GitHub*, <https://github.com/wragge/recordsearch-tools>, accessed June 9, 2015.

29. Tim Sherratt, “Facial-Detection,” *GitHub*, <https://github.com/wragge/Facial-detection>, accessed June 9, 2015.

30. Mayu Kanamori, email message to authors, March 28, 2012. Kanamori’s research is documented in her project blog (<http://aboutmurakami.wordpress.com>) and culminated in a performance called *Yasukichi Murakami: Through a Distant Lens* at the Griffin Theatre in Sydney in early 2015.

31. Peter Binkley, “A Walk with Love and Data,” *Quaedam Cuiusdam*, October 26, 2011, http://www.wallandbinkley.com/quaedam/2011/10_26_a-walk-with-love-and-data.html, accessed May 25, 2015.

32. Tim Sherratt, “The Real Face of White Australia,” *Discontents*, September 21, 2011, <http://discontents.com.au/the-real-face-of-white-australia>, accessed May 25, 2015; Kate Bagnall, “The Real Face of White Australia,” *The Tiger’s Mouth*, September 22, 2011, <http://chineseaustralia.org/the-real-face-of-white-australia>, accessed May 25, 2015.

33. “News—The Real Face of White Australia,” *The Archival Platform*, October 6, 2011, http://www.archivalplatform.org/news/entry/the_real_face_of_white_australia.

34. Tim Sherratt, “It’s All About the Stuff: Collections, Interfaces, Power, and People,” *Journal of Digital Humanities* 1, no. 1 (Winter 2011), <http://journalofdigitalhumanities.org/1-1/its-all-about-the-stuff-by-tim-sherratt>, accessed May 25, 2015.

35. For example: Fiona Romeo and Lawrence Chiles, “The Future of Digital Interpretation: Gallery Objects as Service Avatars” (Museums and the Web 2012, San Diego, 2012), <http://www.museumsandtheweb.com/mw2012/papers/>

the_future_of_digital_interpretation_gallery_o.html; Chris Alen Sula, "Quantifying Culture: Four Types of Value in Visualisation," in *Electronic Visualisation in Arts and Culture*, eds. Jonathan P. Bowen, Suzanne Keene, and Kia Ng, Springer Series on Cultural Computing (London: Springer, 2013), 25–37, http://link.springer.com/chapter/10.1007/978-1-4471-5406-8_3; Kaleigh Bradley, "Exhibiting Race: The Power of Portraiture," *ActiveHistory.ca*, July 19, 2012, <http://activehistory.ca/2012/07/exhibiting-race-the-power-of-portraiture/>, accessed May 25, 2015.

36. James Smithies, "Digital Humanities, Postfoundationalism, Postindustrial Culture," *Digital Humanities Quarterly* 8, no. 1 (2014), <http://www.digitalhumanities.org/dhq/vol/8/1/000172/000172.html>, accessed June 9, 2015; Ted Underwood, "What No One Tells You about the Digital Humanities," *The Stone and the Shell*, December 15, 2011, <http://tedunderwood.com/2011/12/15/what-no-one-tells-you-about-the-digital-humanities/>, accessed June 9, 2015.

37. For example: Todd Presner, "Critical Theory and the Mangle of Digital Humanities," in *Between Humanities and the Digital*, eds. Patrik Svensson and David Theo Goldberg (Cambridge, MA: MIT Press, 2015), 55–68; Anne Cong-Huyen, "#CESA2013: Race in DH—Transformative Asian/American Digital Humanities," *Anne Cong-Huyen*, September 24, 2013, <https://anitaconchita.wordpress.com/2013/09/24/cesa2013-race-in-dh-transformative-asianamerican-digital-humanities/>, accessed June 9, 2015.

38. yiduique [Stephanie Lai], untitled post, *A Penguin of Very Little Brain*, accessed May 25, 2015, <http://yiduique.tumblr.com/post/23480823112/materialworld-via-the-tigers-mouth-the-real>.

39. Kate Bagnall, "Anglo-Chinese and the Politics of Overseas Travel from New South Wales, 1898 to 1925," in *Chinese Australians: Politics, Engagement and Resistance*, eds. Sophie Couchman and Kate Bagnall (Leiden: Brill, 2015).

40. Wendy M. Duff and Verne Harris, "Stories and Names: Archival Description as Narrating Records and Constructing Meanings," *Archival Science* 2 (September 2002): 263–85, doi:10.1007/BF02435625.

41. Sophie Couchman, "In and Out of Focus: Chinese and Photography in Australia, 1870s–1940s" (PhD diss., La Trobe University, 2009), 122.

42. NAA: B6443. Sophie Couchman, "Not So Mug Mugshots: Behind the Portraits of Series B6443," *Crossings* 9, no. 3 (2004), archived at http://pandora.nla.gov.au/pan/13231/20050203-0000/asc.uq.edu.au/crossings/9_3/couchman.html

43. "Wall of Faces | National September 11 Memorial & Museum," accessed March 21, 2015, <http://www.911memorial.org/blog/tags/wall-faces>

44. "Virtual Vietnam Veterans Wall of Faces | The Vietnam Veterans Memorial Fund," accessed March 21, 2015, <http://www.vvmf.org/Wall-of-Faces/>

45. "S21 Victims—The Killing Fields Museum of Cambodia," accessed June 9, 2015, <http://www.killingfieldsmuseum.com/s21-victims.html>; "Hall of Names—The Holocaust History Museum—Yad Vashem," accessed March 21, 2015, http://www.yadvashem.org/yv/en/museum/hall_of_names.asp

46. Jenny Edkins, *Face Politics* (London: Routledge, 2015), 44.

47. Barbara Fister, "'Ambitious, Unfunded, and Possible,'" *Inside Higher Ed*, December 1, 2011, <http://www.insidehighered.com/blogs/library-babel-fish/%E2%80%99C%20ambitious-unfunded-and-possible%E2%80%9D>

48. Mitchell Whitelaw, "Series Browser Screencast," *The Visible Archive*, September 2009, <http://visiblearchive.blogspot.com.au/2009/09/series-browser-screencast.html>

49. Kate Bagnall, "Where Are the Women?" *The Tiger's Mouth*, September 28, 2012, <http://chineseaustralia.org/where-are-the-women/>, accessed June 9, 2015.

50. Tim Sherratt, "Eyes on the Past," *Discontents*, May 24, 2014, <http://discontents.com.au/eyes-on-the-past/>, accessed June 9, 2015.

51. Tim Sherratt, "The Perfect Face," presented at the National Digital Forum New Zealand, October 13, 2015, <https://doi.org/10.6084/m9.figshare.1577665.v1>, accessed September 23, 2017.

52. See NAA: SP42/1, C1914/7447; SP42/1, C1919/9167; ST84/1, 1903/151-160; ST84/1, 1905/251-260; ST84/1, 1908/301-310; ST84/1, 1910/44/11-20; ST84/1, 1911/73/11-20; ST84/1, 1913/138/1-10; ST84/1, 1915/164/81-90.

53. NAA: SP42/1, C1928/5260.

54. Tim Sherratt, "The People Inside," *Discontents*, July 12, 2012, <http://discontents.com.au/the-people-inside/>

CHAPTER 2

Bringing Trouvé to Light

Speculative Computer Vision and Media History

Jentery Sayers

Computer vision has a long social and technical history, dating back to at least the mid-1900s. This chapter surveys that history across disciplines and then provides a case study for “speculative computer vision,” or the use of computer vision to make objects that do not or no longer exist. The case study is grounded in electro-mobile jewelry designed by Gustave Trouvé, a nineteenth-century engineer famous for his work in miniaturization. Although some instances of Trouvé’s jewelry remain in museums, none are functional. This chapter demonstrates how speculative computer vision may be integrated into historical research to model, fabricate, and test working prototypes of Trouvé’s “lost” devices in the present moment. Such integration raises questions about early, hyperbolic representations of electro-mobile jewelry and prompts inquiry into how (and even whether) it worked.

Computer vision programmatically describes patterns in data.¹ Though difficult for people to identify, let alone isolate and examine at scale, such patterns are quite common: the recurrence of a face over time, the repetition of an object in an archive, the persistence of a style across a corpus. Computer vision also automates tasks while expanding the role of algorithms in everyday life. A now-common feature in popular software packages, it sees, hears, finds, and expresses things for people. A. Aneesh links this expansion to “a new kind of power” and governance, or “*algocracy*—rule of the algorithm, or rule of the code.”² Here, programmatic description is embedded

so significantly in infrastructures that algorithms tacitly shape behaviors and prosaically assert authority alongside existing bureaucratic measures. Consequently, routine decisions are delegated (knowingly or not) to computational procedures that—echoing the work of Bruno Latour and many others in science and technology studies—are black-boxed, running in the background as protocols, default settings, or user-friendly interfaces.³

Consider the banal and ubiquitous integration of CAPTCHAs (Completely Automated Public Turing test to tell Computers and Humans Apart) into online submission forms. By prompting you to type obscure letters into a box, CAPTCHAs ask, “Are you human?” In doing so, they regulate information flow and implicitly assume quite a bit about people (e.g., that a “human” reader is not blind, does not have low vision, or is not dyslexic). In many cases, such as Google’s reCAPTCHA initiative, CAPTCHAs also assist computing projects in their digitization, annotation, and machine-learning efforts. When computer vision cannot identify a visual pattern, CAPTCHAs forward it to people for transcription into discrete alphanumeric characters. These characters are then integrated back into collections of data and used to train algorithms to interpret more content. Understood this way, a person’s response to “are you human?” also helps the machine. In economic terms, these quotidian exchanges are mutually beneficial, or at least efficient: CAPTCHA algorithms protect websites from malicious bots and optical character recognition (OCR) attacks, while people provide desired information and verification without consuming much of their own time and effort. Once habituated online, such algocratic tasks become inconsequential in the instance yet extremely meaningful—and strikingly productive of value—in the aggregate.

This chapter is not about algocracy or CAPTCHAs, but it is informed by them both. It encourages humanities scholars, especially scholars of material culture and media history, to *experiment with computer vision techniques that privilege speculation over confirmation*. My particular take on the relevance of computer vision to humanities research does not delegate authority to algorithms. It is not about efficient vision or foregrounding patterns in repositories—approaches that, to be fair, can dramatically influence the very epistemology of historical research (e.g., what it means to see the past through a computer or with an algorithm). Invested instead in the cultural, aesthetic, and embodied dimensions of science, technology, and media studies, speculative computer vision resists dichotomizing media and mediation, digital and analog materials, automated and manual labor, and human and machine phenomenologies. Rather than *capturing or re-presenting*

history, speculative computer vision helps scholars to *make objects that do not or no longer exist* and then argue through them. Put this way, it is highly influenced by the work of Tiffany Chan, Lynn Cherny, Beth Coleman, Johanna Drucker, Devon Elliott, Caitlin Fisher, Lauren Klein, Kari Kraus, Robert MacDougall, Bethany Nowviskie, Daniela Rosner, Mark Sample, and William J. Turkel, building on their research to expand humanities approaches to scanning, modeling, and fabricating media in 3D.

In this chapter, I lean toward the artistic edges of the humanities. I do not dig into the technical specificities of vision science or artificial intelligence. I also do not claim that computers should assist scholars in resolving the long-standing ambiguities of media history. Much is left muddy. Still, my claims for speculation are grounded in an example, namely Gustave Trouvé's electro-mobile jewelry from the 1860s. Later, I elaborate on Trouvé's work. For now, it is important to note that, during the nineteenth century, electro-mobile jewelry was produced in small batches and is now very difficult to study, see, or handle. These obstacles to historical research are in part what render the jewelry fitting for speculative vision. Despite their evasiveness, there is a good chance they can be remade with a pinch or two of conjecture and the perspectives of a computer.⁴ That said, by including "bringing Trouvé to light" in the title of this chapter, I do not imply an Enlightenment paradigm of rationality. I also do not encourage the algorithmic discovery of historical causes that—until now—have rested, ghost-like, beyond the threshold of perception, waiting for scholars to "reveal" them. Instead, I reconsider electro-mobile jewelry by prototyping it with extant materials. First, I survey computer vision research across the disciplines, including some of its ostensible benefits as well as its apparent biases. This survey is intended to give humanities researchers a sense of computer vision's history and various applications. Yet it is also meant to lay the groundwork for speculative computer vision as an alternative to popular methods. After the survey, I articulate an approach to speculative computer vision by way of Trouvé and some prototypes. This approach does not resolve the problems of using computer vision for confirmation. It is not a solution. It merely proposes a way to see sideways with computers—to conjecture with them.

Computer Vision Across the Disciplines

According to various accounts, computer vision research began as early as 1966, during the Summer Vision Project, when Marvin Minsky, Seymour

Papert, Gerald Jay Sussman, and others in the Artificial Intelligence Group (AIG) at the Massachusetts Institute of Technology (MIT) investigated how to use figure-ground analysis to automatically divide pictures into regions based on surface and shape properties. This region description would ultimately act as the basis for object identification, where items in pictures were recognized and named by machines with controlled vocabularies of known objects.⁵ Cameras were attached to computers to achieve this automated description and identification, with an express desire to eventually construct a pattern analysis system that would combine “maximal flexibility” with “heuristic rules.”⁶

Although computer vision has developed significantly since the 1960s and '70s, the AIG's Summer Vision Project marks a notable transition in media history, a moment when researchers began integrating image processing into the development of artificial intelligence, including the training of computers to read, detect, and describe aspects of pictures and visual environments.⁷ During the project, AIG researchers also started asking computer vision questions that, if only tacitly, have persisted: How does computer vision differ from human vision? To what degree should computer vision be modeled on human phenomenology, and to what effects? Can computer or human vision even be modeled? That is, can either even be generalized? Where and when do issues of processing and memory matter most for object recognition and description? And how should computer vision handle conflict and ambiguity?⁸ These questions are at once technical and conceptual, as are many responses to them, meaning that computer vision should not be extracted from the contexts of its intellectual *and* material development.

Today, computer vision has moved, at least in part, from laboratories into consumer technologies. One popular application is answering the question, “Is this you?” or “Is this them?” iPhoto, Facebook, and Kinect users are intimately familiar with this application, where face detection algorithms analyze patterns to calculate a core or integral image within an image, assess differences across a spectrum of light, and view images across scales.⁹ In open-source communities, many practitioners combine the Open Source Computer Vision (OpenCV) library with the Python, C++, and Java programming languages to perform this detection work. These scripts rely on frameworks to train classifiers to detect “objects” (including faces, bodies, and body parts) in images based on cascades of features. To see faces while algorithmically scanning images, OpenCV uses the widely popular Viola-Jones object detection framework, which relies

on “Haar-like” image features for cascading.¹⁰ Similar to iPhoto and other image management applications, OpenCV can be used to identify—often with errors and omissions—the same face across a distribution, even when multiple faces appear in the same image.

But to write computer vision scripts, programmers do not need to know the particulars of Haar cascades or Viola-Jones. Their scripts can simply call trusted cascades (e.g., “Frontal Face,” “Human Body,” “Wall Clock,” and “Mouth”) stored in XML files readily available across the web. Once a computer vision script detects objects in a given cascade, it can extract them from their source files and archive them.

Computer vision techniques may also merge or compare extracted objects with existing objects. Comparisons allow people to confirm or complicate known relationships.¹¹ For instance, when comparing faces, multiple photos of the same person can train algorithms to recognize an “eigenface” (or “unique face”) generated from the principle components of those photos. Although eigenfaces do not actually exist in any lived, social reality, they are fundamental to the processes of face recognition, and data sets with “training face” images for 100 or more people per repository are now common online. One of the most popular sets is the Public Figures Face Database (Pubfig), a “real-world face dataset” that consists of 58,797 internet images of 200 “non-cooperative subjects” that were “taken in completely uncontrolled situations.”¹² While this and other face data sets suggest that training faces are central to big data initiatives anchored in computer vision, humanities practitioners are only beginning to examine the social and cultural implications of treating people’s bodies as big data for vision science. Indeed, more humanities research is needed in this area, especially as it relates to policing and profiling.

For now, it is important to note that computer vision responses to “Is this you?” or “Is this them?” do not stop at recognition or pattern analysis. They also enable predictive modeling and forecasting. For example, in surveillance and forensics industries, snapshots are extracted from video and stitched together to articulate “people trajectories,” which both archive and anticipate people’s movements over time.¹³ Here, the image processing tradition of photogrammetry is clearly linked with artificial intelligence research. As computer vision stitches together a series of objects, it also learns from them and makes suggestions based on them. The programmatic description and reconstruction of the physical world are directed at the past as well as the future. Yet, as is evident in the case of unmanned aer-

Figure 2.1. Real-time detection and tracking using OpenCV, Python, iSight, and a MacBook Pro. Image care of Jentery Sayers. Python script care of Shantnu Tiwari, modified by Jentery Sayers.

ial vehicles (UAVs) or drones, computer vision has real-time intelligence applications, too. (For a still of real-time computer vision, see figure 2.1.)

Many people are familiar with government and private sector investments in UAVs (including military, surveillance, and profiling applications), even if they are not familiar with how UAVs actually work or how computer vision is actually constructed. But, as with OpenCV research, many “extreme hobbyists” are currently building their own UAVs, or “DIY drones,” using low-cost microprocessors, sensors, and actuators.¹⁴ These UAVs can be programmed to follow scripted missions along a series of waypoints, in much the same way one follows directions from a mobile phone when driving a car, allowing the UAV to detect, track, and respond to objects using OpenCV or another library. In doing so, developers directly link recognition to reaction, tightening the loop between pattern analysis and responsive machine behaviors.

Bundled together, the emergence of these techniques raises many questions about how computer vision techniques normalize bodies, environments, and objects and treat them as data, including questions about the relevance of computer vision to privacy and social justice issues. At the same time, many developers are researching computer vision applications in a liminal space between standardized and experimental practice, where the consequences remain uncertain or undefined by policy.

To better understand that liminal space, consider the use of computer vision in the arts, particularly around the question of how standardized machine phenomenology informs experimental network aesthetics. Matt Jones (2011) suggests that computer vision corresponds with a “sensor-vernacular aesthetic” or with “optimised, algorithmic sensor sweeps and compression artefacts.” Somewhere between bits and atoms, a sensor-vernacular aesthetic is “an aesthetic born of the grain of seeing and computation,” with David Berry (2012) pointing to a renaissance of eight-bit visuals, the emergence of “pixel fashion,” and—generally speaking—a widespread obsession with seeing like a machine.¹⁵ Think *Minecraft* and decimated meshes on Thingiverse, or Timo Arnal’s robot-readable world (2012), Martin Backes’s pixelhead (2010), Adam Harvey’s stealth wear (2013), and the machine wanderings of James Bridle’s “New Aesthetic” (2011). Whatever the label or example, such aesthetics have largely revised the notion of technologies as “extensions of man”¹⁶ to suggest that computer vision supplants human vision. In this sense, they are at once humanist, nonhumanist, and object-oriented aesthetics. They throw the very notion of human perspective into relief, understanding computer vision as withdrawn, as beyond human access, as some sort of algorithmic unconscious. At the same time, they demand consensus about what human perception entails in the first place. They need “human” to operate as a stable (or ahistorical, or normative, or universal) category to then displace it with a computer’s phenomenology.

Bruce Sterling suggests that sensor vernacular aesthetics are largely reactive or tactical in character. Their machine wanderings and robot-readable worlds tend to wonder at machine vision—to suspend “Is this you?” from its social dynamics—without systematically intervening or translating it into a meaningful aesthetic category. For instance, in response to the New Aesthetic panel at South by Southwest 2012, Sterling wrote the following in *Wired*: “The New Aesthetic is a genuine aesthetic movement with a weak aesthetic metaphysics. It’s sticky with bogus lyricism” (2012). He also implied that the New Aesthetic is a “glitch-hunt,” adding that “[it] is trying to hack a modern aesthetic, instead of thinking hard enough and working hard enough to build one” (2012). Although Sterling focuses on the New Aesthetic here, his rather frank critiques apply widely to most obsessions with seeing algorithmically, like a machine, in new media arts. And his work on design fiction (2009) and the Internet of Things (2005) are no doubt relevant to the directions of arts-based computer vision practices.

Multiplying Vision

Rather than hacking computer vision, repurposing scripts, or fetishizing machine perspectives, maybe the most pressing challenge for network aesthetics and algorithmic culture is *shifting from a tactical reaction to a strategic articulation of vision infrastructures*. To be sure, this is no small task, especially for humanities practitioners. It might involve interrogating cascading classifiers for their biases and revising them accordingly, much in the way Simone Browne (2010) has approached video surveillance, race, and biometrics. Or it might involve reframing computer vision research to such a degree that mediation is complicated and enriched by not establishing essential, binary ways of seeing.¹⁷ In other words, amid the possibilities of using computer vision for oppressive purposes (e.g., its applications for war and racial profiling), maybe we need vision infrastructures that value partial or contingent vision, much like Donna Haraway's early work (1985) on cyborgs, feminism, and informatics (including, lest that often-overlooked section of "A Cyborg Manifesto" be ignored, her concerns about an informatics of domination).

In digital humanities research, we see some steps toward these vision infrastructures. Although the field has privileged the practical use of OCR to digitize, encode, search, and find texts, it has also pushed machine vision toward some more speculative applications, which allow scholars to interpret—or, better yet, argue—with computers. Here, examples include *Real Face of White Australia* (2011), by Kate Bagnall and Tim Sherratt, and the cultural analytics of Lev Manovich's Software Studies initiative (2009).¹⁸

As noted in Bagnall and Sherratt's chapter, *Real Face* uses a face detection script to foreground non-European Australians not only ignored by the whitewashing of Australia's record-keeping, but also historically subjected to discrimination via the White Australia policy.¹⁹ The project blends an intervention in Australia's archives with a redefinition of vision science. Scripts typically deployed for surveillance and military purposes are instead imagined as mechanisms for critical race studies. At the same time, Bagnall and Sherratt's use of Python and OpenCV for this intervention prompts compelling questions about how race is interpreted as a form (or eigenface) through biometrics and face-detecting algorithms.²⁰ For instance, whose faces are idealized or overlooked by machine vision? How exactly is race a principle component of the training process? When applied to archives as well as everyday spaces (e.g., airports, social networks, and games), what racial biases does computer vision both create

and enable?²¹ By stressing the sociocultural dimensions of face description and reconstruction, these questions—which are still being addressed by scholars—allow us to avoid parsing human from computer vision, and to refrain from delegating interpretative authority and scholarly responsibility to algorithms.

Meanwhile, Manovich's notion of cultural analytics (2009) also resituates computer vision, applying it to the study of art and cultural history as big data. Resonating with Franco Moretti's "distant reading" (2005) of literary history as well as Stephen Best and Sharon Marcus's "surface reading" (2009) of texts, Manovich and his team use OpenCV and other tools to extract features from and visualize patterns in large collections of video, images, and dynamic media (e.g., more than 250 paintings by Mondrian and Rothko).²² Claiming that we have moved from "the stage of 'New Media' to the stage of 'More Media,'" Manovich asserts that this large-scale treatment of visual culture corresponds with the current frequency and ubiquity of new media production.²³ But, beyond the rhetoric of its big data appeals, cultural analytics may be understood as conjectural inquiry, asking an important question of computer vision: What if we read history against the humanist grain, through a combination of automated and manual interpretations? How would such a blended approach force us to reconsider the assumptions we have about art, cultural history, visual media, and perception? While similar questions are commonplace across text-mining and text-analysis projects, very few researchers (aside from the Software Studies research team) have asked them of image and video repositories, where computer vision has arguably the most provocative applications for speculation.

This remark is not to suggest that paradigms for speculative humanities do not already exist. Bethany Nowviskie and Johanna Drucker's speculative computing (2004), Kari Kraus's conjectural criticism (2009), Mark Sample's poetics of nonconsumptive reading (2013), and Devon Elliott, Robert MacDougall, and William J. Turkel's experimental mash-ups (2012) highlight how arguments that "[h]istory is nothing but exteriorities" need not exclude what-if reasoning.²⁴ By positing a variety of ways to perceive, engage, and question the stuff of history, this work collectively foregrounds how scholarship is deeply embodied, why absences or redactions in the record matter as much as inscriptions and patterns, and how computational approaches can be simultaneously procedural and subjunctive. Borrowing significantly from this work, speculative computer vision is not about making truth claims, *per se*. It is about multiplying how scholars see the

past, to include algorithmic procedures without reducing them to algorithmic measures, rational mechanisms for proof, or devices for confirmation. Sample asserts that scholars should, even if counterfactually, “[m]ake the computer model itself an expressive object,” to “[t]urn [their] data into a story, into a game, into art” (2013). The computer model is an argument about how to interpret fiction through techniques that may not correspond with the legacies of close or distant reading. Meanwhile, for Nowviskie and Drucker, the “goal is to place the hermeneut inside a visual and algorithmic system” (2004). This “subjective positioning” (2004) resists proclivities toward overly abstract forms of computation that do not account for embodied knowledge. Similarly, Kraus calls for “the considered manipulation or processing of digital signs with the goal of either recovering a prior configuration or predicting a future or potential one” (2009). Understood this way, a computer is not an innocuous calculator or invisible verification machine; it is a way to at once make, shape, and see the past. And for Elliott, MacDougall, and Turkel, technologies are also creative or experimental media, where the research aim is not the “[e]xact reproduction” or capture of historical artifacts through physical computing or digital fabrication. It is “to create situations in which aspects of the past can be revisited, explored, interrogated, and remixed.”²⁵ Likewise, in the following case of remaking Trouvé’s electro-mobile jewelry, the goal is a speculative vision of how something might have worked, especially when it is neither at hand nor immediately discoverable by a computer.

Moving Jewelry Meets Computer Vision

During the second half of the nineteenth century, Gustave Trouvé prototyped, patented, and built a wide variety of curiosities in France, from electric outboard motors and sewing machines to battery backpacks and telegraphic watches. Informed by a kind of electromagnetic worldview, he also designed a number of battery-powered jewelry pieces, including animated hairpins, blinking stickpins, and other such pieces meant to be worn on the person, often at social gatherings or on stage (e.g., performances of *Faust*). Since he produced this jewelry in small batches, very few pieces remain today. What is more, those that do exist are difficult to access, are quite small, and no longer function as they once did. Still, some of his sketches, patents, and notes have been digitized and are available online via an open-access repository (github.com/uvicmakerlab/trouve) containing more than 200 images cropped by Danielle Morgan from George Barral’s

1891 Trouvé biography, *Histoire d'un inventeur*. Aside from a brief mention in Carolyn Marvin's *When Old Technologies Were New* (1988), Trouvé and his contributions to electricity and magnetism during the nineteenth century are rarely mentioned in science, technology, and media studies. The exact reason for this omission remains unclear, but it might be best attributed to Thomas Edison's prominence at the time, and his continued prominence in the historical record. Additional possibilities include the fact that most of Trouvé's notes and projects were lost during a fire and the few publications about him, including Barral's biography, were written in French and have not been translated.

Still, Trouvé persists through the work of Charlotte Gere and Judy Rudoe. In *Jewellery in the Age of Queen Victoria* (2010), they dedicate approximately three pages of their book to his collaborations with artist Auguste-Germain Cadet-Picard, who manufactured some of the electro-mobile jewelry. At one point, they also mention an important detail: "We have not come across any surviving examples of moving jewels."²⁶ However, immediately after this remark, they issue a peculiar call: "the following contemporary descriptions are given in the hope that they will enable such jewels to be brought to light."²⁷ Here, "brought to light" could be a gesture to today's collectors, an invitation to donate moving jewels to their local memory institutions. But more likely, it implies remaking Trouvé's jewelry based on extant materials, a procedure that poses several challenges: filling in the gaps of material culture, designing today through often-undocumented nineteenth-century techniques, perceiving the particulars of miniatures, reconstructing small objects at scale, and contextualizing the jewelry with attention to its actual use (not just its construction).

With these challenges in mind, Nina Belojevic, Nicole Clouston, Katherine Goertz, Shaun Macpherson, Danielle Morgan, and I applied computer vision practices at work in photogrammetry, laser scanning, and 3D modeling to prototype Trouvé's jewelry and other artifacts like it. Our hope is that, by remaking some of Trouvé's curiosities, we will better understand his contributions to media history, contributions that we believe may be quite relevant to other media studies projects, not to mention current fascinations with wearable technologies.

Currently housed at the Victoria and Albert Museum in London, Trouvé's electro-mobile skull stickpin (1867) is one of the first items we prototyped.²⁸ Originally, the eyes of this stickpin were said to move, and the jaw was said to snap, both when charged by a battery inside the wearer's pocket. (Trouvé is credited with inventing the sealed, reversible battery.) In

this sense, the stickpin might be described as an early “wearable,” or part of what Susan Elizabeth Ryan calls a “dress act” that connects garments and meaning, fashion and technology.²⁹ Yet remediating this stickpin in today’s computational environments revealed its own share of problems. First, the stickpin is quite small (height 9.2 cm, width 1.5 cm, depth 1.6 cm), meaning it would likely require a high-resolution laser scanner, instead of photogrammetry, to become a 3D model on a screen.³⁰ The stickpin’s size also means people may need a microscope or similar vision instrument to reverse-engineer the tiny mechanisms inside the skull. Also, given the lack of extant electro-mobile jewelry in London, Paris, or elsewhere, the stickpin is considered to be somewhat precious, making it risky to handle and arguably impossible to take apart. Finally, even if the pin were laser-scanned, the digital model would probably not scale well for editing and study on a screen. Nevertheless, after extensively studying the pin and Trouvé’s work through various publications in English and French, including patents, drawings, and Barral’s biography, we decided to remake the skull by hand, sketching it in pencil and then carving it into basswood instead of relying on Trouvé’s original materials of gold with enamel (see figure 2.2).

As our prototypes demonstrate, the results are speculative from start to finish. They are not exact reproductions of Trouvé’s stickpin, yet they do allow us to better grasp the stickpin’s composition as well as the challenges Trouvé and Cadet-Picard probably faced when making such a small, detailed piece of electro-mobile jewelry during the 1860s. In their digital and analog formats, models of the pin also prompt us to think through component parts, which are conducive to assembly, rapid prototyping, repair, and editing. (See the digital models in figure 2.3.)

Such conjectural prototypes are fundamental to better understanding features occluded by museum displays, where the object cannot be handled; 2D images in publications and archives, where visual representations are framed by a person’s particular ways of seeing; and the artifacts themselves, the insides of which are difficult to perceive without breakage or damage. In the case of Trouvé’s stickpin, one of the most important occlusions is the internal electromagnetic mechanism, which helped the skull move its eyes and open its jaw. A hidden component of the pin that was probably “black-boxed” to maintain its aesthetic or status as jewelry, this mechanism did not remain in working condition over time, and we wonder whether it was ever reliable. Rather than concluding with this conjecture, however, we treat it as motivation for envisioning a functioning stickpin.

Figure 2.2. Sketch of Trouvé's skull stickpin, in pencil, by Shaun Macpherson.

Part of our research involves prototyping mechanisms that are based, with some obvious anachronisms (e.g., Energizer batteries), on notable electromagnetic technologies from the nineteenth century: the telegraph and the solenoid. These mechanisms can be integrated into the various stickpins we are building with computer vision and then read alongside existing histories of science and technology. For example, our tests to date prompt us to speculate that Trouvé and his team may have been the first to remediate interrupter bells and telegraph sounders into jewelry. However, more research is required to verify whether this is the case.

Once the skull's components were digitized using computer vision techniques, we edited them in "gravity-free" environments such as Rhinoceros, a popular computer-aided design (CAD) application. We could then alter

Figure 2.3. High-resolution digital model of a carved wooden skull based off stickpin sketch in figure 2.2, made using FlexScan3D and an LMI HDI 120 laser scanner.

the scale of each part; zoom for a degree of detail that is impossible without a vision instrument; return, where necessary, to the handmade for more carving; determine whether the digital model is watertight (i.e., conducive to manufacture by a computer numerical control machine); and, throughout the process, examine surface particulars both on-screen and off. In short, we constantly oscillated between digital and tactile media, shifting from this format to that one, from one substance to another, documenting each decision along the way.

Such design and editorial decisions demand tactile media. If the approach were strictly digital, then it would neglect the grain of the handmade that is obvious and even assertive in Trouvé's jewelry. Our aim, then, was to use computer vision to distinguish the skull of the 1867 stickpin from a popular or abstract notion of a skull, employing a combination of digital and analog techniques, or a speculative blend of historical and present-day making, based on extant materials.

We also annotated all digital models for viewing in 3D, letting us precisely describe specific aspects of the stickpin and allowing others to see and interact with all annotations in a single space (as opposed to, say, a distribution of TIFFs or JPGs). These models can be converted into text (G-code) or exported in STL, SBP, or a similar format for fabrication in plastic, metal, wax, or wood. Fabrication allowed us to manufacture the stickpin's parts at or near their original size, even if we hand-carved and

prototyped them at roughly ten times that size. Another appeal of fabrication is that the stickpin could be easily assembled and disassembled during rapid prototyping to test the design prior to batch manufacturing and circulation online or by post. Using a numerical control machine also helped us to manufacture these small objects with more precision and consistency than doing so by hand, even if manual labor was inevitable (and often significant) in the manufacturing process, too.

These appeals of fabrication, together with the various analog-digital convergences at play in our workflow with the stickpin, suggest how computer vision techniques may apply to media history in particular, and humanities research in general. Furthermore, they highlight how the entanglement of human and computer vision,³¹ as opposed to an instrumentalist binary of people and machines, may result in inquiries that are speculative in character, whereby algorithms do something other than confirm history, explain away its causes, graph its relationships, or automate its reconstruction. They can help scholars to materially engage history with a palpable sense of honesty about how we constantly remake the past in the present.

Next Steps

Science, technology, and media studies tell us how histories of technologies brim with contingencies and conjectures that are erased, ignored, or refreshed after the fact. At a minimum, our source material is always resourced material.³² Although computer vision is no exception, a speculative approach to it encourages scholars to foreground their decisions, investments, and uncertainties. This way, scholarly biases and conclusions are not too easily displaced onto technologies, even if some displacement is inevitable. In the particular case of Gustave Trouvé's electro-mobile jewelry, a speculative approach raises numerous epistemological questions about doing media history today, especially through emerging vision technologies. While detailing these issues is beyond the scope of this chapter, I will conclude by outlining them here.

First, when seeing the past through computer vision, how should scholars conjecture about materials—such as inaccessible, broken, or dead media—not at hand? In other words, when and on what grounds do we have license for speculation, and how are we held accountable for it? Once articulated, such a license could spark more than a mere repurposing of new gadgets for historical inquiry; it could support a legible methodology for doing science, technology, and media studies with computer vision.

Second, a growing concern about remaking old media is how to better translate speculative computer vision into cultural criticism. In many ways, these concerns rehearse long-standing debates about “critical” distance and “uncritical” immersion in scholarly research. Where early electro-mobile jewelry is concerned, we must ask how the very act of remaking now cannot be completely detached from material conditions then, including how jewelry visibly marked class, reified gender, and sourced both its materials and labor from Europe’s colonies. Although Trouvé’s jewelry was clearly satirical and never intended as fine art, it was not somehow divorced from the patriarchy, militarism, racism, and inequalities of its time, either. That said, does immersion in material culture and technological processes help scholars to better understand such social and cultural forces, or does it fortify our complicity with them? At the moment, my argument for speculative computer vision involves a combination of immersion and distance, without the assumption that we can ever perceive like anyone did back then.³³ Still, the details and consequences of such a hybrid method need to be elaborated and tested, especially as they relate to media history as a form of social history. Until then, remaking old media may fail to convince critics of science, technology, and culture.

Finally, as computer vision becomes ubiquitous across the academy and popular culture, scholars must continue to explore the effects of its habituation through digital devices—to tease out when decisions are routinely delegated to it, how, under what assumptions, and to what effects. Without such attention, we risk running culture in the background.

Notes

Research for this chapter was conducted on the traditional territory of the Lkwungen-speaking peoples, with support from the Social Sciences and Humanities Research Council (SSHRC), the Canada Foundation for Innovation (CFI), and the British Columbia Knowledge Development Fund (BCKDF) and in collaboration with Nina Belojevic, Nicole Clouston, Katherine Goertz, Shaun Macpherson, and Danielle Morgan. Timothy J. Compeau, Devon Elliott, Edward Jones-Imhotep, Kevin Kee, Kari Kraus, and William J. Turkel provided feedback on drafts.

1. Richard Szeliski, *Computer Vision: Algorithms and Applications* (New York: Springer Science and Business Media, 2010).

2. A. Aneesh, *Virtual Migration: The Programming of Globalization* (Durham, NC: Duke University Press, 2006), 5.

3. Bruno Latour, *Science in Action: How to Follow Scientists and Engineers Through Society* (Cambridge, MA: Harvard University Press, 1987); Alexander R. Galloway, *Protocol: How Control Exists After Decentralization* (Cambridge, MA: MIT Press,

2004); Wendy Hui Kyong Chun, *Programmed Visions: Software and Memory* (Cambridge, MA: MIT Press, 2011).

4. Jentery Sayers, "Prototyping the Past," *Visible Language* 49, no. 3 (2015): 157–77.

5. Seymour Papert, "The Summer Vision Project" (July 1966), <http://dspace.mit.edu/handle/1721.1/6125>; Daniel Crevier, *AI: The Tumultuous History of the Search for Artificial Intelligence* (New York: Basic Books, 1993); Margaret Boden, *Mind as Machine: A History of Cognitive Science* (New York, Oxford University Press, 2006); Szeliski, *Computer Vision*.

6. Crevier *AI: The Tumultuous History*; Boden, *Mind As Machine*; Papert, "Summer Vision Project," 6.

7. Szeliski, *Computer Vision*, 10–13.

8. Marvin Minsky, "A Framework for Representing Knowledge," 1974, <https://web.media.mit.edu/~minsky/papers/Frames/frames.html>

9. Szeliski, *Computer Vision*, 5–10.

10. Paul Viola and Michael Jones, "Robust Real-Time Object Detection," *International Journal of Computer Vision* 57, no. 2 (2001): 137–54.

11. John Resig, "Using Computer Vision to Increase the Research Potential of Photo Archives," *John Resig* (2015), <http://ejohn.org/research/computer-vision-photo-archives/>, accessed January 20, 2015.

12. Columbia University. "Pubfig: Public Figures Face Database" (2010), <http://www.cs.columbia.edu/CAVE/databases/pubfig/>

13. Simone Calderara, Andrea Prati, and Rita Cucchiara, "Video Surveillance and Multimedia Forensics: An Application to Trajectory Analysis," in *Proceedings of the First ACM Workshop on Multimedia in Forensics*, MiFor '09, New York, NY (2009): 13–18, doi:10.1145/1631081.1631085.

14. www.diydrones.com

15. David Berry, "What Is the 'New Aesthetic'?" *Stunlaw* (2012). <http://stunlaw.blogspot.ca/2012/04/what-is-new-aesthetic.html>

16. Marshall McLuhan, *Understanding Media: The Extensions of Man* (New York: Routledge, 1964).

17. John Berger, *Ways of Seeing* (New York: Penguin Books Limited, 1972).

18. Tim Sherratt, "The Real Face of White Australia," *Discontents* (2011), <http://discontents.com.au/the-real-face-of-white-australia/>; Lev Manovich, "Cultural Analytics" *Software Studies* (2009), <http://lab.softwarestudies.com/p/cultural-analytics.html>

19. Kate Bagnall and Tim Sherratt, "Invisible Australians: Living under the White Australia Policy" (2011), <http://invisibleaustralians.org/>

20. Ibid.

21. Simone Browne, "Digital Epidermalization: Race, Identity and Biometrics," *Critical Sociology* 36, no. 1 (2010): 131–50, doi:10.1177/0896920509347144.

22. Franco Moretti, *Graphs, Maps, Trees: Abstract Models for a Literary History* (New York: Verso, 2005).

23. Manovich, "Cultural Analytics."

24. Jonathan Sterne, *The Audible Past: Cultural Origins of Sound Reproduction* (Durham, NC: Duke University Press, 2003), 19.

25. Devon Elliott, William J. Turkel, and Robert MacDougall, "New Old Things: Fabrication, Physical Computing, and Experiment in Historical Practice" *Canadian Journal of Communication* 37, no. 1 (2014): 127. <http://www.cjc-online.ca/index.php/journal/article/view/2506>
26. Charlotte Gere and Judy Rudoe, *Jewellery in the Age of Queen Victoria: A Mirror to the World* (London: British Museum Press, 2010), 213.
27. Ibid.
28. collections.vam.ac.uk/item/O115814/stick-pin-cadet-picard-auguste/
29. Susan Elizabeth Ryan, *Garments of Paradise: Wearable Discourse in the Digital Age* (Cambridge, MA: MIT Press, 2014), 10.
30. Whereas photogrammetry algorithmically stitches a series of 2D images into a 3D model, laser scanning rapidly detects the surface and shape properties of objects and converts that data into a model. Photogrammetry also tends to require a stable object, whereas laser scanning allows for some movement or repositioning during the conversion process.
31. Karen Barad, *Meeting the Universe Halfway: Quantum Physics and the Entanglement of Matter and Meaning* (Durham NC: Duke University Press, 2007).
32. Chun, *Programmed Visions*.
33. Sterne, *The Audible Past*.

CHAPTER 3

Seeing Swinburne

*Toward a Mobile and Augmented-Reality
Edition of Poems and Ballads, 1866*

Bethany Nowviskie and Wayne Graham

The drive to digitize the print culture of the past often obscures the fact that the history of a book is as much a product of material culture as it is of the ideas it contains. In this chapter, Bethany Nowviskie and Wayne Graham present research into the complex history of an infamous collection of poetry subjected to generations of piracy, forgery, and unauthorized printings. How can layers of error and misprints change the text, and can they be peeled away and made intelligible? Nowviskie and Graham share their experiments in applying new methods of digital visualizations to help readers understand the malleable nature of print culture and to assist in the work of literary criticism.

In a literary-historical context, scholarly editing involves the painstaking visual and tactile examination of scores of printed and hand-written documents, to identify their independent bibliographical qualities and to compare sometimes-minute textual differences among them. This is done to understand and represent, for scholarly consideration, change over time. In the case of critical editing—a more interventionist form of textual scholarship—the goal is even *to attempt to reverse* transformations wrought by bygone agents and processes on poetry and prose. Such transformations range from the intentional (such as minor revisions to a manuscript, or “stop the presses!” authorial corrections) to the accidental (such as freaks

of inking and binding, traces of damage to plates, or unauthorized changes in spelling and punctuation that may appear when letterpress type is manually reset). Errors and alterations are also often cumulative, generating near-genetic “family trees,” or textual stemmata, as they are handed down through generations of reprinting. The effect is that different audiences, over time, may receive markedly different literary works—a fact that can tempt us to view textual history as a chained sequence of abstractions. But this sequence is not merely conceptual; it is, rather, always anchored in physical objects. In other words, textual history is as much *material and embodied* as it is social and intellectual.

So, too, is the work of the bibliographer or textual critic, whose forensic findings are typically *predicated on*, and often ultimately *re-mediated into*, manipulable book objects—even today, when a good portion of textual scholarship is conducted and expressed online. Digital archives and editions have made it possible for editors to share an exponentially greater number of high-quality facsimiles of their original sources than would be economically feasible (or even physically possible) through printed works alone. But even when scanned page images are available, users of online editions rarely see or sense the visual and haptic engagement with our textual past that has resulted in the scholarly findings on their screens—the process of unraveling errors and mapping change. Contemporary digital editions thus seem frozen in time and lost in space: divorced from textual scholarship’s basis in book design, textual materiality, bibliographical examination, and optical collation—and stuck in trite, skeuomorphic page-turning interfaces or 1990s click-and-scroll paradigms of design interaction.

Our project takes the physicality of bibliography and textual criticism as its subject and invites student readers and fellow textual historians and critics to see differently, reach out, and step away from their desks. An in-progress scholarly and teaching edition of Algernon Charles Swinburne’s scandal-rocked and fundamentally unstable *Poems and Ballads* (1866) becomes a playground for touch- and tablet-based presentation, augmented-reality interaction with printed book objects, temporal modeling, and experimentation with textual collation and the work of scholarly editing using abstract visualization and immersive virtual reality. This chapter sketches for the first time the extraordinarily textual history of Swinburne’s book (our problem set), surveys the design literature and theoretical landscape for our technical work, and describes four experimental interfaces—three of which are currently in progress at the Scholars’ Lab of the University of Virginia Library—that we feel would bring sophisticated, readerly

scholarly editions to fresh audiences and make the *material work* of textual scholarship visible in new ways.

Poems and Ballads: A Textual History

In the spring of 1866, Algernon Charles Swinburne was preparing for the publication of his first collected book of lyrics by the highly respected house of Edward Moxon and Company. Moxon had, in the past year, issued Swinburne's verse-dramas *Chastelard* and *Atalanta in Calydon* to great critical acclaim, and had asked the poet to edit a volume of Byron, which had appeared in February. To all appearances, Swinburne's relations with his publisher and reading public were excellent. His star was on the rise: his work was known to Queen Victoria, who was later to hint that Swinburne might follow Tennyson as poet laureate.¹ Moxon and Company were taking such care with *Poems and Ballads* that, although proofs had been pulled and authorial corrections made once already, the publisher consented to many costly last-minute changes Swinburne requested—some very minor. These were made by hand and with the means of cancel-leaves (that is, by slicing pages with rejected text out and gluing new ones in), even after the book had gone to press and to the binder. The hand corrections were meticulous. Nowvickie's bibliographic examination of approximately twenty volumes of *Poems and Ballads* even shows evidence of commas being turned into periods at Swinburne's request, by means of scratching at the ink with a blade. Publishers do not undertake such labor for a writer unless they are serious about the integrity of that writer's work.

On July 25, 1866, *Poems and Ballads* was made available to the Victorian reading public. The book was widely advertised and, by all accounts, sold well, despite (or more likely, because of) its controversial subject matter. However, Moxon's heirs, still reeling from an unconnected blasphemy suit, quickly realized that a few initial negative reviews of *Poems and Ballads* were the first symptom of a growing literary scandal, centered around Swinburne's unconventional themes of lesbianism, hermaphroditism, wild sado-masochism, vampirism, and outright sacrilege. By August 10, the poet had been summarily informed that Moxon meant to withdraw the book from publication and sell remaining sheets as scrap paper.²

Swinburne's friends in the Pre-Raphaelite brotherhood—a group of collaborating artists and writers including the Rossetti family—sprang into action. A second, rather less reputable publisher was located, and it was soon arranged that John Camden Hotten (whose titles more commonly

leaned toward railway novels and Victorian erotica) would take over all of Moxon's interests in *Poems and Ballads* as well as in Swinburne's previous and forthcoming work. In September, Hotten reissued *Poems and Ballads*, thriftily using the Moxon volumes that were already prepared—only replacing the original title page with his own. The poems continued to sell at a rapid clip, and Hotten was soon reduced to a stock consisting of faulty, not-yet-corrected books and loose proof sheets that had not received the careful emendations Swinburne ordered before publication. In other words, nothing was left but bad, error-filled copies and loose, error-filled sheets. These texts, too, were bound up, labeled as Hotten publications, and placed on booksellers' shelves. They quickly sold as well, and before the year was out, Hotten was forced to typeset a new edition of the book from scratch—a fact he seems to have hidden from Swinburne in order to avoid paying the royalties due on a new edition.

Bibliographic evidence suggests that Hotten's typesetters used the only copy of *Poems and Ballads* remaining to them—one last text in the uncorrected Moxon state. This means that, between late July and mid-November of 1866, Swinburne's *Poems and Ballads* was put out in at least five versions: the Moxon Uncorrected state, a Moxon Corrected state, a Hotten-labeled publication of Moxon Corrected stock, a Hotten publication of Moxon Uncorrected stock and, finally, a Hotten Reprint. Only the last of these is a *true* new edition set from a Moxon Uncorrected edition and additionally introducing, through the complete resetting of the type, a whole new collection of accidental errors.³

In the meantime, Swinburne and his friends began to feel that his many detractors should be given some response, and October saw the publication of William Michael Rossetti's *Swinburne's Poems and Ballads: A Criticism*, the first independent work of scholarship on the poet, and Swinburne's own defiant pamphlet, *Notes on Poems and Reviews*. Interest in the controversy was building in America as well, and under a previous agreement with Moxon, the Carleton publishing house of New York issued yet another a modified version Swinburne's text, also dated 1866, but now entitled *Laus Veneris: and Other Poems and Ballads*.

Swinburne considered this book a piracy. The New York edition shifts the order of the poems, Americanizes many of Swinburne's spellings, and regularizes punctuation in unauthorized ways. Furthermore, in addition to numerous accidental and intentional compositors' variants that pepper the text, all of the errors originally detected and corrected by Swinburne in the original Moxon state reappear. This means that Carleton (like Hot-

ten) typeset the new edition from a Moxon Uncorrected printing—and Nowvieskie's examination of multiple copies of the 1886 Carleton New York edition has yielded independent textual variants *among them* as well.⁴

Hotten continued to reprint *Poems and Ballads* until his death in 1873, when his publishing house was taken over by an apprentice, Andrew Chatto. The Chatto and Windus editions that follow stretch on into the twentieth century, and they form the basis for modern reading-texts of Swinburne, both in print and electronic form. For instance, the widely used and reproduced Chadwyck-Healey English Poetry database presents Hotten-derived Chatto and Windus texts, as does the sole scholarly online edition of Swinburne, the diplomatic *Swinburne Project*, edited by John Walsh. This means that our most common twenty-first-century conceptions of *Poems and Ballads*—the poetic language we receive today online and in modern printed editions—is not as Swinburne intended it. It has been sculpted by the events of 1866.

The drama of textual variation in *Poems and Ballads* deepens, however, because the first scholar of the textual history of Swinburne was a fascinatingly bad man. Thomas J. Wise, the most eminent bibliographer and book collector of his time, was also a forger and a thief, a person later shown to have capitalized on the trust and admiration of the literary world by sneaking into the British Museum's library with a razorblade in his pocket, which he used to cannibalize title pages and other key leaves from books. These loose pages later played a role in Wise's profitable and private deceptions with a glue-pot. It seems to have been easy to make a valuable first edition out of, say, a fifth edition, if only one had the proper title page and the absolute trust of late Victorian and Edwardian literary London. Soon Wise was drawn to imagine the possibilities of access to a printing press of his very own.

Wise studied and collected the works of some Victorian poets while they were still alive, and he befriended many of them. Swinburne, who had long since retired as a recovering alcoholic to the home of his admirer and caretaker Theodore Watts-Dunton, was especially vulnerable. When Wise appeared on Swinburne's doorstep one day in 1890 with a smile and a hot-off-the-presses forged pamphlet version of *Laus Veneris*, a key poem from the 1866 volume, it didn't take much to confuse the aging and infirm poet into validating the forgery as an extreme rarity.

Swinburne quickly went from having no memory of an 1866 pamphlet, to a firm belief that Moxon, whom he had never trusted, had deceived

him by printing a portion of his manuscript without permission. In seeming generosity, Wise gave the forged document to Swinburne, and it went directly into the poet's private library. Soon after Swinburne's death, the culprit returned, calling on Watts-Dunton with an offer to buy the poet's entire collection. Later, when Wise literally *wrote the book* on all of Swinburne's publications (*A Bibliography of the Writings in Prose and Verse of A. C. Swinburne*), he could cite Swinburne's own letters and commentary about the forged pamphlet, establishing it as genuine. (In fact, Wise embellished a little more, to say that it seemed to have been printed as a test publication, on Moxon's part, to ascertain whether the public would be scandalized by the content of the larger *Poems and Ballads* project). Wise could also state truthfully that his own copy of the extremely rare "Laus Veneris" pamphlet had been found among the poet's personal effects. It was therefore made even more valuable by personal association. Interest was piqued and demand grew. Before long, Wise announced that he had miraculously located a small cache of the pamphlets, gathering dust in a pensioner's attic.⁵ He might be induced to part with a few for the right price.

The fraud was so successful that Wise continued it with other poets. It was not until 1934 that John Carter and Graham Pollard, a bookseller and a graduate student respectively, became suspicious enough of Wise's regular pamphlet discoveries to apply modern—at that point, pioneering—techniques of analytical bibliography to their paper, ink, and typefaces. Chemical and type analysis, published as *An Enquiry into the Nature of Certain Nineteenth-Century Pamphlets*, proved not only that Swinburne's "Laus Veneris" was a forgery, but that many widely accepted first editions of noted authors such as Tennyson, Kipling, and the Brownings were absolute fakes.

The event of Moxon's hasty withdrawal of the first edition of *Poems and Ballads* and the campaign of Thomas J. Wise's deliberate interventions into its reception hover in our story like twin dark stars, exerting force not only on the history and documentary evolution of the work, but on interpretative scholarship being undertaken by modern critics of Swinburne as well. *Poems and Ballads* has never been adequately edited, and Nowiskie's work is geared toward untangling this textual conundrum. In terms of Swinburne studies, because the poet fell into disfavor and was not much examined by bibliographers in the twentieth century, most extant biographies, works of literary criticism, and even rare-book-dealer's pricing guides contain false information about the actual texts and events of 1866. They repeat the story of the "test balloon" edition of "Laus Veneris," reprint words and

phrases Swinburne had once deliberately and painstakingly replaced, misidentify various editions and states, and draw false interpretive conclusions from this textual history about the poet's life and art.

Such a vexed textual history lends itself to experimentation in the form of a digital edition, which might simultaneously emphasize the textual complexity and fascinating materiality of *Poems and Ballads* and creatively place it in relation to alternate timelines and imaginings.

State of the Art: Theoretical and Design Frameworks

The production of such alternative and critically remediated texts would participate in what Kari Kraus has termed “conjectural criticism,” a practice that exploits computational “tools of reconstruction and forecasting” for applications ranging from “the recovery of lost readings in classical texts and the computational modeling of the evolution of a literary work . . . [to] the descent of a natural language.”⁶ Conjectural criticism, in Kraus's terms, is concerned—like scholarly editing and philology—“with issues of transmission, transformation, and prediction (as well as retrodiction),” but also requires an algorithmic and semiotic “computational model of textuality” to position itself “as a counterweight to the material model of textuality that now predominates.”⁷ Our work on the Swinburne project, however, is meant to demonstrate that that weighty textual materiality is itself a fruitful ground for speculation and retrodiction—or even re(tro)vision. We feel the time is right for this work. Heather Love, writing on the empirical revival of digital humanities, the hermeneutics of suspicion, and the “descriptive turn,” notes a pronounced “disengagement” in most sectors of bibliography and material text studies from “critical hermeneutics—and, more generally, from the kind of speculative and abstract thought so common during the heyday of literary theory.”⁸ This disengagement from speculation is problematic and not in line with our understanding of the ends of textual scholarship.

We strongly agree with Kraus that “a cogent theory of conjecture is a desideratum of textual studies,”⁹ and suspect—based on promising results from researchers such as Natalie M. Houston and Ted Underwood—that such a theory can best emerge from and be tested in large-scale analysis. In Underwood's terms, we “don't already understand” some things that have been taken as foundational in the study of poetry and prose: even “the broad outlines of literary history.”¹⁰ And while so-called “distant reading” is most commonly associated with the processing of texts divorced from

their instantiation as typeset pages, or gatherings sewn into bindings—as in Underwood’s typical practice—Houston smartly applies computer vision and optical character recognition techniques to a macroanalysis of thousands on thousands of instances of “the visual page.” This work allows her to see with new eyes “the unremarked material aspects of ordinary books of [nineteenth-century] poetry.”¹¹ “To look at a book of poems through data visualization,” Houston writes, “foregrounds visual and bibliographic codes obvious in the material object but often overlooked in its digital surrogates.”¹² Houston, of course, is not looking at a solitary book, but at as many as her repositories and methods will allow. Our project argues that the ability to apply visual methods and emphasize textual materiality is equally important in the context of a single, hand-held scholarly edition—so long the basic unit of humanistic inquiry. Such an approach can (in the terms of this book) better help general readers and students of textual criticism “see the past.”

In 1965, computer graphics pioneer Ivan Sutherland imagined what he called *The Ultimate Display*:

If the task of the display is to serve as a looking-glass into the mathematical wonderland constructed in computer memory, it should serve as many senses as possible. So far as I know, no one seriously proposes computer displays of smell, or taste. Excellent audio displays exist, but unfortunately we have little ability to have the computer produce meaningful sounds. I want to describe for you a kinesthetic display. . . . The ultimate display would, of course, be a room within which the computer can control the existence of matter. A chair displayed in such a room would be good enough to sit in. Handcuffs displayed in such a room would be confining, and a bullet displayed in such a room would be fatal. With appropriate programming such a display could literally be the Wonderland into which Alice walked.¹³

By the late 1960s, together with Robert Sproull, Sutherland had made rudimentary gestures toward such a fully materialized “display,” by introducing a system that would prove foundational to future interaction designers’ and technologists’ work toward the compelling use of virtualized and augmented 3D environments.¹⁴ Sutherland and Sproull’s initial design, however, had serious technical limitations. Their device—like most modern systems—consisted of a helmet with attached, partially transparent goggles

that projected images to the user. Because of the need to track head movements to calibrate the perspective of displayed wire-frame cubes, this head-mounted display (HMD) required sensors to be physically attached to a long arm suspended from the ceiling. The system was affectionately named the “Sword of Damocles.”¹⁵

Techniques have vastly improved in the nearly fifty years since. Head-mounted displays have not only been miniaturized, but have become commercial products that a growing number of consumers can afford. Facebook’s Oculus Rift creates a fully immersive three-dimensional environment, and with key hires such as John Carmack and Michael Abrash, the company is poised to enter the business of 3D gaming. Systems such as Google Glass, in contrast, allow developers to overlay small amounts of requested or contextual information on a tiny screen, mounted to a lightweight pair of eyeglasses. Both of these platforms leverage the hardware computing available from cellular phones, and the Oculus Rift’s display itself is derived from Android-based phones. Augmented-reality applications, such as Yelp’s Monocle app, which leverages open data and the user’s location to overlay information on the world as viewed through a cell-phone’s camera, have grown in popularity and become increasingly robust due to advances in mobile chip technology and device battery life. And creative uses of virtual and augmented reality by independent developers, such as book artist Amaranth Borsuk, have begun to focus academic and artistic interest on third-party peripherals, like the gesture-based LEAP Motion controller.¹⁶ Tools like these allow users to operate in a post-WIMP manner—that is, to work in gestural and spatially enabled ways beyond the long-established computing paradigm of windows, icons, menus, and pointing devices.¹⁷ We can now imagine the book differently, and perhaps know it better through that imagining.

Taken together, we hope a focus on the material and conjectural, the haptic and the conceptual, the possible and the real, will allow us to participate in Donald Norman’s notion of reflective design—particularly its “subjunctive perspective . . . the possibility space of the *what if?*” so helpfully modeled by the work of Charity Hancock et al. in “Bibliocircuitry and the Design of the Alien Everyday.”¹⁸ We use modern, digital, and augmented-reality reading devices both as platforms for experimentation and as defamiliarizing technologies to allow us to see historical texts alternately and anew—valuing, in the early stages of our work, “critical inquiry over usability and exploratory prototyping over fully-realized productions.”¹⁹ The theoretical, technical, and design frameworks we’ve described here inform

our four experiments toward a digital edition of Swinburne's *Poems and Ballads*.

In a nutshell, we are prototyping and refining:

1. The primary interface to the edition, meant to convey Nowviskie's original textual research and advance usability of TEI-based editions on tablet and mobile devices, at the same time that it teaches users about textual materiality and the work of scholarly editing.
2. A simple, clean, print-on-demand reading text of *Poems and Ballads*, enhanced with footnotes and a wealth of supplementary material, not visible on the page but rather made available through the augmented-reality viewport of a camera phone.
3. A set of interactive timelines that not only convey the textual history of *Poems and Ballads*, but experiment with the alternate and conjectural histories that have been posed by Swinburne's bibliographers, forgers, and critics.
4. A proof-of-concept optical collation mechanism that allows users to compare minute textual differences in two versions of *Poems and Ballads* using the head-mounted virtual-reality display of the Oculus Rift.

All four experiments are described below.

Seeing Swinburne: Four Experiments

1. The Digital Edition

The goal of this experiment is to rethink the basic paradigm for electronic scholarly editions, creating a touch-friendly, contextual, "sliding pane" interface better suited to tablets and mobile devices, but also highly usable on a desktop display. The edition is meant to include textual and scholarly notes, a demonstration of collation techniques in 3D, and a "reveal codes" view of the underlying TEI-XML markup in which all known 1866 versions of *Poems and Ballads* have been encoded. Our emphasis is simultaneously on providing a fresh take on the design of digital editions and on emphasizing the physical, optical, and bibliographic-historical processes that go into the creation of any scholarly edition.

The first step toward a digital edition involved translating Nowviskie's

textual apparatus, which represented the collation, analysis, and categorization of approximately twenty copies of *Poems and Ballads* undertaken over several years. Assisted by our UVa Library colleague Tyler Magill, who built on separate XML documents first created by John Walsh at the University of Indiana and augmented by UVa graduate students Keicy Tolbert and Rob Stilling, we encoded various key groupings of textual witnesses using the Text Encoding Initiative's standards for collation. A TEI technique called "parallel segmentation" allowed us to nest variants and sped the encoding of the various textual witnesses, now combined into a single XML document. We are also encoding explanatory literary and historical information in the form of footnotes written by Nowvskie and Jerome McGann (provided with permission from his 2004 Yale collected reading edition of Swinburne, coedited with Charles L. Sligh).

While the TEI provides a convenient, consistent, and sustainable method of encoding the texts associated with the various printings and assemblages of *Poems and Ballads*, it presents certain technical challenges to facilitating our desired user interactions, both on desktop computers and mobile devices. A typical technique for displaying a TEI-encoded document is to transform its data through various extensible stylesheet language transformations (XSLT) to produce a version suitable for the Web. Editors have typically chosen to render such content as long, scrolling pages corresponding to an entire encoded work. Scholarly editions using this technology commonly either preprocess an entire project into static HTML, or employ a server-based page-rendering engine that can create responses on the fly for a given subset of the underlying XML document. The latter approach treats the XML document as a data store, but lacks many developer conveniences and opportunities for optimization that a more modern data store provides.

To build as flexible and tablet-friendly an edition as possible, we designed a hypermedia application-programming interface (API). This is an information architecture that adheres to representational state transfers of data (REST) based on URL hierarchies, but enhances programming capacities by explicitly providing the URL path any given state can perform. For instance, in a simple state, one may want to return a particular page from the primary reading text of *Poems and Ballads*. In the body of the response from the server, the application explicitly links to other actions that are available to this state, such as the line groups contained on the page, marked stanzas, any textual variants, images, and scholarly textual notes—all of which are also individually addressable through a web URL.

This technical approach allows us the flexibility to abstract how the data is stored from various client applications of the data. By preprocessing our TEL, we can create highly optimized data structures for the various clients we are building to consume, minimizing latency and maximizing the responsiveness of the overall system to the end user. Clients of this data can then use this information in a way that makes the most sense for their particular platforms, but that are still solidly grounded in the structural, scholarly arguments being made by the editor of the edition—in this case, by Nowvskie.

Our first experiment in rethinking a modern scholarly edition, therefore, started by rethinking the role of the internet itself as a delivery mechanism. The last several years have seen a trend away from generating statically accessible web content (such as web pages or downloadable PDF and EPUB files) in favor of more powerful interactions with content in the form of “web applications.” These interactions have been made possible by the work that browser vendors have done in the JavaScript runtime environment to ensure that JavaScript applications can take advantage of all of the central and graphical processing units available in modern computing platforms. Our project uses the most current thinking on developing for mobile devices first, and then on scaling and enhancing the application to take advantage of the additional computing power and desktop display space.

By treating our scholarly edition as an application rather than a series of web pages, we are able to better allow the user to interact with various components of its underlying information. A major goal of this application is to provide a haptic interface, natural to tablet-based interaction, through which users can “slide” additional panes of information over the primary reading text of *Poems and Ballads*. In this way, additional scholarly information, normally relegated to footnotes, is made immediately available within the context of the poem currently being read, but it can also be easily dismissed. Another such panel allows users to access the underlying data structures themselves, revealing the mechanisms and data representations through which our edition has been constructed. In other words, this skeuomorphic sliding pane reveals the project’s underlying code. Still another panel offers an experiment in simulating optical collation through superimposition of red and blue text from differing witnesses. Simple, cheap 3D glasses of the sort provided at movie theaters and in the cereal boxes of our youth reveal floating variants through anaglyphic filtering in much the same way that a bibliographer employing a

Lindstrand Comparator device will observe textual differences as variations in a visual plane.²⁰

Nowiskie's *Poems and Ballads* is, therefore, imagined as a teaching edition about the process of bibliography and about its own construction. We want users to *see* things normally hidden in electronic editions, or even information that has been "hidden in plain sight" in printed books—in the way that students new to descriptive bibliography suddenly notice page signatures (those little, marginal letters and numbers that serve as cues to bookbinders) *everywhere* in eighteenth- and nineteenth-century volumes. By composing the information from various textual witnesses to *Poems and Ballads* through a hypermedia API, we are able to provide contextual, as-needed data, and to recreate different editions of the work on an ad hoc basis. Each bibliographic state of the text may be called on demand, allowing the edition to be used as a teaching tool through which students may compare historical or editorial/eclectic versions of the work. Our goal is to offer a new, mobile-ready, distraction-free, and touch-screen-first design paradigm for modern scholarly editions, while simultaneously emphasizing the physical, optical, and bibliographic-historical processes that go into the creation of any edition. This, too, is a material text.

2. The Printed Edition

Augmented-reality technologies provide unique opportunities for working with physical, printed scholarly and reading editions. Book publishers have begun to experiment with augmented reality to better produce compelling editions of classic books. Penguin Books, for example, recently partnered with Zappar to create eye-catching augmented-reality covers for several novels in English, including *Moby Dick* and *Great Expectations*. With the Zappar application loaded, users can point their mobile devices at a physical book's cover to reveal and provoke an interactive experience. To date, though, such augmented-reality experiments are little more than sales gimmicks. ("Thar she blows!") More compelling user experiences are currently found in the gaming realm. Consider the Harry Potter–licensed *Wonderbook: Book of Spells*, released in 2012 for the PlayStation 3, which works with the gaming station's motion controller and eye camera, allowing users to wave a wand to interact directly with a printed book. And Steven Feiner's work with the Boeing Corporation, toward a system for enhancing real-world construction and repair manuals, leverages augmented-reality technologies to help users accomplish critical tasks more quickly and safely.²¹

In a scholarly context, an augmented-reality client could easily supplement a physical, printed edition. Our second experiment lies in augmented-reality interaction with a simple, uncluttered, printed or print-on-demand “reading text” of *Poems and Ballads*, which can be made to show textual and explanatory footnotes, related artworks and 1866 page images, and other bibliographic or scholarly features through a webcam or mobile phone camera. In addition to suggesting a new relation between serious textual criticism and lightweight and inexpensive classroom reading texts of the “Dover Thrift” variety, this experiment allows us to engage in book design and—appropriately to Swinburne’s Morris-influenced Pre-Raphaelite circle—to be more respectful than digital editions usually are of the “opening” of a book, or its two-page spread, as a visual unit of analysis and display.²²

Anecdotal evidence from creators and readers of digital collections such as the Rossetti Archive suggests that, for classroom use, students and instructors cleave to print. Yet how might teaching practices change if they had ready access through their mobile devices—slim, attractive printed books in hand—to all of the bibliographic and scholarly features of a serious critical edition, and to more supplementary material than could be feasibly printed, even in multiple volumes and the heftiest of tomes? How might augmented reality help bibliographers and textual scholars to work differently with primary sources in archives? Can augmented-reality techniques more effectively bring digitized page images from a far-flung archive into conversation with a newly encountered, undigitized physical object? The same techniques that allow users of our printed “reading text” to access digital facsimiles for reference may open up additional possibilities for optical collation—the visual comparison of one printed page with another—through cellphone cameras. We have a working prototype for this experiment and are currently deciding between the creation of a dedicated “app” for the project, like Zappar, or the design of another responsive web application, similar to our primary electronic edition, using existing JavaScript frameworks.

3. Swinburne in Time

Most of the underlying data for our *Poems and Ballads* edition has been derived from bibliographic examination and takes the form of identified variants related to the production history of the work. These data imply

an evidence-based textual stemma, or branching timeline of variant texts. Such a timeline, however, would be greatly enriched by an expression of the social, cultural, and biographical circumstances of the work's production: its Victorian milieu. In addition to presenting a production history of *Poems and Ballads*, could we imagine producing additional timelines that offer subjective *reception* histories—the story of reviews, responses, and interventions or conjectures? Drawing on original research and the unpublished University of Cambridge doctor-of-philosophy thesis of Clive Simmonds, on conceptual work by Johanna Drucker and Nowviskie—and later by Stan Ruecker et al.—on the *Temporal Modelling Project*, and on the technical planning process for a 3.0 version of the Scholars' Lab's *Neatline* tool, we are wire-framing possible interactive timeline expressions of the history of *Poems and Ballads*.

We can also plot the various attacks on the content and style of *Poems and Ballads* penned in 1866 on an imagined timeline, as well as some immediate defenses by the Rossettis and by Swinburne himself. Similarly, circumstances of Swinburne's personal biography and his later evolution as a poet could be mapped out. And, alongside the scholarship and creative interventions of T. J. Wise, we can temporally position small revivals and moments of critical attention to Swinburne (by Georges Lafourcade in the late 1920s and Jerome McGann in the early 1970s, among very few others), and works by contemporary reviewers, as traced by Clive Simmonds.

In his growing 1870s conservatism, Swinburne himself expressed a desire to reissue *Poems and Ballads* in a new order, rearranging some works and leaving many of the most controversial lyrics out. This means that one could plot the expurgated text that “could have been” alongside the now-established history of *Poems and Ballads*, and creatively map either Swinburne's imaginings of the impact of such a thing, or our own. And other alternate-timeline games could be played. Rikky Rooksby, who wrote the major modern biography of Swinburne, opens his study of the poet's life by imagining Swinburne's funeral—not, as it happened, in 1909 after forty years of declining power in the protective custody of Watts-Dunton, but as if it had occurred in 1866, at the height of the literary scandal of *Poems and Ballads*, and of Swinburne's poetic vigor. If Swinburne had died in 1866, Rooksby asks, would he enjoy the reputation of a Byron today?

All of these later wrinkles—from Wise's fictional production timeline to Swinburne's proposal for an expurgated edition and Rooksby's alternate historical imaginings—extend the already hugely complicated story of the

printings and reprintings of 1866 and seem perpetually to prompt Swinburne and his later editors, critics, and bibliographers to posit alternative chronologies and networks of intent, cause, and effect. One question, then, for our own editorial project might be whether conjectural timelines such as these could be reconciled in productive ways in a visual environment in which the user is *also* modeling a textual stemma meant to have some relationship to the “truth” of the matter.

4. Swinburne in Space

Another avenue for the Swinburne data set applies advances in inexpensive HMDs for immersive virtual reality to the age-old problem of comparing or collating two slightly variant texts. To date, the bibliographic process of optical collation has either taken place unaided by instruments (that is, by simple back-and-forth comparison of two books held in the scholar’s hands), or with the use of specialized optical and mechanical collation devices. In the 1940s, Charlton Hinman built a wardrobe-sized system that employs lights and mirrors to superimpose views of two documents, resting on small tables to the left and right of the user. Where differences between the documents are present, the Hinman Collator causes variant words to appear to blink. Similarly, a 1970s-era Lindstrand Comparator allows the user to align two books housed in a large wooden box with mirrors, so that each eye views one separate image. When the user’s brain reconciles the disparity between the images through binocular convergence, differences are perceived in terms of depth of field, creating the optical illusion that any area of variance floats off the page. In contrast, experiments in computer vision for optical collation such as like the Sapheos Project have focused on automated superimposition of two images on a desktop display, with use of color to indicate regions of difference.

Optical collation machines have typically been quite unwieldy and are either impossible to transport or difficult for scholars to set up in the archives they visit. Even relatively lightweight versions (the McLeod Portable Collator and the Hailey’s Comet) are only “portable” in the sense that they can be transported in large suitcases. HMDs, by comparison, weigh a few pounds and can be easily carried in a protective case. Perhaps even more interestingly, their use is not limited to physically present books. Recently, several companies have announced prototypes of lightweight, cheap, immersive virtual-reality systems. Among these are Face-

book's Oculus Rift, Sony's Morpheus, Samsung's Gear, and Google's Cardboard devices. HMDs for virtual reality depend on stereoscopic displays that render seemingly three-dimensional variants in depth of field, in much the same way as a Lindstrand Comparator. Inspired by these technical developments, we are exploring new options for virtual-reality-based optical collation of texts.

Most libraries and archives now have the capacity to digitize a rare book on request. Because digital surrogates can also be easily created with modern cellphone cameras—themselves increasingly welcome in rare book libraries—bibliographers and textual critics can quickly create collections of page images of various witnesses to a text. For these images to be viewed on an HMD like the Oculus Rift, however, they must be processed so that they are both aligned vertically and corrected to an identical focal length. Once calibrated to one another, the images must then be made to account for the specific barrel distortion of a given HMD's rounded lenses. Yet with the correct distortion in place, page images can be projected to the user for optical collation.

One well-known drawback to working with HMDs is that they separate users cognitively from the world around them. A scholar might quickly forget where his or her keyboard and mouse are located, breaking work flows for textual collation. We are experimenting with a potential work-around for this problem, integrating a third-party 3D motion controller, such as the LEAP Motion or Microsoft Kinect, to provide a gestural interface to our collation application in the form of a set of virtual “hands” for users' interaction with documents. Possible virtual-reality annotation applications include noting differences in the text, adding flags for further investigation, changing the documents or “witnesses” visible on the screen, or virtually turning the leaves of the books.

Ultimately, virtual-reality approaches may provide a more congenial user environment for textual collation, and one particularly suited to our age of mass digitization, allowing scholars to make better use of witnesses from far-flung archives. New page images could be digitized and processed to verify their similarity to a particular edition, or even help discover alternate texts. In keeping with the pedagogical aims of our other Swinburne edition experiments, an interface like this might also serve as a teaching aid, offering experience in bibliographical techniques using tools that students are—in the long run—much more likely to have access to, than to a rare, unwieldy (if beautiful) out-of-production mechanical and optical collation device.

Fault Lines and Failures

“If,” as digital humanities pioneer John Unsworth claimed in 1997, “an electronic scholarly project can’t fail and doesn’t produce new ignorance, then it isn’t worth a damn.”²³ We expect our four experiments to prove their full worth in this regard, wishing for failures and problems relating not only to the technical implementation of the work (such as image-recognition challenges, or the calculation of optical distortions in augmented-reality collation), but also along the natural fault lines that always exist between a human reader or user and her material texts (cf. Kari Kraus’s contribution to the present volume). Our project sees this particular design intersection as a place of research opportunity, because we are working to emphasize materiality and study Victorian textuality—the physical transformations of Swinburne’s work due to human error, intention, and implementation of mechanical processes of remediation such as printing and binding—at the same time that we are remediating historical texts using twenty-first-century technologies and into twenty-first-century *things*. How will the materiality of a new computational object (a head-mounted display, a glossy tablet, the portal of a cellphone) articulate with the historical textual materiality being studied and expressed in our digital edition? How can we best explore and teach the conceptual and physical *disassembly* of poetic texts and book objects using contemporary mobile devices and algorithmic procedures that—as most users experience them—lack hackable or exploitable seams?²⁴ We know that every act of textual transmission is a material, embodied, and historically situated re-mediation. This applies equally to those we study and those we enact. We hope our attempts at seeing Swinburne’s textual condition—complete with successes and failures, past and present—will help us better see and understand our own.

Notes

1. Edmund Kemper Broadus, *The Laureateship: A Study of the Office of Poet Laureate in England* (Oxford: Clarendon Press, 1921), 197.
2. Joss Marsh, *Word Crimes: Blasphemy, Culture, and Literature in Nineteenth-Century England* (Chicago: University of Chicago Press, 1998), 83–90.
3. These conclusions are based on Nowviskie’s research and use the terminology of her in-progress scholarly edition. She has recently discovered a sixth state, as evinced by three previously unexamined witnesses—but hasn’t figured out its place in this sequence—yet! Stay tuned.
4. American texts from 1866 are not represented in our online scholarly edition but will form part of a later study.

5. John Carter and Graham Pollard, *An Enquiry into the Nature of Certain Nineteenth-Century Pamphlets* (London: Constable & Co.; New York: C. Scribner's Sons, 1934), 272ff.
6. Kari Kraus, "Conjectural Criticism: Computing Past and Future Texts," *Digital Humanities Quarterly* 3, no. 4 (2009): 4. <http://www.digitalhumanities.org/dhq/vol/3/4/000069/000069.html>
7. Ibid., 4–5.
8. Heather Love, "Close But Not Deep: Literary Ethics and the Descriptive Turn," *New Literary History* 41, no. 2. (Spring 2010): 382.
9. Ibid., 66.
10. Ted Underwood, "We don't already understand the broad outlines of literary history," *The Stone and the Shell* blog, 8 February 2013, <http://tedunderwood.com/2013/02/08/we-dont-already-know-the-broad-outlines-of-literary-history/>
11. Natalie M. Houston, "Toward a Computational Analysis of Victorian Poetics," *Victorian Studies* 56, no. 3 (Spring 2014): 505.
12. Ibid., 507.
13. Ivan E. Sutherland, "The Ultimate Display," *Proceedings of the Congress of the International Federation of Information Processing IFIP*, vol. 2 (1965), 506–8.
14. See for instance a video of Sutherland's 1968 "Head-Mounted Three Dimensional Display" at https://www.youtube.com/watch?v=7B8aq_rsZao, described in *Proceedings of AFIPS* 68, 757–64.
15. Howard Rheingold, *Virtual Reality* (New York: Summit Books, 1991).
16. Matthew Kirschenbaum and Sarah Werner, "Digital Scholarship and Digital Studies: The State of the Discipline," *Book History* 17 (2014): 446–47.
17. Van Dam, "Post-WIMP User Interfaces: The Human Connection," *Communications of the ACM* 40, no. 2 (February 1997): 63–67.
18. Charity Hancock et al., "Bibliocircuitry and the Design of the Alien Everyday," *Textual Cultures* 8, no. 1, (2013): 76.
19. Ibid., 72.
20. See section 4 for more information on optical collation.
21. Steven Feiner, Blair Macintyre, and Dorée Seligmann, "Knowledge-Based Augmented Reality," *Communications of the ACM* 36, 7 (July 1993): 53ff.
22. William Morris, "The Ideal Book" [1893], in *The Ideal Book: Essays and Lectures on the Arts of the Book*, ed. William S. Peterson (Berkeley: University of California Press, 1982), 70.
23. John Unsworth, "The Importance of Failure," *The Journal of Electronic Publishing* 3, no. 2. (Dec. 1997): 1, <http://hdl.handle.net/2142/192>
24. We thank Bill Turkel and Devon Elliot for this insight, offered in their helpful response to our paper at the workshop that inspired this volume, Niagara-on-the-Lake, November 2014.

CHAPTER 4

Mixed-Reality Design for Broken-World Thinking

Kari Kraus, Derek Hansen, Elizabeth Bonsignore, June Ahn,
Jes Koepfler, Kathryn Kaczmarek Frew, Anthony Pellicone,
and Carlea Holl-Jensen

The past often only survives in fragments. The work of historians, antiquarians, and archaeologists has traditionally involved trying to reconstruct history, to put the past back together, and augmented reality is providing exciting new possibilities to assist in these efforts. Kari Kraus and her research team argue that augmented reality can also be used to explore the knowledge gained from broken, exploded, and unfinished objects. In this chapter, Kraus et al. discuss how they used these ideas in the 2015 augmented reality game DUST, and the potential these ideas hold for helping audiences see the past in new ways.

Within library, archives, and museum settings, augmented-reality (AR) projects are often motivated by an ethos of restoration and repair, resulting in an abundance of tools that digitally reconstruct cultural antiquities ravaged by time and circumstance. Such virtual reconstruction techniques have been used to recreate faces on murals, architectural details on temples, missing words from ancient manuscripts, lost shards from broken vases and vessels, and cityscapes destroyed by earthquakes, among a host of other applications.¹ The theme of *broken-world thinking* is thus central to the concerns of this chapter. Coined by Steven Jackson in “Rethinking Repair,” broken-world thinking “asks what happens when we take erosion, breakdown, and decay, rather than novelty, growth, and progress” as “start-

ing points” in our understanding of old tech and our design of new.² In this chapter we argue that the completionist tendencies prompted by AR in cultural heritage contexts—such as the digital restoration of a broken vase—should be balanced by fragmentary approaches favored by various schools of design and consistent with the historical truth of fragmentation and decay.

Based in part on Kraus’s research into how individuals identify the constituent parts of objects—including broken, obsolete, and semantically ambiguous objects—we distill a set of AR design principles addressed to historians, game designers, and library, archives, and museum professionals. We show these principles at work in two contexts: (1) A mechanical drawing of a telegraph—loosely modeled after exploded-view diagrams—that documents the material origins of Samuel Morse’s invention. The larger aim of this proof of concept is to open a window onto alternative history-enriched design paradigms for AR in science and technology museums. (2) An educational alternate-reality game focused on the deep-time sciences, which ran for two months in early 2015 (see the trailer at <http://fallingdust.com>). Funded by the National Science Foundation, DUST incorporates AR apps and 360-degree panoramas that have been conceptualized in part as *non-finito products*: deliberately unfinished things intended to be completed by users.³

Parts and Wholes: Exploded Views for Science and Technology Museums

The popularity of AR as a technique for virtual reconstruction within museums is understandable given the material barriers to restoration that otherwise exist. Unlike a damaged manuscript whose text can be conjecturally recovered in a new edition that is physically separate from the original, the mending and repair of paintings, statues, and other works of art are carried out in situ, thus potentially confounding any distinctions between the contributions of the original artist and those of the conservator. Within cultural heritage contexts, the established principle of *discernibility* has provided a practical workaround to this problem: any intervention must be visually distinct from the original and yet, paradoxically, harmoniously integrated with it.⁴ In practice this may be accomplished through a variety of means, including the application of thin, striated brush strokes known as *tratteggio*,⁵ or even by creating a recessed zone on a canvas that can function as a kind of safe harbor for experimenting with more audacious conjectures.

The promise of AR is that the art object can be *visually* altered without being *materially* modified. Researchers affiliated with the European Virtual Engineering Foundation, for example, have experimented with projecting a polychromatic digital layer over the stone sculpture of the Madonna displayed in the portico at St. Mary's Cathedral in Spain, thus aesthetically reintegrating lost color into the gothic statuary.⁶ While this approach has much to recommend it, it nonetheless risks propagating an atemporal view of cultural heritage by attempting to experientially reverse the effects of aging, privileging an originary moment over subsequent transformational states, and valuing stasis over change. In this chapter, by contrast, we are interested in exploring alternative paradigms for AR that are not bound by the logic of restoration. Inspired by Daniela Rosner et al.'s analytic category of *material traces*, we embrace ideas of provenance, breakage, wear, endurance, and fragmentation.⁷ "Time," writes Rosner, "has been under-represented in the conceptualization of artifacts and their design."⁸ Drawing on the prevalence of exploded-view 3D models and mechanical drawings in AR marketing design, particularly in the automotive sector, we adapt and reconceptualize them for a museum context.

Like a freeze frame of a detonation, an exploded diagram represents its subject as if its components were hovering in midair, torn asunder yet still perfectly preserving their order and relationships. The object is thus visually suspended between assembly and disassembly, between cohesion on the one hand and physical disintegration on the other. In this split-second time frame, individual parts preserve their allegiance to one another while simultaneously being forced apart. The exploded view is thus a generative metaphor for thinking through the dynamic role of parts in technological design, breakdown, and invention.

The initial impetus for Kraus's research study was a literature review on the subject of technological change. Historically we know that many new technologies have inherited parts from prior technologies. The skateboard remediated the surfboard; the camera pilfered from the gun; the telephone harvested batteries and wires from the telegraph; and early models of the phonograph borrowed the treadle mechanism of the sewing machine. In each of these instances, the logic of part-whole relationships governs the design. "Many of a technology's parts are shared by other technologies," notes Brian Arthur in *The Nature of Technology*, "so a great deal of development happens automatically as components improve in other uses 'outside' the host technology."⁹

To better understand this process, Kraus decided to take a closer look

at the role of parthood in creativity and design. She recruited thirty research subjects at the University of Maryland who were asked to examine six three-dimensional artifacts and complete a written questionnaire about them. These artifacts ranged from the familiar (a book) to the unfamiliar or imaginary (a 3D-printed steampunk variation on the medieval astrolabe) to the broken or visually distinctive (a fork with a bent tine and a rock that appears to have a face, a phenomenon known as *pareidolia*).¹⁰ The questionnaire asked participants to identify each object or hazard a guess about it if they didn't know what it was (*object recognition*); analyze it into its constituent parts (*component analysis*); identify the most significant, distinctive, or salient part (*significant properties*); and describe the process by which they performed these tasks and any challenges they faced (*self-reflection*).

One of the more curious findings of the study was the degree to which individuals showed tolerance for deviation from a prototype. In the case of the broken fork, for example, it was often viewed as whole and redeemable as long as—and this is key—it was also perceived as still useable. In such instances test subjects seemed to mentally normalize the object, suggesting that something like a reparative psychology was at work. Indeed, in three separate cases, the bent tine was never mentioned at all, despite its glaring conspicuousness. For those who concluded the fork was no longer functional, however, the response to the bent tine took on a very different character. In one notable instance, the reaction was spectacularly hostile and indignant: “One [tine] is BENT!! This is no longer a fork!!!! Ahhhh *hack hack.* I can't tell you how distressing it is to look at a fork that is bent out of shape like this. It is not a fork—not a fork—not a fork. . . . I had to cover it back up again. When I uncovered the ball, there was delight—the fork was distress, and it makes a horrible noise on the table. Owwww.”

“Not a fork”: for participants who reached this conclusion, the broken tine functioned as a transformational lever, allowing them to begin to see the utensil as *something else*. The distorted prong “can behave differently than a straight piece of metal,” one test subject observed. It “changes the purpose of the object,” echoed another. The fork became weaponized for one participant, presenting the possibility of “non-trivial nasal penetration”; and it morphed into a trident for someone else, that classic three-pronged spear we associate with Neptune. The invocation of the trident suggests that broken parts may be particularly susceptible to what cognitive scientist Jennifer Freyd calls “representational momentum,” in which individuals project a hypothetical future state onto an object.¹¹ In this case the test subject mentally followed the trajectory of deterioration to its logi-

cal end, specifying a path along which the tine eventually transitions from a merely bent state to a fully detached one, leaving behind something that is resolutely “not a fork.”

The tine of the fork has a counterpart in the form of the needle on the astrolabe.¹² The needle isn’t broken, but it is moveable, a feature that was often singled out as important by test subjects. The way they responded to the needle suggests that the parts of a technology most emblematic of function—of working order and design, such as the minute hand of a clock or the spout of a teakettle—are often simultaneously (and paradoxically) the parts that appear most readily breakable, snapable, and detachable.

The astrolabe, an ancient instrument used by mariners to navigate by stars and planets, provided an interesting case. Nearly 70 percent of participants—a large majority—identified the pointer as the most important part. The measuring blade came in at a distant second, with just 20 percent identifying it as most important. And tied for last place were the stem and eyepiece combo and the cross-hairs, each of which received just under 7 percent of the vote. Test subjects tended to use the pointer diagnostically in an attempt to figure out what the astrolabe was. These diagnoses were conjectural and nearly always inaccurate, but governed by an analogical imagination that lends them plausibility. Like the broken tine of the fork, the pointer triggered the invocation of other objects, including a game spinner, a compass, a clock, an old-timey elevator-floor indicator, a sextant, a gauge for measuring rain, and a steamboat propeller. Seen as the most salient part of the astrolabe, it was also conceptually detached from it, making it cognitively available for grafting onto and visualizing in other objects.

When Kraus presented on this research at the AR workshop in fall 2014, Bill Turkel insightfully remarked that knowing when to suppress the urge to repair may be vital to any culture of invention. Building on Turkel’s point, we’d add that the ability to locate the fault lines in an object, to reason combinatorially with its parts, to conceptually move those parts into and out of different artifactual contexts, allows us to creatively make and unmake the world.¹³

Consider the telegraph as a case study. In 1844, Samuel Morse, a one-time professor of arts and design at New York University, sent a sequence of electromagnetic pulses over wire from Washington, D.C., to Baltimore, Maryland. Those pulses transmitted the first message sent via telegraph. The content of that first coded message—“What hath God wrought!”—was chosen by Annie Ellsworth, the daughter of Morse’s friend.¹⁴ If there was a current of divinity barreling through that wire, as

Ellsworth's choice implies, then there was also one of common humanity: Morse's original telegraph was cobbled together from odds and ends lying about his artist's studio and his brother's print shop. Type slugs, an artist's canvas stretcher, wire from a paper mould, a compositor's stick, and old clock-works were salvaged from the jetsam to create a device that would chart a new communications course for civilization.¹⁵ The telegraph was assemblage art, hacked together from found objects. Or call it sampling and remixing. However one chooses to label it, it is the very banality of Morse's act that makes it what it is: a parable of invention.

Drawing inspiration from this assemblage model of design, we have mocked up an experimental mechanical drawing of a telegraph (figure 4.1), with the aim of suggesting how such models in general, rendered in 2D or 3D when a user captures a museum artifact on her video display, might be implemented as a mobile AR app in a science or history museum. In addition to showing the relationship of the different parts of the technology, as is characteristic of exploded-view diagrams, we have also sketched in shadowy representations of the antecedent technologies from which they inherit components. These ancestral ghosts in the machine give the lie to *ex nihilo* theories of invention and creativity. They also serve as a form of what has been called "material citation," visually documenting the provenance of the parts of the telegraph that make up the whole.¹⁶ Within the field of evolutionary biology, phylogenetic trees are used to visualize such patterns of descent with modification; one potential advantage of our design is that it represents not only *sources*, but also *units* of inheritance (e.g., the pendulum of the clock, but not the clock face, suspension spring, weight arm, or other mechanism). It thus serves as an affirmation of Tim Ingold's dictum that "to understand materials [and artifacts] is to be able to tell their histories."¹⁷

Non-Finito Design in an Educational AR Game

A joint endeavor between Brigham Young University and the University of Maryland in partnership with NASA and Tinder Transmedia, DUST (figure 4.2) is an alternate-reality game designed to appeal to youth aged thirteen to fifteen. Sponsored by the National Science Foundation, the game gives teens the opportunity to learn and apply STEM principles in an informal learning context. DUST is focused on the deep-time sciences: those sciences—including astronomy and evolutionary biology—that deal with processes that occur over millions or billions of years, such as the

Figure 4.1. Mechanical drawing showing the component parts of Morse's telegraph. A painting, grandfather clock, and letterpress type have all been ghosted in using a brown-colored wash to suggest the source artifacts and domains from which Morse harvested parts. The oil painting—one of Morse's own—is titled *The Chapel of the Virgin at Subiaco* (1830). Illustration of Morse's telegraph by Lisa Kong; Copyright Lisa Kong. Reproduced by permission. This illustration does not fall under a Creative Commons license.

Figure 4.2. DUST poster image. (Copyright 2015 BYU. Reproduced by permission. Artist credit: Melissa Manwill. Released under a CC BY-NC-ND license. <https://creativecommons.org/licenses/by-nc-nd/4.0/>.)

formation of the galaxies, the evolution of species, or the continental drift of the earth. DUST is thus a “Big History” game whose fictional events unfold not in human time (measured in years, decades, or centuries), but in cosmic or geologic time (measured in eras, eons, and supereons).

DUST ran live for nine weeks, from January to March 2015, and a re-playable version is currently undergoing play-testing. The game centers on the mysterious collapse of adults worldwide who fall into a coma-like state following a cataclysmic meteor shower. Players “hack” into NASA research databases and engage in collaborative play and inquiry across multiple media platforms to search for answers that will save their parents’ lives. Eventually they discover that the meteor dust is full of microscopic extremophiles whose DNA contains the records of a lost alien civilization, a voice from the DUST, whose world was destroyed. It is these alien microbes that have inadvertently (or perhaps purposely) “infected” the adults, knocking them out, and leaving traces of these records in their brains. By painstakingly reconstructing the ZNA of an alien species, scrambled in the act of genetic transmission, players simultaneously repair the fragmented narrative, scattered—like all good transmedia fictions—across multiple media.

While the game is overtly fictional—encouraging players to inhabit the possibility space of the “what-if”—it nonetheless offers the opportunity to engage with authentic scientific principles and theories, including the chemistry of life, the properties of habitable worlds, deep time and planetary time scales, and scientific instrumentation, such as receivers and antennas, for detecting and analyzing cosmic signals. Thus the fictional narrative serves as a compelling hook to attract teens who otherwise may have no interest in engaging with such scientific phenomena.

DUST’s AR apps are intended to help embed learning activities into the story world and provide tools with which teens can explore the fundamentals of scientific inquiry, such as data collection and analysis. Several of the apps trade on the notion of DNA as a long-term storage medium for archival records. The microscanner, for example, allows players to scan their real-world environments for (fictional) microorganisms, including the alien microbes whose ZNA contains messages to posterity. Two complementary tools—the microbe app and the Dream Decoder—give them the tools to (1) genetically engineer tiny frankencreatures capable of fighting off the aliens by manipulating nucleotide and protein sequences, and (2) decrypt rogue ZNA fragments incorporated into the adults’ native genome by the microbes, revealing the enigmatic message. The health scanner app, the second of DUST’S two mobile AR apps, simulates fMRI

technology, which can be used to perform brain scans, revealing sites of alien microbe infestation, later cured by the supertardigrades genetically engineered by players.

DUST also includes 360-degree panoramic environments created with the Unity Web Player plug-in. There are three simulated 3D environments for players to explore: the International Space Station, a NASA science lab, and an astrobiologist's office. Although these environments began as austere, seemingly uninhabited spaces outfitted with only the most basic supplies, office furniture, and lab equipment, DUST undergraduate codesigners at Brigham Young University and the University of Maryland created "artifacts" for them that helped flesh out the game's characters, themes, plot points, and scientific principles. Acting as a cross between interior decorators and mad scientists, they produced posters for walls and ceilings; books for bookshelves; screen savers, emails, and other documents for the computer monitors; voice mail and audio dictation for cellphones; and lab specimens for workstations.

While each of the environments is born-digital, many of the artifacts populating them began life as physical objects that were digitized. *Spike Field*, for example, is a creative reinterpretation of a concept design originally published in a report commissioned by the Department of Energy in 1993 on warning future civilizations about the hazards of radioactive waste.¹⁸ Created by a talented undergraduate,¹⁹ the new work adorns the wall of the office environment, which has been loosely inspired by the real office of an astrobiologist who served as a consultant on DUST and was a member of the scientific team responsible for the original *Spike Field* design. The artist used pencil to sketch and draw the layout of the scenery, followed by pastel pencils for blending and charcoal and sharpie markers for outlining. The physical work was then digitized for inclusion in the virtual environment.

Like the AR mobile apps, the 3D panoramic environments represent a combined view of physical and virtual reality.²⁰ But whereas the mobile apps use the analog world as a canvas on which to overlay digital information, the panoramic environments use a digital world on which to overlay analog information, such as the artist's drawing. The mobile and web apps thus function as inverse types of mixed reality.

In broad strokes, DUST can be understood as a *non-finito product*. Described in the research and design literature as an unfinished work meant to be completed by users, non-finito products are valued for their ability to open up a design space for user creativity.²¹ Unfinished, fragmented,

sketched, and ruined works have all traditionally fallen within the purview of the non-finito.²² Examples include a novel without an ending (unfinished work); the bust of a statue missing arms or head (fragmented work); a loosely executed drawing not intended as the basis for a final work (sketch); and a castle fallen into oblivion and decay (ruin). Recently the field of human-computer interaction (HCI) has borrowed and adapted the term, extending its reach and influence into the domains of software development and interface design. Within this larger context, Facebook’s “poke” feature has been cited as a quintessential example of non-finito design.²³ Ambiguous and open-ended, it depends on users to establish its purpose and conventions.²⁴

ARGs almost by definition fall squarely within the category of the non-finito. Players must work collaboratively to piece together and advance an adaptive narrative that is integrated into real-world media and spaces, including museums, social media sites, print novels, text messages, and the web. The features that make DUST a non-finito product can be found at all levels of design. Each week a new graphic novel chapter—or “story beat”—is released and accompanied by a call to action that players must respond to before the release of the next story beat, such as posting hypotheses, adding evidence, and collecting and analyzing data. Players have also initiated their own projects in the form of spreadsheets, Google forms, and—in one instance—a computer program for calculating star temperature and mass, all in the service of the game. In addition to being non-finito, DUST is an open-ended ARG, meaning that not every detail about the plot or storyline has been scripted in advance. Some of the best-known education, public media, and museum-focused ARGs have been open-ended, including *World Without Oil*, *Urgent Evoke*, *PHEON*, and *Superstruct*. The designers deliberately cede some control and authority to players, whose contributions help determine the shape and character of the game, as well as final outcomes. DUST’s AR apps and panoramic environments are at the heart of realizing this open-ended vision and fostering user creativity.

Central to DUST’s approach, for instance, are clearly delineated units of information that players can contribute or systematically manipulate. In the case of the microbe experimentation app, players combine, rearrange, and mutate the chemical letters of DNA to produce new forms of microbial life capable of attacking the invasive species affecting the adults. Taken together, the microscanner and experimentation apps reinforce through design the lesson that the genotype (the “inside data” as IRIS, DUST’s resident NASA AI, explains it) codes for the phenotype (the “outside phe-

nomena”). Every living organism, in other words, is “the outward physical manifestation of internally coded, inheritable information.”²⁵ Through combinatorial creativity, players thus help bring resolution to the game by creating the circumstances under which the adults can be woken up.

The mobile and web apps also work in conjunction with each other to structure player experience across media channels. Data collected through the microscanner, for example, is analyzed and manipulated using the microbe app. The player is thus pulled across the so-called “Reality-Virtuality Continuum” as she engages in scientific inquiry:²⁶ a view of the real world is first captured by the player’s smartphone camera and then overlaid with fictional information about the microbes putatively lurking on objects, humans, and animals in the environment. These microscopic organisms are then transferred to the web app, where players analyze and genetically modify them. The progression from reality to mixed reality to virtual reality thus roughly correlates with the different stages of scientific investigation (see note 17 for a definition of *virtual reality* as used in this chapter). Each stage, incomplete in itself, is one step closer to completion when combined with the next stage. While the process is decidedly iterative rather than linear, there is an overall forward momentum. As data is passed from one domain of reality to the next, its value is enhanced, its meaning enriched, its secrets gradually revealed. Players eventually decrypt it with the help of the Dream Decoder, a web app introduced in the final stage of the game. The Dream Decoder lies at the far end of the reality-virtuality continuum, the culmination of all the data collection and experiments performed further upstream with the help of the other mixed-reality apps. Taken together, the mobile and web apps are what allow players to collectively unravel the mystery.

In the case of DUST’s intentionally unfinished 3D environments, the contributions take the form of curated artifacts that add texture and realism to the story world and expand it in whimsical, suspenseful, and potentially unforeseen directions. We have also included user-created content in the environments as a way to recognize player achievements and contributions. Two long-term thinking activities introduced by Violet Cannon, one of our fictional protagonists, for example, resulted in drawings and sketches of warning systems to posterity, one of which was featured on the astrobiologist virtual office wall.

The artifacts might also function as clues or red herrings, depending on what is produced and how the artifacts are received and interpreted. In 1979, Kit Williams, author of the best-selling treasure-hunt book *Mas-*

querade introduced the term “confirmers” to denote a technique whereby multiple clues in an ARG all point to the same solution, thereby assuring players they are “on the right track and following more than a string of chance coincidences.”²⁷ The strategically placed poster and separate infographic on cryptobiosis in the ISS environment function in just this way, reinforcing the message that the hibernation-like state some extremophiles enter into under adverse circumstances is critical to solving the mystery of DUST. By obliquely cross-referencing other mentions of cryptobiosis in the graphic novel and co-lab posts, these artifacts collectively assert the importance of the scientific concept to the game.

Coda

DUST is, both literally and figuratively, a fractured world that requires players to intervene in the molecular machinery of life to salvage it. By piecing together genetic codes, players by extension stitch the seams of the story world. This *seamful design*, to borrow a term from Matthew Chalmers and Ian MacColl,²⁸ creates a unified mixed-reality composition whose sutures are intentionally visible. Like the patchwork design of the telegraph, they accentuate the relationship of parts to parts, drawing attention to sites of player creativity, engagement, invention, and inquiry.

In Japanese culture, *kintsugi* is an ancient method of repairing broken pottery using a special adhesive mixed with gold. This “golden joinery” emphasizes flaws and cracks rather than attempting to disguise them. The idea behind *kintsugi* is that the imperfections in everyday objects bear witness to human and environmental activity. By dusting them with gold, artists preserve the material traces of the past and affirm their cultural value. In this chapter we have similarly advocated for a design paradigm that enhances fault lines, proposing digital overlays inspired by exploded-view diagrams and alternate-reality games that draw on an aesthetics of fragmentation. Rather than trying to repair reality with virtual reality or invisibly stitch them together, mixed-reality design for broken-world thinking invites designers, viewers, and players to fill the cracks and seams with gold.

Notes

The short descriptions of Kraus’s mereology study and the plot of DUST draw on language Kraus and the research team have used in other talks and publications. We thank the students who participated in the mereology study and the students,

teachers, and librarians who served as codesigners and game-runners for DUST, our NSF-funded alternate reality game. We also thank the NSF for supporting this research (NSF AISL award numbers 1323787 and 1323306) and Western Lights and NASA for partnering with us.

1. See the examples listed in Florin Gîrbacia, Silviu Butnariu, “Virtual Restoration of Deteriorated Religious Heritage Objects Using Augmented Reality Technologies,” *European Journal of Science and Theology* 9, no. 2 (2013): 223–31.

2. S. J. Jackson, “Rethinking Repair,” in *Media Technologies: Essays on Communication, Materiality and Society*, eds. T. Gillespie, P. Boczkowski, and K. Foot (Cambridge, MA: MIT Press, 2014), 221.

3. See Jin-min Seok, Jong-bum Woo, and Youn-kyung Lim, “Non-finito Products: A New Design Space of User Creativity for Personal User Experience,” *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York: ACM, 2014).

4. Salvador Muñoz Viñas, *Contemporary Theory of Conservation* (London: Routledge, 2004), 5.

5. For a description, see Magdalena Grenda, “*Tratteggio* Retouch and Its Derivatives as an Image Reintegration Solution in the Process of Restoration,” *CeROArt* 2010, <http://ceroart.revues.org/1700>

6. Rosa Peral, Diego Sagasti, and Sara Sillaurren, “Virtual Restoration of Cultural Heritage through Real-Time 3d Models Projection” Mark Mudge, Nick Ryan, and Roberto Scopigno, eds., *VAST 2005: 6th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage* (Pisa, Italy: Eurographics Association, 2005).

7. Daniela K. Rosner et al., “Designing with Traces,” *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York: ACM, 2013), 1649–58. ACM Digital Library. Web. Accessed Apr. 21, 2015.

8. Rosner et al., “Designing with Traces,” 1657.

9. W. Brian Arthur, *The Nature of Technology: What It Is and How It Evolves*, (London: Penguin Group, 2009).

10. We thank students in Kraus’s 2013 Classification Theory course, who participated in an informal version of this study as a class activity and made recommendations on what types of objects to present to test subjects, including the suggestion that Kraus include an object exhibiting the characteristics of pareidolia.

11. Jennifer Freyd, “Dynamic Mental Representations,” *Psychological Review* 94 (1987): 427–38.

12. To view the 3D-printed astrolabe by the artist Whystler that was used in the study, visit <https://www.shapeways.com/product/FBS8F8973/simple-astrolabe>

13. For a discussion of how counterfactual thinking entails the discovery of “fault lines in reality,” see Ruth Byrne, *The Rational Imagination: How People Create Alternatives to Reality* (Cambridge, MA: MIT Press, 2007), 5.

14. Alexis C. Madrigal. “The First Long-Distance Telegraph Message, Sent This Day in 1844: ‘What Hath God Wrought?’” *The Atlantic*, May 24, 2013. Web, accessed Oct. 7, 2015.

15. Bernard Finn, ed., *Exposing Electronics*, vol. 2 (Boca Raton: CRC Press, 2000), 136.

16. John Chapman and Bisserka Gaydarska, *Parts and Wholes: Fragmentation in Prehistoric Context* (Oxford, UK: Oxbow Books, 2006), 20.
17. Tim Ingold, "Toward an Ecology of Materials," *Annual Review of Anthropology* 41, no. 1 (2012): 434.
18. Kathleen Marie Trauth, Stephen C. Hora, and Robert V. Guzowski, *Expert Judgment on Markers to Deter Inadvertent Human Intrusion into the Waste Isolation Pilot Plant*. No. SAND—92–1382. Sandia National Labs, Albuquerque, NM (United States), 1993.
19. See <https://goo.gl/48xXyz> for a reproduction and artist's credit.
20. In this chapter, we are using the term "virtual reality" to refer to online simulated environments—such as DUST's 3D panoramic models—rather than simulations that rely on special electronic equipment such as VR headsets or gloves.
21. Jin-min Seok et al., "Non-Finito Products: A New Design Space of User Creativity for Personal User Experience," CHI 2014, Toronto, Canada.
22. Seok et al., 695, 701.
23. Seok et al., 697–698.
24. Seok et al., 697.
25. John Blamire, "BIODotEDU," 2001. http://www.brooklyn.cuny.edu/bc/ahp/LAD/C20/C20_Genes.html
26. Paul Milgram, H. Takemura, A. Utsumi, and F. Kishino, "Augmented Reality: A Class of Displays on the Reality-Virtuality Continuum" (pdf), *Proceedings of Telemanipulator and Telepresence Technologies* (1994), Vol. 2351, retrieved March 15, 2007.
27. Kit Williams, *Masquerade: The Complete Book with the Answer Explained* (New York: Workman Publishing Company, 1993).
28. Matthew Chalmers, Ian MacColl, and Marek Bell, "Seamful Design: Showing the Seams in Wearable Computing" IEEE Xplore, Conference Proceedings, (2003): 11–16.

CHAPTER 5

Faster than the Eye

Using Computer Vision to Explore Sources in the History of Stage Magic

Devon Elliott and William J. Turkel

Stage magicians amaze audiences with feats of sleight-of-hand and misdirection that are faster than the human eye can follow. In this chapter, Devon Elliott and William Turkel delve into the magical arts, applying computer vision techniques that are themselves “faster than the eye,” to find the secrets buried in the thousands of images contained in late-nineteenth- and twentieth-century magic periodicals. Computers and humans see differently, Elliott and Turkel explain, and what seems to degrade an image to human eyes can make it crystal clear to a seeing computer. Elliott and Turkel harness this effect to help them see the past in new ways and on a scale far larger than was once possible.

Computer vision provides new ways to analyze primary sources that complement and supplement the historian’s traditional techniques of image analysis. In this chapter we discuss the use of optical character recognition (OCR), face recognition, and other image-processing techniques on a large collection of early-twentieth-century periodicals for stage magicians. Using these techniques, we automatically extracted, classified, and visualized 40,000 advertisements, drawings, and photographs, providing a synoptic view of the changes over time in the imagery used by a particular community. Some of the visual features that we extracted were specific to stage magic in various periods: the recurring use of imps and devils, for example,

or the appropriation and creation of orientalist motifs. Others shed light on the modes of instruction and pictorial cues used to teach methods and techniques of magic; a single image might convey the essence of an action that would be difficult to describe in words, such as the proper way to unobtrusively palm a coin or ball. By exploring the change in imagery over time, we gain insights about a range of phenomena that are poorly documented in textual representations, if at all.

We begin with a description of Elliott's research on the history, technology, and culture of stage magic. By 1895, stage magic had become a popular form of entertainment in America. Successful performers earned public fame touring stages across the country with their shows. At the same time, they were revered as stars by a growing community of amateur magicians. These enthusiasts seldom performed their own magic on stages for public entertainment, but they were successful in other professional fields, such as medicine, science, or business, and they had a deep interest in learning and practicing magic as a hobby.

Magicians were and are known for their secrets. To create an illusion, their methods had to remain hidden, intentionally obfuscated, simply to keep the methods of their craft from public attention. Yet among magicians at this time, knowledge and techniques were widely shared. Sharing their secrets with other magicians would give individuals credit for their creations and innovations, earn them reputations in the magic community, and if all went well, give them a bit of notoriety.

To learn the techniques of creating magical effects, an amateur or aspiring stage performer could turn to books that were published in the mid to late nineteenth century. Several early successful magicians, for example, wrote memoirs that explained some of their feats.¹ One such magician was Jean-Eugene Robert-Houdin, who ushered in the use of popular evening attire as the costume of a magician. Considered "the father of modern magic," Robert-Houdin was also known for his innovative use of electrical technology in apparatus, and for the ways he rhetorically framed the performance of illusions. Another popular instructional book for enthusiasts was *Modern Magic*,² written by the English barrister Angelo Lewis and published in 1877 under the pseudonym Professor Hoffmann. Covering Robert-Houdin's methods extensively, Hoffmann created a series of instructional texts that became popular because they focused on the methods and apparatus necessary to create particular effects; he also included descriptions of stage dress and deportment for the professional magician.

Other documents of the time included a memoir published by Signor

Blitz in 1872 (stage name of the Englishman Antonio van Zandt)³ and a travelogue attributed to the famous American stage magician Harry Kellar, published in 1891.⁴ Fictionalizing and romanticizing accounts of the magicians' success to varying degrees, these works typify the printed material hobbyists were using to learn magic at the end of the 1800s.

At the end of the nineteenth century, however, another form of print was becoming popular with magicians: trade magazines. Edited by amateur magicians, these periodicals focused solely on magic and offered a more dynamic format for magicians to publish their secrets than instructional books and memoirs. The pages featured innovations and variations on methods and effects. These works allowed a wider set of contributors to submit their own articles, which editors could approve for publication. These magazines also performed social functions for the community of magicians. They provided a forum where magicians of all abilities could offer updates on magicians' formal and informal gatherings; editorialize and comment on events and performances; and, as magicians formed trade-oriented organizations and clubs, air club politics. In doing so, these magazines not only maintained a contemporary feel, they helped build and sustain a network of professional, hobbyist, and aspirational magicians.

These magazines also showcased the latest wares on sale for magicians. Advertisements and reviews for new magical apparatus and books increasingly appeared on the pages, and some stores and magic companies began publishing their own magazines to help readers not only learn magical secrets, but know where to find the materials that promised to make them better magicians. Mail-order options extended the reach of these commercial outlets.

A growing network connected amateur, hobbyist magicians to one another through the pages of these magazines and the circulation of magic paraphernalia. Very often these amateur magicians were successful professionals in other fields, and were thus uniquely situated to build and contribute to these networks. They had time to devote to their hobby, and the income to afford magazine subscriptions and magic paraphernalia. Burton Bledstein, writing on the rise of professionalism during this period, has emphasized the roles and functions of command over print media in achieving success in professions.⁵ Having achieved success in their own professions, they applied similar techniques within their hobby to become known, and sometimes renowned.

Of the magic magazines available at this time, two stand out: *Mahatma*, which was founded in 1895 and ran for a decade before folding in 1906,

and *The Sphinx*, which appeared monthly from 1902 until 1953. Both were edited and published in the United States by amateur magicians rather than professional stage magicians. They offer insight into this emergent form of publishing for magicians and the world of stage magic at a time when amateur magicians were rising in prominence among professional performers.

Yet even these publications had limited circulation, because they were often available only to the membership of specific magic organizations, and the clubs themselves were small enough that entire membership lists could be printed in the magazine. Nevertheless, issues survive, and many have been scanned to create digital editions. The most complete collection is held by the Conjuring Arts Research Center⁶ in New York, and we are indebted to their generosity and support in making these sources available. Adam Matthew Digital, in their collection of Victorian popular culture, also has digitized versions of the early issues of *The Sphinx*, and provided page transcriptions rather than the results of optical character recognition, thereby allowing us access to more accurate digital texts for analysis.⁷ Some copies of these magazines are also held in private collections, generally those of magicians, which limits access to them.

There are two basic ways to get a digital copy of pages in a magazine. One way is for a human to read the image and transcribe it into a text editor or word processor. Though typically the most accurate process, this can be slow and expensive. The second way is to use optical character recognition (OCR), which uses automated image processing to attempt to recognize each character and produces a digital version. For a clean, recent image of printed or typescript text, this method is fast and inexpensive; it can also be surprisingly accurate, but uneven lighting, poor print or reproduction quality, unusual fonts, multiple or inconsistent columns, skewed or blurred images, handwritten annotations, watermarks, page damage, and the like all reduce OCR accuracy—sometimes to the point of near gibberish. Sometimes these errors can be corrected with postprocessing,⁸ as is evident in some large digital humanities projects, such as *The Old Bailey Online*, that use human editors to clean up the OCR'd text before providing a further layer of markup in XML. Other projects, including Google Books and its Ngram Viewer,⁹ the Internet Archive,¹⁰ the Hathi Trust,¹¹ and commercial repositories such as JSTOR, all point to the scalability and utility of OCR for providing access to vast amounts of scanned print material, even if the collections continue to correct OCR errors after publishing

data. Despite the drawbacks, OCR is the most direct route to creating vast amounts of textual representation of digital images with text.

The text of millions of books and hundreds of thousands of serials is now accessible for both keyword searching and more sophisticated methods of text mining. Recent projects such as “Viral Texts: Mapping Networks of Reprinting in 19th-Century Newspapers and Magazines,”¹² by Ryan Cordell, Elizabeth Maddock Dillon, and David Smith, have shown the value and insight that can come from using computational methods of analysis on vast amounts of text from digitized primary sources.¹³ As in these projects, having access to a digital version of the text in early magic magazines allows us to subject them to various kinds of computational analysis.

Applying the OCR process to *Mahatma* and *The Sphinx* proved problematic in various ways. First, the page layouts—most in closely spaced, two-column format—were seldom recognized accurately by the OCR process. Instead of parsing the two columns into separate fields of text, OCR often spread the results across columns; consequently, even if the character recognition was nearly error-free, the context of the entire document was not maintained. Although still useful for measuring word frequencies and for text mining based on the bag-of-words approach, these errors produced data that would be unreliable for higher-level analysis such as automatically determining parts of speech and for other context-dependent methods of analysis. For future work, we are planning to do some preprocessing on the page images to extract columns before running OCR.

The publications’ fonts were also difficult for our OCR software to analyze. Although the main body of the text used a single font, article titles were often printed in other fonts. Similarly, articles that began with a stylized large character were often mistranslated by the OCR software, as was the hand-lettered, misaligned text accompanying drawings. Such inconsistencies introduced further errors into the text results generated by OCR.

Another problem with OCR textual representations is that it overemphasizes the importance of text relative to a source in its entirety. While text is a major source of information in these magazines, a quick perusal of the pages reveals that text is not the only interesting thing we could be analyzing. The pages are full of visual elements, too, such as drawings, photographs, and advertisements, but these images merely contribute to OCR errors. Considering that OCR is one part of a larger set of computational techniques known as image processing—a field that

uses computational methods to measure, represent, and transform digital images—we began to examine other techniques that would help us automatically extract images from magazine pages, to both improve the quality of OCR and to create a set of visual representations for traditional study and automated image mining.

At the most basic level, digital images are often represented as a two- or three-dimensional matrix of pixels, or picture elements. Each of these elements can be analyzed by itself (for values such as brightness, saturation, and hue), in the context of its immediate neighbors (to establish image features such as edges, lines, boundaries, or “blobs” of pixels with similar values), or in the context of the statistical properties of the image as a whole (to extract shape, shading, depth, geometry, or motion, or to recognize entities such as faces, people, vehicles, buildings, and so on). Although most of these low-level image-processing techniques do not readily correspond to visual analyses that humanists traditionally use, they will be familiar to anyone who has worked with image-processing software such as Adobe Photoshop; such tasks as brightening or darkening an image, changing the color space, blurring it, detecting edges, selecting a region with the magnetic lasso, and applying a filter are all based on image processing. By coupling machine learning with image processing, we can create systems that learn to sort or classify new images based on training or experience with previous images, and computer vision offers one accessible way to do so.

Image-processing techniques that are applied to camera or sensor data in real time are known collectively as computer vision, and OpenCV, an open-source library for computer vision,¹⁴ provides a way to experiment with various image-processing methods using images from *Mahatma* and *The Sphinx*. Just looking at these magazines quickly reveals the strong visual components of each page—the photographs, drawings, and advertisements—but what, exactly, does a computer see in these pages? Computer vision techniques are based on statistics, and like OCR, they often contain a degree of error. Nevertheless, by using these techniques we can take advantage of machine seeing to look for new things in large volumes of visual data.

Consider, for example, some work presently being done in this area. Chris Adams is using OpenCV to automatically extract images from scanned newspapers in Library of Congress holdings.¹⁵ Using the OpenCV library for Python, he converted each page image to grayscale, eroded it (an image-processing technique that removes pixels on the boundaries of image elements), and then binarized it (rendered each pixel as black or

white based on a threshold value); this work clustered similar regions in the page image, causing them to flow together into blobs, as it were. He then used another OpenCV function to define regions of connected contours and, using that data, to identify the portions of the page that were most likely to be images so he could clip out those elements for further analysis. Adams's results were consistent and promising on the newspaper images he tested, and since the images Adams was working with are similar in many ways to the page images of *Mahatma* and *The Sphinx*, we applied a similar set of techniques to the magic journals, using code written in Processing with OpenCV (Elliott) and in Mathematica (Turkel).¹⁶

Initially created by Ben Fry and Casey Reas at the MIT Media Lab, Processing is a Java-based programming environment intended to help artists and designers explore and create with code.¹⁷ Well-suited to working with visual images, Processing is predicated on the idea of generating “sketches” with code—just as artists might sketch to initiate and develop their work. To make more complex functions accessible to the sketches in development, software libraries can be added to the Processing environment. Greg Borenstein developed such a library in OpenCV, one that allows for flexible access to the OpenCV toolkit within the Processing development environment.¹⁸ Elliott's initial sketches with the Adams method were promising. Using scanned images from Albert A. Hopkins's *Magic*—a nineteenth-century book that explains magic illusions, stage effects, and technological innovations¹⁹—the Processing technique consistently isolated visual images within pages of the book and exported those regions as digital images in their own right.²⁰ The method was also scalable. After proving to be successful on a handful of pages, we readily generated a folder containing images extracted from the entire book.

When we applied this method to the pages of *Mahatma* and *The Sphinx*, we could define drawings, photographs, and advertisements within pages; automatically extract over 40,000 images; and easily export and save those elements as independent images. Just as OCR could export text results, this technique could generate collections of images extracted from magazine pages, providing a new data set of scanned, digital images.

Parameterization of variables is a common method employed in the fields of creative coding²¹ and generative design.²² We created interactive applications that allowed an image to be viewed and the binarization threshold value adjusted with a slider, with results displayed in real time. This had the advantage of allowing us to see which portions of the image were defined by the software given particular threshold values. When

appropriate regions were highlighted, the resulting value was applied across a directory containing images of all pages of that publication. Significant binarization threshold values varied across different documents, but seemed consistent within a given document. We also parameterized the erosion of the image. With some scans erosion improved results, and with others it did not, so it was handy to be able to turn it on or off with a checkbox.

This process of tuning the image-processing settings for specific documents not only results in better results for the extraction, but the interactive and responsive interface offers users more insight into the process of what computer vision does with the documents; in doing so, it allows a user to play around with the settings to attempt to achieve more useful results.²³ That engagement also makes the underlying document transformations more visible, suggesting other possibilities for manipulation and analysis.

In this method, digital images are actually simplified in various transformations. A full-color image is reduced to greyscale. The greyscale image is binarized so each pixel is either black or white. The fine details of the image are then eroded, resulting in a simpler version with less discrete edges. The results of these transformations are, arguably, less useful to human eyes, but they allow computers to “see” and take actions, which can then be applied on a large scale to generate meaningful results.²⁴ In other words, a computer cannot see things as we do, nor should we expect it to, but the unique ways computers can be made to see things are not necessarily less useful ways of seeing.

While tools do not yet exist for humanities scholars to work extensively with this sort of information, there are parallels here with the integration by scholars of OCR and methods for text analysis. Ideally, the process of image extraction will lead to a set of digital representations that can be mined for various features that will not only be interpreted in their own right (giving us new insights into cultural analytics), but that will lead to new and better methods of extracting and handling digital visual sources. In our own work, we are trying to feed the results of our analyses back into the initial steps of image extraction, and the following examples demonstrate this work in progress.

Applying the techniques above to the two sets of magazines, we produced over 40,000 extracted images, giving us enough data to attempt to draw some meaning from them. One way to see the images is linearly—one image after another. In figure 1, time runs left to right across the x-axis,

Figure 5.1. Images extracted from *The Sphinx*, arranged by median brightness across all pages.

from early to late; brightness is plotted on the y-axis, with light images at the top and dark at the bottom.

Although *The Sphinx* was published for many more issues than *Ma-hatma*, resulting in substantially more images, the linear depictions reveal some patterns. For example, a dense row of bright thumbnails appears along the top of each overview image, and closer inspection reveals those to be primarily advertisements and drawings. Advertisements had extensive text, with white space around each character and word. Images within advertisements were also bordered with white space. We use OCR on the text in advertisements to separate them from the line drawings, which typically had the highest median brightness of all. If you do OCR on just about any digital image, you will get some results, but these are usually meaningless strings of random characters. However, when we OCRed an image that has some text in it (such as an advertisement), we found some of the words were recognized correctly, and some of these words will appear in a dictionary. So if dictionary words exist in an image, it is probably an advertisement. The main drawback here is that OCR is a time-intensive process, and each set of OCR results has to then be tested against a dictionary.

Nevertheless, having a synoptic view of advertising across the entire run of a journal offers interesting insights. For example, Welch's grape juice is block-advertised in early issues of *The Sphinx*, but eventually, all advertising is magic-related, a tradition that continues today—only magic products are advertised in magic magazines. The reason for this shift warrants more study, as it could reflect the advertising priorities and trends of the time, or perhaps the number of magic vendors increased to the point where there was simply no advertising space available for advertisers outside of the community. Alternately, it might reflect the personal connections that various editors had to the companies that advertised early on.

More varied than advertisements, drawings pose a greater problem for automated classification. In some cases the drawings are imagery used to represent elements of instruction (e.g., “palm the card like this”). Sometimes they showcase products for advertising. Sometimes they are company logos. Given the importance of card tricks in the history of magic, we have begun experimenting with computer vision methods to recognize the card pips, so that we can classify an image as a (probable) card trick if it contains clubs, hearts, spades, or diamonds.

Machine learning and classification offers the most potential for clustering images automatically. As we isolate particular sets of images, such as advertisements or images of playing cards, we can use these to train a machine learner to quickly “look through” vast sets of automatically extracted images in an attempt to locate more instances of the same kind.

We are exploring the use of machine learning for automated face detection to pull out portraits of people, and also as a means of sorting images into other categories of various kinds, such as drawings or charts.²⁵ The photographs are particularly useful for learning about the magicians these magazine editors and their readers valued and about trends in the profession. For instance, the vast majority are men, but the costumes change over time. Additionally, the focus gradually shifts from successful professional magicians to include a wider variety of amateurs. Such findings show the shift of authority and legitimacy in the magic community over time.

Based on our experience, Mathematica provides the fastest and most accessible route to this sort of work with versatile and robust functions for image analysis and computer vision built into the platform. The distinctions made between advertising, drawings, and photographs could be used as initial classes, and a classifier built around examples of each of those types.

Although much more work remains to be done with data sets of im-

ages that have been automatically extracted from digitized documents, we believe that the field shows enormous promise. Tim Sherratt's work using face detection on archival collections demonstrates how computer vision can make collections more accessible and can highlight themes within those collections.²⁶ Most photographs in *The Sphinx* and *Mabatma* include people, so face-detection algorithms seem useful for collections of images.²⁷ The British Library has also shown interest in the imagery contained in their document holdings, and they have released one million images for public use.²⁸ More recently, millions of images extracted from digitized publications in the Internet Archive's collection have been added to Flickr, with accompanying text from nearby pages making a searchable form for those images.²⁹ The *New York Times* has also launched a crowd-sourced effort to isolate the advertising elements within its collection for researchers.³⁰ At an individual level, too, humanists can liberate the images that lie in digital collections that they have created for particular projects. Each of these collections hints at the untapped potential of the images that lie within.

Notes

Special thanks to William Kalush, executive director, and Dr. Alexis Lynne Pavénick, former head librarian, at the Conjuring Arts Research Center for support and access to their digital collections for this research. The Conjuring Arts Research Center is a not-for-profit organization dedicated to the preservation and interpretation of magic and its allied arts: <http://conjuringarts.org>. We are also grateful to the Social Sciences and Humanities Research Council of Canada for funding to support this project.

1. Jean-Eugene Robert-Houdin, *Memoires of Robert-Houdin: Ambassador, Author and Conjurer*, ed. Dr. R. Shelton Mackenzie (Philadelphia: Geo. G. Evans, 1859), <https://archive.org/details/memoirsofroberth00robe>

2. Professor Hoffmann, *Modern Magic: A Practical Treatise on the Art of Conjuring*, 2nd ed. (London: George Rutledge and Sons, 1877), <https://archive.org/details/modernmagic00hoffgoog>

3. Signor Blitz, *Fifty Years in the Magic Circle; Being an Account of the Author's Professional Life; His Wonderful Tricks and Feats; With Laughable Incidents, and Adventures as a Magician, Necromancer, and Ventriloquist* (Hartford, CT: Belknap & Bliss, 1871).

4. Harry Kellar, *A Magician's Tour Up and Down and Round About the Earth: Being the Life and Adventures of the American Nostradamus* (Chicago: Donnohue, Henneberry & Co., 1891), <https://archive.org/details/magicianstourupa00kellrich>

5. Burton J. Bledstein, *The Culture of Professionalism: The Middle Class and the Development of Higher Education in America* (New York: Norton, 1976).

6. Conjuring Arts Research Center, New York, <http://conjuringarts.org/>

7. Adam Matthew Digital, Victorian Popular Culture Collection, <http://www.amdigital.co.uk/m-collections/collection/victorian-popular-culture/>
8. Ted Underwood, "The Challenges of Digital Work on Early 19c. Collections," <http://tedunderwood.com/2011/10/07/the-challenges-of-digital-work-on-early-19c-collections/>
9. Google Ngram Viewer, <https://books.google.com/ngrams>
10. Internet Archive, <http://www.archive.org/>
11. Hathi Trust, <http://www.hathitrust.org/>
12. Ryan Cordell and David Smith, "Viral Texts: Mapping Networks of Re-printing in 19th-Century Newspapers and Magazines" (2017), <http://viraltexts.org>
13. David A. Smith, Ryan Cordell, and Elizabeth Maddock Dillon, "Infectious Texts: Modeling Text Reuse in Nineteenth-Century Newspapers," in IEEE Computer Society Press for the Proceedings of the Workshop on Big Humanities, <http://viraltexts.org/infect-bighum-2013.pdf>
14. <http://docs.opencv.org>
15. Chris Adams, "Extracting Images from Scanned Book Pages," <http://chris.improbable.org/2013/08/31/extracting-images-from-scanned-pages/>
16. See Jones-Imhotep and Turkel in this volume.
17. Processing Foundation, <https://processing.org/>
18. Greg Borenstein, <https://github.com/atduskgreg/opencv-processing>
19. Albert A. Hopkins, ed., *Magic; stage illusions and scientific diversions, including trick photography* (London: Sampson Low, Marston and Company, 1897), <https://archive.org/details/magicstageillusi00hopk>
20. Ibid.
21. Casey Reas, Chandler McWilliams, and Jeroen Barendse, *Form+Code in Design, Art, and Architecture* (New York: Princeton Architectural Press, 2010), <http://formandcode.com/code-examples/parameterize-chair>
22. Hartmut Bohnacker, Benedikt Gross, Julia Laub, and Claudius Lazzaroni, *Generative Design: Visualize, Program, and Create with Processing* (New York: Princeton Architectural Press, 2012).
23. Stephen Ramsay, "The Hermeneutics of Screwing Around," <http://www.playingwithhistory.com/wp-content/uploads/2010/04/hermeneutics.pdf>
24. Mark Sample, "Notes Towards a Deformed Humanities," <http://www.sample reality.com/2012/05/02/notes-towards-a-deformed-humanities/>
25. See Jones-Imhotep and Turkel in this volume for details.
26. Tim Sherratt, "The Real Face of White Australia," <http://invisibleaustralians.org/faces/>
27. See Sayers in this volume.
28. The British Library, "A million first steps," <http://britishlibrary.typepad.co.uk/digital-scholarship/2013/12/a-million-first-steps.html>
29. Robert Miller, "Millions of historic images posted to Flickr," <http://blog.archive.org/2014/08/29/millions-of-historic-images-posted-to-flickr/>
30. *The New York Times*, *Madison*, <http://madison.nytimes.com/>

CHAPTER 6

The Analog Archive

Image-Mining the History of Electronics

Edward Jones-Imhotep and William J. Turkel

Few of us ever stop to consider the symbols all around us, their origin, or the authority we grant them. Yet they are vital to how we live and think, and they are especially important for understanding complex systems. In this chapter, William Turkel and Edward Jones-Imhotep explore a crucial period in the history of technology, when engineers debated how to represent and comprehend new advances in circuitry. Jones-Imhotep and Turkel discuss the deeper concerns and anxieties embedded in this technical debate about symbols and meaning, and present their method for seeing the past by collecting and analyzing the vast quantity of historic schematics and designs available on the internet. By applying computer vision and machine learning to thousands of schematics produced and published in the mid-twentieth century, they hope to uncover the deliberations, disputes, and the rationales behind the evolution of electronics.

William Higginbotham specialized in creating virtual worlds. Early on in World War II, he had worked at MIT's Radiation Laboratory, developing electronic circuits that would trace out radar echoes from airborne, shipborne, and land-based radars on cathode ray displays. Before the end of the war, he moved from radar to atomic weapons and became an electronics group leader at Los Alamos, designing pulse circuits to monitor and trigger the world's first atomic bomb. But as he grew increasingly critical of the atomic weapons he had helped create, he pursued his long-standing inter-

ests in representing the world electronically, working at Brookhaven National Laboratory, the nuclear research facility created in 1947 to explore the peaceful uses of atomic energy.¹ In October 1958, as the head of its Instrumentation Division, Higginbotham sketched a now-legendary electronic circuit, often seen as a precursor to millions of other virtual worlds to come. Eager to demonstrate the broad appeal of the laboratory's work, he used his division's new analog computers to create a game for the public. In two hours, he had sketched the circuits that would make it possible, laying out the hardware for one of the world's first video games—Tennis for Two. He would later add controls for gravity, so that players could simulate game play on Jupiter or the moon. But our interest here is not so much in the fascinating history of Higginbotham's virtual worlds. Instead, we are interested in uncovering some of the practices he used to draw the electronic circuits that helped generate them.

Higginbotham's sketch hides as much as it reveals. Behind some of the diagram's most innocuous icons—the triangles denoting the operational amplifiers that Higginbotham helped develop—raged a fierce debate over how to draw electronics, particularly transistors and vacuum tubes. That debate reached its peak just as Higginbotham was sketching his circuits, and it spawned dozens of competing schemes that represented now-lost ways of thinking about electronics and about the power and peril of images in the middle of the twentieth century.

In this paper, we do several things. First, we use the example of the visual culture of electronics in the mid-twentieth century to describe a more general, large-scale digital project designed to support collaborative research in the history of electronics and computing. Recognizing that the work we describe here is very much a rough work in progress, we stress that all of our statements should be taken as preliminary.

On the methodological side, we describe how techniques of web-spidering, machine learning, text-mining, image processing, and computer vision can be used to compile a multiterabyte collection of digital historical sources, which can then be subjected to both supervised and unsupervised analysis. (In the former, a human analyst provides the computer with feedback about the correct way to understand or classify a source.) In particular, we focus on using computer vision and image-processing techniques to automatically extract and classify images (e.g., photographs, schematics, graphs, drawings, and equations) that appear in the mid-twentieth-century literature on electronics, thereby allowing us to better identify nonstan-

dard ways of representing electronic components and recapture some of the historical variation in electronic visual culture in the twentieth century.

To emphasize the value of that historical project, we discuss how the history of circuit diagrams is not only a history of visual and material practices, but part of a broader social and cultural history. For technologists in the early Cold War, the seemingly simple question of how circuits should be drawn entailed much deeper concerns: How did people understand the workings of electronics? How should technical workers be commanded? Who could be trusted to design properly? What dangers arose as line drawings transformed into artifacts? The various and competing methods for drawing electronic circuits in the 1950s and 1960s formed part of a set of Cold War anxieties over the possibilities and shortcomings of human reason. Within that context, technologists portrayed now-defunct rules and symbols as structures that would safeguard the reliability of both humans and machines.² Our project explores one way to recover the visual culture at the core of those concerns.

Our larger argument relates to the methodological issues of image-mining and classification. On one hand, it focuses primarily on the roles of pictorial information in the practice of electronics; on the other, it emphasizes (digital) historiography in the age of big data and machines that can see. Thus we begin with a brief history of circuit diagrams that shows their importance in fields beyond electronics. We concentrate especially on the issues that surrounded them in the early Cold War, when drawing became an area of intense concern. We then move to a detailed discussion of the digital techniques and machine vision that can help explore this forgotten visual culture and the practices that produced it. Finally, we turn to a discussion of the context in which this kind of project becomes possible and even necessary in an age of digital photography, drone warfare, self-driving cars, and the war on terror.

The Practice of Illustration

Unlike mechanical engineering diagrams, which developed sophisticated ways of rendering the precise physical layout of machines through isometric drawings and projective geometry, circuit diagrams evolved to show functional relationships. Electrical and electronic diagrams have their origins in depictions of experimental apparatus, where they were used to show interconnections between various components. Sometimes

portrayed as stripped-down reinterpretations of the beautiful isometric drawings of electrical instruments going back to the eighteenth century, these diagrams provided possibilities for replicating experimental setups in electrical science.

Into the late nineteenth century, and especially in electromechanical systems, these diagrams often mixed elements of realism and abstraction in hybrid depictions to show the specific physical characteristics of mechanical devices alongside the functional role of electrical components, as evident in figure 6.1 where the symbol P represents a battery while the rest of the mechanical system is rendered isometrically. The individual symbols that made up these diagrams were generally designed to “suggest” the physical form of components and—particularly through a set of mechanical analogies and hydrodynamical metaphors such as “potential” and “flow”—to tie that macroscopic physical form to the actions of electrons, and by extension, to electronic function. The symbol for the vacuum tube, for instance, evolved from a sketch of the device’s more realistic patent drawing, trading on the portrayal of the device as a “valve,” restricting or permitting the “flow” of electricity through portions of the circuit. By the 1890s, the most common electronic components—resistors, capacitors, inductors—had more-or-less standardized symbols (the earliest being the symbol for the battery or electrical pile), originating in this mix of abstracted form and “suggested” function and, in the decades that followed, virtually every icon in widespread use followed suit.

Paper Tools and Practice

The diagrams formed instruments themselves by the early twentieth century. Extending the late-nineteenth-century attempts to unify science under electromagnetism, circuit diagrams came to represent mathematical relationships that, in an age of analogies, could be applied much more broadly and powerfully to a host of explicitly nonelectrical phenomena. Because energy was conserved, rather than being created or destroyed, physical phenomena could be transduced into electrical currents, voltages, or charges by a variety of sensors. Since the nineteenth century, an ever-expanding array of sensors was developed to convert almost any imaginable phenomenon into an electrical signal. These electrical signals then served as analogs for the phenomenon itself, allowing a change in the phenomenon to be modeled by manipulating its electronic analog.³ Because those phenomena could be modeled by electronics, they in turn began to be rendered as circuits.⁴

Figure 6.1. Early example of a transmitting telephone. From Théodore du Moncel, *Le Téléphone, le Microphone et le Phonographe* (Paris: Hachette, 1878), 51.

Already in the 1890s, in areas such as electroacoustics, researchers had begun depicting hybrid mechanical-electrical devices as fully electrical systems. In areas such as telephony, the close proximity of sound and electrical phenomena, as well as the telephone's status as a transducer, encouraged analogies between sound and signal propagation.⁵ In the period between the world wars, researchers at the Norwegian Institute of Technology used circuit diagrams to represent sound transmission through a double wall (see figure 6.2), while researchers at Bell laboratories represented the vibration of vocal cords using equivalent circuits (see figure 6.3). Even the iconic London Underground map was developed while its creator, a part-

Figure 6.2. The electric equivalent-circuit diagram of a double wall as an oscillating system with damping. Berg and Holtsmark, “Schallisolation von Doppelwänden,” 1935 (cit. n. 12), 75.

time electrical draftsman named Harry Beck, was sketching an electrical circuit diagram.

Through that expanded use of electronic schematics, the practices and conventions for drawing the diagrams became essential to making sense of phenomena in a wide range of disciplines. And as they became more common, their production and use increasingly entailed a specific division of labor between those who sketched the diagrams, using them as paper tools

Figure 6.3. The self-exciting radio-tube transmitter as an equivalent circuit for the human voice. Trendelenburg, *Fortschritte der physikalischen und technischen Akustik*, 1934 (cit. n. 73), 76.

for the conceptual work of design, and those who drafted them and built the devices they represented.⁶ It is this role as powerful paper tools, as the locus of a set of visual and conceptual practices that go beyond straightforward questions of representation, that we want to focus on for the rest of the chapter.

Errors and Drawing Techniques

Following World War II, circuit diagrams became sites of intense anxiety over the relationship between drawing practices and the individual circuit designer. Some of these concerns focused on the public function of the drawings, their role in communicating and coordinating among the various groups implicated in the design and construction of electronic artifacts.

Schematics had sketched out the electronic apparatus at the center of the recent war, and throughout the 1940s, they delineated increasingly complex devices (now created collaboratively with separate technicians contributing individual circuits). As they did so, the practices of drawing

schematics took on a moral cast. The British physicist L. H. Bainbridge-Bell, who worked on radar development in the 1930s and early 1940s, argued that circuit diagrams stood apart from virtually all other technical drawings. They were, according to Bainbridge-Bell, highly specialized thought tools that were not to be confused with, say, layout drawings, which merely gave spatial relations. Arguing for a cultivated distance from the realism of mechanical drawings, and rejecting the complex history of drawing practices that had mixed the mechanical and the isometric with the electrical and the symbolic, he targeted “those who spoil a circuit diagram by trying to combine it with a layout drawing and end by producing something which is suitable for neither purpose.”⁷ Like many others, he would argue for special drawing conventions that characterized the images. He called them “table manners,” and in doing so, he gave voice to a growing vision that cast the electronic arts in terms of a wider set of social and, particularly, moral practices.⁸

Alongside growing concerns about the public role of diagrams were anxieties about their role in private, individual acts of reasoned interpretation. Beginning in the interwar period, deeply shaken by the apparent irrationality of the Great War and concerned about persistent cultural traditions in multiethnic societies, observers across a broad range of disciplines had begun pointing to a “crisis of reason” in Western democracies that potentially called into question even the possibility of self-governance. After the war, social scientists would struggle with those questions, and by the late 1940s, experts in fields from psychoanalysis to anthropology, sociology, and economics had recast humans as “imperfectly rational.”⁹ Thus, while officials cast the public’s wariness about technology on subjects from automation to atomic weapons as emotional, ignorant, and irrational,¹⁰ researchers within the social sciences turned increasingly to a view that rationality was not an inviolable trait, but a product of structures and systems—operations research, management science, and other tools of the modern liberal state.¹¹

All electronic failures, one observer suggested, could be traced back to the fallibility of people—the operators, maintenance personnel, and designers of electronic equipment. The detailed failings of each of these groups is a story too complicated to discuss here; cutting across them, however, was the idea that unreliability and failures in electronics arose from a series of disjunctures between the functioning of electronics and the capacities of humans—the slowness of our reactions; the incompatibility of our senses; and, most disturbing for our purposes, the divide between the often

paradigmatically rational functioning of electronic devices and the subtly flawed human reason that tried to make sense of them through drawings.

In the same publication where Bainbridge-Bell held forth on “table manners,” a fellow wartime radar researcher likened circuit diagrams to the system of Arabic numerals, calling them “a practically indispensable aid to thought; simple and effective.” So effective, in fact, that they “created the danger of handing over to them too much of our reasoning powers.”¹² In electronics, those concerns transformed the material and visual culture of electronics in the 1950s in ways that mirrored the drive in social science toward rationality as a product of arrangements and systems. The concerns over maintenance and repair, for example, led to systems of color-coded circuits, functional groupings of electronic components, marked signal paths, and the introduction of teaching machines designed to make technicians think in terms of the algorithms of the machines they maintained.

The concerns over drawings temporarily transformed the visual culture of the electronic arts as well. To realign human rationality with the actual functioning of machines, circuit designers (particularly those involved in teaching electronics during the war) proposed what they openly described as new “systems” of drawing practices that would guide the reasoning powers of circuit designers through the two-tone scheme of the diagrams.

Greatest Concern over Transistors

One of the great foci of those concerns over drawings was the preferred way to symbolize operational amplifiers. The operational amplifier (or op-amp) is an abstract electronic device, a black box, described in terms of the actions it performs on its electrical inputs. Originally made with vacuum tubes, op-amps are now more commonly implemented with integrated circuits or transistors. Whatever the case, an op-amp is represented using a single symbol with two inputs, one of which is inverted, and one output. The op-amp operates based on the difference between the inputs. If the inputs are identical, no output is produced. If the inputs are different, the output multiplies that difference, producing a gain (it is an amplifier, after all). Although an ideal amplifier (one that cannot be physically realized) has infinite gain, real op-amps have a large but finite gain. Moreover, feeding some of the output of an op-amp back into one or both of its inputs allows one to generate an astounding variety of other useful electronic devices: circuits that add, subtract, invert, multiply, take logarithms, compare values, and compute derivatives and integrals. As such, op-amps serve as a uni-

versal building block for (analog) computers and other signal-processing devices, and thus are ubiquitous in analog electronic design.

Hiding inside the black box of the op-amp symbol was a fierce debate about how transistors should be drawn. Between 1947 and 1957, an explosion of competing schemes for drawing transistors swept across journals, blackboards, and lab notebooks. Frustrated observers pleaded for a single “symbol language” in place of the emerging dystopia of organization-specific symbols that distinguished circuit diagrams at Bell Labs, for instance, from those at RCA.¹³ The proliferation was particularly driven by the way that a new version of the transistor introduced in 1951—the junction transistor—broke with both the idioms and the ontology of electronic diagrams, entities that could be said to exist in the world of the drawings.¹⁴ In the interests of space, we are putting aside the details of those dilemmas to simply say that, in the face of these breaks, proposals for a new transistor symbol split along two principal lines that highlighted the historical tensions in circuit diagrams. One approach embraced the new ontology and emphasized questions of structure and physical form, particularly the alternating materials that made up the device. The other approach rejected questions of form in favor of symbols that gestured toward function; within that emphasis on function, proponents split once again into those who, thinking about the transistor as an analog of the vacuum tube (or thermionic valve), sought to invoke those similarities in drawings, and those who believed such visual and conceptual practices were dangerously misleading.

That debate, often heated, played out in the pages of leading journals in the late 1950s, where the thermionic emphasis sought to draw visual genealogies between valves and semiconductors. Marcus Scroggie, who had talked ominously about the power of diagrams, led that emphasis. In an unrelenting series of examples, Scroggie explicitly translated semiconductor behavior and circuits into their vacuum-tube equivalents. Current flow in diodes moved from disembodied charges to the interchangeable forms of vacuum tubes and their crystal analogues; a discussion of basic transistor arrangements erected solid-state circuits on a foundation of thermionic counterparts. “There is, I believe, a school of thought that deprecates likening a transistor to a valve,” Scroggie explained. “Personally I hold that transistors have so much in common with valves, as regards function, methods of use, and to some extent internal workings, that it is futile not to note the similarities.”¹⁵

Scroggie’s personal choice for a transistor symbol was the one used at RCA (see figure 6.4). The two symbols on the left represent the dif-

Figure 6.4. Scroggie's transistor symbol used at RCA.

ferent polarities of transistor; the symbol on the right is the standard vacuum-tube icon. Its visual similarity to the vacuum tube symbol, Scroggie claimed, would ground the work of circuit designers in the familiar, avoiding the perplexing solid-state physics involved in transistor action, thus clarifying transistor function. At a conference on nuclear instrumentation in 1957, the Swiss physicist Ernst Baldinger echoed this functional similarity through the powerful equivalent diagrams that had first begun to dominate circuit design in the 1930s¹⁶ (see figure 6.5). Again and again, Baldinger and others drew their lines of relation, their “histories” of electronic components, visually from transistors to valves and back again. To judge from their drawings, nothing functioned so much like a transistor as a vacuum tube did.

As Scroggie acknowledged, though, others thought differently. At the same conference where Baldinger produced his equivalent circuits, Norman Moody was deeply worried about the reliability of transistors in nuclear instrumentation. He argued that those very connections caused circuit designers to misunderstand what transistors were, and that the practice of drawing visual and therefore functional analogies between the two devices misled engineers, causing them to reason (falsely) that transistors were a straightforward replacement for vacuum tubes. Moody himself had helped develop instrumentation for Canada's nuclear laboratories after the war; he had worked on British nuclear weapons, designing high-speed circuits to

Figure 6.5. Equivalent circuit of a triode and of a junction transistor for low and medium frequencies. Typical values for a low-frequency transistor (type OC 71; collector current 2 mA) are indicated. Ernst Baldinger, “Transistors in Nuclear Instruments,” *Nuclear Instruments* 2 (1958): 193–202: 194. Amsterdam: North-Holland Pub. Co. (Courtesy of Elsevier.)

measure gamma ray fluxes from atomic bombs before the electronics were destroyed by bombarding neutrons. In that work, he focused on the special properties of transistors that would make their circuits especially robust under punishing conditions.

In no area were engineers more prone to error, nowhere could their faculties be more easily misled, Moody reasoned, than in moving to transistors from vacuum tubes. For many technologists the transistors and vacuum tubes shared a critical functional similarity—they amplified power. For Moody, the similarity was trivial, but its implications were dire. The transistor, he explained, was “different from the vacuum tube in all important respects. . . . you must learn an entirely new circuit technique to use it satisfactorily. Even then, you must avoid a tube-by-tube replacement by transistors. You must re-design on a functional basis.”¹⁷ He drove his point home with his own equivalent circuit, completely incompatible with Baldinger’s, that likened the transistor to a diode. The new design practices it required would take at least a year to master and the current ways of drawing transistors were an obstacle to acquiring them.

Visually evoking the similarity of the two devices, symbols such as Scroggie’s effaced the crucial features of transistors that set them apart

from tubes—properties that were critical to the sensitive pulse circuitry at the heart of both Moody’s research and Cold War electronics. “To designers who are familiar with vacuum tubes,” Moody explained, “this trivial resemblance [of power amplification] . . . has been a trap. Many of their failures are strewn along the path of progress.”¹⁸ In this remarkable statement, Moody proclaimed that the first step to making transistors reliable was to draw them differently, prompting him to reveal a new transistor symbol that he used at a conference earlier that year, where the audience complained “with intensity” about the lack of a standard representation.¹⁹ For Moody, however, the new icon formed part of an expandable iconographic system—a “graphical nomenclature”—that rationalized the drawing of circuits by evoking visually the place of the devices in the larger universe of electronic parts and operations.

Compared to Scroggie and Baldinger, Moody’s genealogies ran orthogonally—not from disembodied electron flows to vacuum tubes and transistors, but rather from elemental solid-state devices to more evolved forms.²⁰ They captured visually the logical and material connections of electronics, thereby guiding the thinking of engineers as they traced out their designs.

These were only two proposals for drawing transistors in the late 1950s, each with its own set of anxieties and understanding of electronic devices, but how many more lie buried in the massive volumes of electronic literature from those years? By the early 1960s, competing systems had largely disappeared from public view, swept away by a standardization that would ultimately make circuit diagrams even more authoritative. Yet being able to recapture those images and the debates that surrounded them would reveal the continuities and ruptures, the pedagogical and technical communities, the practices of work and production that surrounded electronics precisely at a time when the stakes for them were immense, and when those same practices were being applied to systems and phenomena far beyond the electronic arts. In other words, recapturing those images could help us understand how the practices of electronics that underwrote everything from atomic weapons to video games were bound up in the larger histories of trust during the Cold War.

Gathering the Sources

Cataloging this variation in how electronics were drawn throughout the twentieth century sheds light on the wider sets of practices and preoc-

cupations that surrounded them, and on their creators. The vast published literature on electronics provides one particularly fruitful path into these questions. Emerging from late-nineteenth-century work on electrical physics and growing exponentially since, thanks in large part to Moore's Law,²¹ electronics-oriented literature is particularly attractive for computer analysis because substantial numbers of sources about electronics are (1) already digitized or born-digital; (2) typeset, facilitating optical character recognition (OCR), a familiar kind of image processing; and (3) freely available and readily findable online, which makes automated downloading straightforward. Furthermore, those involved in developing the internet and laying the foundations for the World Wide Web came from and participated in a culture of extensive technical documentation in digital form, and they often collected older technological artifacts and their associated documentation as a hobbyist practice. Yet while this literature may be widely and readily available, it also presents interesting challenges for text- and image-mining.

Practices of technical drawing are an integral part of the practice of electronics writ large: it is rare to find a source on computing that does not contain some images, and almost impossible to find such a source on electronics. Taking a computational approach to electronics literature thus requires the (semi)-automated handling of photographs, schematics, exploded mechanisms, graphs, block diagrams, timing diagrams, network diagrams, drawings of atomic structure, charts, equations, and hundreds of other examples of visual culture—each of which appeared in a variety of competing forms and changed substantially over time, as we will discuss shortly.

In addition, although the visual idioms of electronics and computing are often unchanged over decades and even centuries (in the case of the battery), the physical infrastructure they refer to has changed, sometimes dramatically, over time. The operational amplifier, for example—an essential building block for electronic instrumentation, analog computers, and signal processing—was implemented using vacuum tubes in the 1940s, transistors in the 1950s, and integrated circuits in the 1970s. As another example, computer memory can be implemented with cathode-ray tubes, relays, magnetic cores, magnetic tape, plated wire, thin film, optical disks, and many other technologies, even though its basic function remains the same.

In light of these challenges, our first step has been to assemble a database for this visual culture—a multiterabyte-scale archive of technical

literature—by spidering it from the open web. In some cases, this has been very easy. The Internet Archive, for example, hosts entire runs of important journals (e.g., *Bell System Technical Journal*, 1922–1983),²² collections assembled by other researchers (e.g., the Arpanet archive compiled by Katie Hafner),²³ hobbyist collections (e.g., the BITSАVERS.ORG collection of over three million pages of computer-related documentation),²⁴ and books (e.g., the Folkscanomy Computer and Folkscanomy Electronics collections contain over 7,000 titles).²⁵ Furthermore, the Internet Archive supports mass downloading of their collections and is even happy to assist researchers who are not able to figure out how to do it on their own.²⁶

Another important source of images has been electronics manufacturers. Companies archive white papers, data sheets, application notes, schematics, and a wide assortment of other documents about their own products, with the expectation that these materials will often be downloaded in bulk. Some companies also provide free access to journals produced in-house, such as the Analog Dialogue (1967–present) of Analog Devices, Inc.²⁷ The terms and conditions of some of the large online electronics retailers allow their websites to be mirrored or spidered for personal, noncommercial use.

Many online databases of data sheets can also be utilized. Some of these contain information that is primarily of historical interest. Frank's Electron Tube Pages, for example, hosts more than 13,000 data sheets for electron tubes dating from the early- to mid-twentieth century.²⁸ More sophisticated web-crawling techniques turn up large collections of preprints, open-access literature, out-of-print sources, digitized archival holdings, photographs, and so forth.

Since large-scale text- and image-mining are still unfamiliar to most historians, it is worth saying something about what it is like to use systems such as these. Prior to the widespread availability of full-text keyword searching in digital sources, finding potentially interesting documents was a slow and laborious process. General-purpose search engines now throw up thousands or millions of potential matches in a fraction of a second, but the process of separating out valuable results can still be tedious. When one searches in a large-scale, machine-curated collection, results tend to be much more topical: the signal-to-noise ratio is a lot better. To take an example from the history of electronics, an engineer named George Wilson invented an important transistor-based circuit now called the Wilson current mirror. The name George Wilson is relatively common, and searches on the web tend to turn up fictional characters, politicians, and sports figures. Searches in a large database of sources related to the history of elec-

tronics, however, return documents that are overwhelmingly relevant. The day-to-day use of such technologies gives one a feeling of almost continual serendipity.

The process of collecting and managing these sources is ongoing. Since many of these materials are provided online for personal, noncommercial use only, we do not intend to make them freely available (and legally we cannot), but we are always happy to talk to potential research collaborators.

Image-Mining Workflow

Some of our spidering is accomplished with Bash scripts and off-the-shelf tools, but the majority of our computational work is done in Mathematica.²⁹ Given a set of digital sources, our workflow is as follows.

Burst documents. We start by splitting multipage documents into individual pages, each identified by its source and position in the original. In addition to providing manageable chunks for text and image processing, this also makes fine-grained searching possible. In general, we want to know that a particular book not only has some information that we are interested in, but we want to see the page(s) where that information appears.³⁰

Text mining. We extract text and subject it to standard techniques such as indexing, frequency analysis, TF/IDF, named entity recognition, and so on. Indexing provides a map from a given term (or n-gram) to all of the locations where it appears in the corpus. Frequency analysis determines how often a particular term (or n-gram) appears. TF/IDF is a measure commonly used in information retrieval; the terms that are relatively rare in the corpus overall, but frequent in a particular document, are the ones that best represent what that document is about. Named entity recognition is the process of identifying the people, places, organizations, and other elements that the sources mention. These text-based techniques, and others, can be used for a range of familiar tasks in the digital humanities, including keyword searching, topic modeling, and unsupervised clustering.

Automatic image extraction. Using image-processing techniques applied to document pages, we identify regions that are probably images and automatically extract those to new files. This process (described in more detail by Elliott and Turkel in this volume) provides a clear visualization with a thumbnail of the page in the first column, a second column that shows orange boxes around the images that have been automatically iden-

Figure 6.6. A 1948 schematic diagram after processing with computer vision algorithms.

tified, and a third column showing the last image that has been automatically extracted from each page. This technique works surprisingly well.

Automatic image classification. Next, we use machine learning techniques to “teach” the computer to distinguish between different kinds of images. We provide learners with a small training set of images that have been labeled by a person and then ask the learners to provide their own labels for images that were not in the training set. Learners were able to differentiate schematics, block diagrams, textboxes, drawings, and photographs with surprising precision.

Schematic understanding. The next step is to use computer vision to try to break schematics down into smaller components (e.g., transistor) and/or design idioms (e.g., current mirror). The methods we use here depend on the task at hand, and the problems that result are much more difficult to solve. We have just started working on the task of parsing component images out of complete schematics. Figure 6.6, for example, shows a

schematic diagram's automatic parsing into components, represented with different colors. It is obviously not perfect yet, but it is a start.

Further processing. Under some conditions it may be possible to parse schematics in such a way as to automatically create a simulation of the circuit using SPICE.³¹ At this point, we do not foresee the need to go quite that far, but it is an interesting possibility for future systems.

Taken together, these techniques can help recapture the rich visual, material, and practical worlds of twentieth-century electronics. As the early history of electronics and the debates over schematics demonstrate, illuminating that history means bringing to light the pervasive social and cultural role of electronics, not just as a collection of material devices, but as a source of powerful metaphors and flexible practices for organizing and making sense of the world. And recovering the legacies of those historical practices becomes increasingly important for understanding our own historical moment.

War and Historiography in the Age of Seeing Machines

We conclude with a brief discussion of contemporary features that both allow (and even require) the practice of this kind of research, and that raise interesting parallels with our Cold War case study. Techniques of image-mining—and attendant technologies of digital photography, image processing, computer vision, and machine learning on visual representations—have expanded rapidly in the post-9/11 world, in part due to the relentless growth captured by Moore's Law. Take streaming video. The first commercial products were introduced around 1995. YouTube was founded a decade later, and within a couple of years accounted for 10 percent of all internet traffic. Video streaming currently accounts for almost 80 percent of US internet traffic, and 100 hours of video are uploaded to YouTube every minute. The first commercial digital camera was introduced in 1990, but digital cameras did not outsell film cameras for at least a decade. The vast majority of photography is now digital, and done primarily with smartphones. Every few days, a billion images are uploaded to Facebook alone.

Research in computing with images is driven in part by consumer applications, in part by more exotic projects that may soon appeal to consumers (self-driving cars come to mind), and in part by less social or even antisocial goals. A camera built into an advertising display can track the eye movements of people as they look at the ad. Marketing firms use image-mining on “selfies” to identify brand logos on clothing and target potential

consumers. For security applications, machines are trained to analyze faces, facial expressions, gaze, gait, fingerprints, and many other biometric signals to identify people, and to answer questions ranging from whether people appear to be sick to whether they have a bomb strapped on under their clothing. A wide range of cameras are deployed specifically for surveillance, in automated teller machines, on street corners and highways and ports, in satellites, in drones, and in a million other places . . . and to these can be added the possibility of surreptitiously hijacking cameras built into computers and smartphones.

In our world, images and seeing machines play a central role in the creation of new modes of trustworthiness and reliability and give rise to many new anxieties. As always, knowing something about the past can help us make better decisions in the present. Our view of the past, however, will be increasingly dominated by the exponential growth of still and moving images in digital form, and making historical sense of these will require that we come to terms with machines that can see.

Notes

1. Ronald Sullivan, "William A. Higginbotham, 84; Helped Build First Atomic Bomb," *New York Times*, November 15, 1994.
2. Edward Jones-Imhotep, "Icons and Electronics," *Historical Studies in the Natural Sciences* 38, no. 3 (2008): 405–50.
3. William J. Turkel, *Spark from the Deep: How Shocking Experiments with Strongly Electric Fish Powered Scientific Discovery* (Baltimore: Johns Hopkins University Press, 2013).
4. As Eccles would explain in 1929, the advanced state of analytical studies of electricity and "electrical vibrations" in the early twentieth century furnished a ready analytical apparatus for other disciplines in which vibrations (or any other phenomena describable by differential equations) were seen to be central, such as acoustics. This was done with one-to-one correlations between electrical variables such as voltage, current, and charge on one hand and force, speed, and displacement in the case of things such as sound. See *Osiris* 28: *Music, Sound, and the Laboratory from 1750 to 1980* <http://www.ece.rice.edu/~dhj/paper1.pdf>; and D. H. Johnson, "Origins of the Equivalent Circuit Concept: The Current-Source Equivalent," *Proceedings of the IEEE* 91 (2003): 817–21, <http://www.ece.rice.edu/~dhj/paper2.pdf>. On Steinmetz, see Ronald Kline, *Steinmetz: Engineer and Socialist* (Baltimore: Johns Hopkins University Press, 1992), 108–12.
5. Carl Steinmetz, for instance, was one of the first to use graphical methods and complex numbers in heavy-current engineering, and routinely used analogies between mechanics and electricity. See Don H. Johnson, "Origins of the Equivalent Circuit Concept: The Voltage-Source Equivalent," 91 (2003): 636–40, See

for example the work of Hahnemann and Hecht in Wittje, 56. Wittje, notes that “By the 1930s electric circuit diagrams, even in the case of representations of non-electric systems, had become a lingua franca of the new acoustics.” Trendelenburg referred to two scientists from Bell Laboratories, Irving Crandall and Raymond Lester Wegel, the latter of whom had discussed representing the vibration of the vocal cords by an electric equivalent circuit.

6. Here we borrow the term “paper tool” from the work of Ursula Klein and David Kaiser. See U. Klein, “Techniques of Modeling and Paper-Tools in Classical Chemistry,” in *Models as Mediators: Perspectives on Natural and Social Science*, eds. Mary Morgan and Margaret Morrison (New York: Cambridge University Press, 1999), 146–67; U. Klein, “Paper Tools in Experimental Cultures,” *Studies in the History and Philosophy of Science* 32 (2001): 265–302; U. Klein, *Experiments, Models and Paper Tools: Cultures of Organic Chemistry in the Nineteenth Century* (Stanford, CA: Stanford University Press, 2003); and D. Kaiser, *Drawing Theories Apart* (Chicago: University of Chicago Press, 2005), chapter 1.

7. L. H. Bainbridge-Bell, “Circuit Symbols: Notes on the New British Standard,” *Wireless World* 54 (1948), 437.

8. Those practices extended to the material circuits. Frederic Williams, for example, the English engineer reputed to have developed the first operational amplifier and, with the physicist P. M. S. Blackett, an automatic curve follower, stressed the elegance, efficiency, and reliability of his circuits, calling well-behaved circuits “sanitary.” On manners and morals, Sarah Buss has explained that although most philosophers treat the two separately—what does moral duty have to do with politeness?—most people draw no sharp distinction between them. See Sarah Buss, “Appearing Respectful: The Moral Significance of Manners,” *Ethics* 109, no. 4 (1999): 795–826.

9. Hunter Heyck, “Producing Reason,” in *Cold War Social Science: Knowledge Production, Liberal Democracy, and Human Nature*, eds. Mark Solovey and Hamilton Cravens (New York: Palgrave Macmillan, 2012), 102.

10. Jacob Darwin Hamblin, “Exorcising Ghosts in the Age of Automation: United Nations Experts and Atoms for Peace,” *Technology and Culture* 47, no. 4 (2006): 734–56.

11. Heyck “Producing Reason,” 100.

12. Cathode Ray, “Invisible Components: What The Circuit Diagram Fails to Show,” *Wireless World* 54 (1948), 459.

13. Allen F. Pomeroy, “Universal Engineering Shorthand,” *Proceedings of the IRE* 40 (1952): 771.

14. Jones-Imhotep.

15. Marcus Scroggie [Cathode Ray], “Transistor Graphic Symbols: A Critical Analysis of Existing Ideas and Conventions,” *Wireless World* 63 (1957), 194–95.

16. Ernst Baldinger, “Transistors in Nuclear Instruments,” *Nuclear Instruments* 2, no. 2 (1958): 193–202.

17. N. F. Moody, “The Present State of the Transistor and Its Associated Circuit Art,” *Nuclear Instruments* 2, no. 2 (1958): 182–92.

18. *Ibid.*, 182.

19. Howard E. Tompkins, "Foreword to Transistor Papers," *IRE Transactions on Circuit Theory*, CT-4 (1957), 173.
20. Philip M. Thompson and Jack Bateson, "Semiconductor Symbols: Logical System for Diodes, Transistors and Other Junction Devices," *Wireless World* 63 (1957), p. 525.
21. Paul E. Ceruzzi, "Moore's Law and Technological Determinism: Reflections on the History of Technology," *Technology and Culture* 46, no. 3 (2005): 584–93, <http://plato.acadiau.ca/courses/comp/dsilver/2903/Readings/A2-2010%20-%20Moore's%20Law.html>
22. *Bell System Technical Journal* [archive, 1922–1983], <https://archive.org/details/bstj-archives>
23. Arpanet [archive], <https://archive.org/details/arpanet>
24. BITSAVERS.ORG Documents Library [archive], <https://archive.org/details/bitsavers>
25. Folkscanomy Computer [archive], https://archive.org/details/folkscanomy_computer; Folkscanomy Electronics [archive], https://archive.org/details/folkscanomy_electronics
26. Internet Archive Blogs, "Downloading in Bulk Using wget," April 26, 2012, <http://blog.archive.org/2012/04/26/downloading-in-bulk-using-wget/>
27. Analog Dialogue [archive, 1967–2014], <http://www.analog.com/library/analogdialogue/archives.html>
28. Frank's Electron Tube Pages [archive], <http://www.tubedata.org/> and <http://tubedata.milbert.com/index.html>
29. <http://www.wolfram.com/mathematica/>
30. See W. J. Turkel, "Burst Documents for Fine-Grained Matching," March 27, 2011, <http://williamjturkel.net/2011/03/27/burst-documents/>
31. A proof-of-principle example is given in Donald Bailey, Andrew Norman, and Giovanni Moretti, "Electronic Schematic Recognition," *Proceedings of the 4th Electronics New Zealand Conference*, ENZCon '97, Massey University, Auckland, New Zealand (September 1997), http://sprg.massey.ac.nz/pdfs/1997_ENZCON_147.pdf. For SPICE see <http://bwrcs.eecs.berkeley.edu/Classes/IcBook/SPICE/>

CHAPTER 7

Learning to See the Past at Scale

Exploring Web Archives through Hundreds of Thousands of Images

Ian Milligan

What can historians do with the millions of images found on the Internet? Can we make sense of the seemingly infinite sources? In this chapter, Ian Milligan presents his method of exploring early online communities and top-level domains, and his search for cultural meaning among hundreds of thousands of images. The countless portraits, pictures, and photographs accessible online are far from being a jumbled mass of chaos; Milligan explains how to find patterns by combining distant reading techniques with image-mining software.

Since its advent in 1990–91, the World Wide Web has been generating billions of pages of data, reshaping society in general and the historical profession in ways we are only now beginning to grasp. Web archives containing everything from personal home pages to professional or academic websites offer the ability to reconstruct large-scale swathes of the recent past. Yet the sheer number of these primary sources presents significant challenges—a trend that accelerates as we continue to preserve data at an ever-increasing rate. Due to shortcomings of technical ability, interest, and awareness, the scholarly use of web archives remains quite low, and work to date has largely focused on either text or link/network analysis.¹ But there are ways we can use web archives, especially graphical images that are increasingly available, to “see” the past on a previously impossible scale.

Imagine an historian staring up at a mountain of these pictures and wondering what on earth she can do about it. Some historians might just begin sifting through them one by one, as if they had been found in a large archival box. Yet rather than carrying out a close reading of individual websites, I am interested in what images can tell us about these archives as systems. Therefore, I have chosen to use computer vision to examine two data sets—9,605,749 images extracted from the now-defunct GeoCities, an easy-to-use web platform and online community of the 1990s, and 4,346,682 images the Internet Archive scraped from the Web in 2011—to see how historians and archivists might be able to make sense of such extreme data sets.

I approach these large collections with several questions in mind as examples of a way forward, demonstrating the potential that large image collections present for the study of history. First, by taking the “neighborhoods” that comprise GeoCities and the national top-level domains that comprise the Wide Web Scrape, I ask whether we can use images to answer questions about community and consistency. How are images used, and how can we use computational methods to see patterns in these large data sets? Second, by comparing data largely from the late 1990s between an array of top-level domains, can we track some broad changes in the World Wide Web? If, for example, we can take wide web scrapes, visualize the interconnections between various clusters of images, explore montages of tens of thousands of images, and discover and explore meaningful and important recurring images, we might be able to compare and contrast top-level domains so that historians and archivists of specific subdomains—say the Canadian .ca domain—can see what might have made this section of the World Wide Web different throughout a certain time period. We can similarly do that with mid-1990s-era GeoCities archives.

Because one cannot read every website in these collections oneself, analyzing these image collections as a whole gives us new insights, complementing research findings garnered from distant textual analysis and occasional close reading. This approach is not better than normal methodologies, but complementary: computer vision is augmenting the historian’s craft, and in particular source abundance, providing a way to explore the visual content of digitized archives.²

A central concept in approaching these images is that of “distant reading,” which I have discussed in relation to historical scholarship in the *Journal of the Canadian Historical Association*.³ This idea comes from English literature scholar Franco Moretti, who used it to help understand the rise

of the Victorian novel. To understand these novels as a genre, traditional literary scholars were basing their studies on a traditional corpus of approximately 200 books. Though this corpus offered an impressive number of books, to be sure, it was only 1 percent or so of the total output of novels during that period, making conventional models of analysis impractical. Moretti argued that a different approach was necessary:

[C]lose reading won't help here, a novel a day every day of the year would take a century or so. . . . And it's not even a matter of time, but of method: a field this large cannot be understood by stitching together separate bit[s?] of knowledge about individual cases, because it isn't a sum of individual cases: it's a collective system that should be grasped as such, as a whole.⁴

Historians might consider and approach images in web archives in a similar way. This does not mitigate the need to actually closely read a few hundred images, or to look at individual websites themselves; rather, it suggests a new way to approach historical research.

Images can be contextualized, and taking them in their context—such as a GeoCities neighborhood or top-level domain—to examine their similarities, overall structure, and contours is one way into these immense born-digital collections of data. Suddenly, nine million images can look a bit less like an insurmountable mountain and more like an opportunity.⁵

The Two Collections: The Top-Level Domains of the Wide Web Scrape and the Neighborhoods of Geocities

As noted, I use two major collections. One draws from the end-of-life holdings of the popular GeoCities.com website-hosting service, with websites dating between 1996 and 2009. The second uses the collections of the Internet Archive and amounts to a snapshot of web activity from 2011.

The first data set was GeoCities, which by the early 2000s had thirty-eight million websites.⁶ Opening in May 1995, GeoCities became an overnight sensation because it allowed users to create their own web pages in an easy-to-use, intuitive, and—most importantly—free environment (it had similar free competitors, including tripod.com and Angelfire). What arguably set GeoCities apart from other website providers, however, was its explicit emphasis on forming communities through familiar space- and placed-based metaphors and rhetoric. It was presented as a “cityscape,”

with streets, street numbers, and recognizable urban and geographical landmarks (for instance, one might live on a virtual Fifth Avenue or on a festive Bourbon Street), and this setup was emphasized in company press releases and user communications. Users were “homesteaders,” who settled various parts of the web space, and spatial metaphors were employed.

When users arrived to create their site, they were presented with a list of neighborhoods where their site might belong. For example, pages on “[e]ducation, literature, poetry, philosophy” would be encouraged to join the Athens neighborhoods, those interest in politics could set up shop in the Capitol Hill, small businesspeople in the aptly named Eureka area, and so on. Some neighborhoods came with overt restrictions and guidance, such as the more protective and censored Enchanted Forest for children. Each neighborhood had its own apparatus of support (much like its real-world equivalent): community leaders, coding guidelines, web rings, standards, and so forth. Content standards were maintained by the aptly named “Neighborhood Watch,” centrally managed by the GeoCities administration: “[i]f you notice any of your neighbors not following our policies, please let us know,” they asked volunteers.⁷

Neighborhood membership was mandatory. If you wanted to “homestead” in GeoCities, your first step was to review the list of neighborhoods and their small descriptive blurbs.⁸ The largest was Heartland. The smallest were more specific, such as the advertising-focused Madison Avenue or the car-focused Motor City. Each area had a home page, maintained by GeoCities itself. As with any real estate transaction, a potential homesteader needed to see if they were a good fit for the area. Each neighborhood invited users to “join this neighborhood,” where they would be given a longer list of appropriate topics. In Heartland, for example, this included pages about families, pets, hometowns, genealogical research, organization, and other local events. Popular neighborhoods began to fill up quickly, necessitating a move to the suburbs, such as Heartland/Plains and Heartland/Hills.

While it was easy to join GeoCities, there was a lively scholarly debate about whether it amounted to anything more than just a place to park a website.⁹ My work on GeoCities explores whether it amounted to a community or not.¹⁰ Communities need to be enacted, received, and perceived by members, and certainly the Geocities website aimed to create community by allowing members to bypass conventional online access providers—most notably America Online and CompuServe—and form a unique neighborhood system of users simply chatting with each other, even

as it offered users the familiarity of space-based metaphors to perhaps keep them anchored. But is this enough to be considered a community? In part by examining thousands of images, we can see that for a substantial minority of users, GeoCities did indeed offer a meaningful virtual community.

My second data set comes from the Internet Archive. On October 26, 2012, the Internet Archive announced that it would make an entire crawl of the World Wide Web, 80 terabytes of information, available to interested researchers. This Wide Web Scrape was released in its raw form: 85,570 one-gigabyte WebARChive (WARC) files, an ISO-standardized format that combines all of the information necessary to reproduce a website and puts it into one big package. While WARC files can normally be navigated by the Wayback Machine, which lets users see a website as it would have appeared in the past, this system limits the user to seeing only one page at a time, and I wanted to distantly read millions of pages simultaneously.

To get around this obstacle, I combined a sample of twelve top-level domains (TLDs) to use as a methodological starting point. Three were generic: .com, .org, and .net. The rest were geographically specific: United States education (.edu), United States military (.mil), United States government (.gov), as well as Canada (.ca), the United Kingdom (.uk), France (.fr), China (.cn), India (.in), and South Korea (.ke). I selected these domains for a variety of reasons: .ca was the primary focus of my research funded by the Social Sciences and Humanities Research Council of Canada, and I selected the others to provide comparative data: the generic domains (.com, .org, .net) to see how the .ca TLD differed from these “norms” within the web; the American education, military, and government domains to explore distinctive, topic-circumscribed corpuses; and the other geographic domains to see how each compared to relatively large and active TLDs with potentially different approaches to web development. I do not claim to have compiled a representative sample of the World Wide Web, but rather, a comparative data set that allows us some vision into different sets of the 2011 Wide Web Scrape.

To give a sense of the contours of this data set, I collected 29,219,706 .com addresses (2.32 percent of the total of 1,260,409,874); 10,268 .mil addresses (0.92 percent), and 622,365 .ca addresses (7.31 percent), resulting in a total of 4,346,682 extracted images. Each of those addresses might point to multiple things: HTML pages themselves, text files, PDFs, or images. For the analysis that follows, I do not seek to be exhaustive or definitive about the contents of these TLDs but, rather, point toward various

comparative metrics that we can use and the historical value we can glean from these sources. Furthermore, while the different languages pose difficulty for textual analysis, this is not a present problem with the image analysis I am conducting.

For both data sets, I extracted the images in a fairly simple manner. I decompressed the files and then used a script to extract all images, sequentially number them, and keep them in directories specific to where they came from (so all .ca images were kept together, for example, as were all Heartland or EnchantedForest images). To find files, I relied on the work of other digital preservationists. At the UK Web Archive, Andrew Jackson examined a corpus of some 2.5 billion files from the UK web domain between 1996 and 2010, looking at the usage of common image formats and how they changed over time.¹¹ I used this as a seed for my own data analysis, looking for images with the extensions .gif, .jpg/.jpeg, .tif/.tiff, .png, .jp2, .j2k, .bmp, .pict, .wmf, .emf, .ico, and .xbm. The results were rewarding.

Visualizing the Past: The Montage as First Stop

While at times it was fun to sit back and just look at the pictures one by one, these were extremely large data sets. This prompted me to scale up my methods to extract meaning from so much information. A useful initial stage before examining metadata and wholesale computer vision was to pull back our gaze from individual images and look across tens of thousands of them simultaneously, and the easiest way to “distantly read” images, if you still want to have them on your screen—albeit in much-reduced size—is the montage.

Lev Manovich and colleagues at the Software Studies Initiative have done interesting work with large data sets of images. In one study, for instance, they visualize one in every three pages of *Popular Science* to see how the publication has gradually evolved over 125 years, moving toward more pictorial content and different page layouts, for example.¹² They also have taken *Time* magazine covers and detected different color hues based on who was in political power.¹³ While they have robust longitudinal data, their basic methodology gives us a way to learn about the past. Even more promising, they provide online instructions to easily show users how to do this themselves with their own collections of digital images. Even so, we need to use these techniques with caution: the arrangement of images in a montage does not relate to their significance, especially in the case of web

Figure 7.1. A montage of 119,080 images in the .ca TLD.

archives that do not have a longitudinal element. Unfortunately, viewers tend to privilege up-down relationships over left-right relationships, even if they are identical.¹⁴

To illustrate, I will take examples from my data sets. Only a few will be presented in this book, because they do not lend themselves well to print reproduction. Figure 7.1 displays sample of images from the .ca top-level domain. (A full-color version is available at <http://ianmilli.files.wordpress.com/2014/01/50kmontage.jpg>). Similar to the procedure in programs such as Preview, Photoshop, or Windows Photo Viewer, users can rapidly move between the levels of the overall top-level domain (or the 7 percent I am using) down to the individual images themselves, simply by zooming in and

out. In looking at such montages, we can make a few quick observations about the shape of this domain. First, there is a lot of digital photography, and within these images, we see a remarkable number of faces (facial recognition algorithms note that almost a quarter of the pictures have at least one human face). Second, there is remarkable coherence in batches of images (e.g., lines of yellow images or white images), suggesting a large website here or there with visual consistency: these are often catalogs, e-commerce portals, or large institutions with brand guidelines (universities, for example).

To get a sense of the method's utility, I compared this montage to those created with the images of two other top-level domains: .mil and .cn.¹⁵ The .mil montage offers a rather unique confluence of images. First, there are lots of black-and-white images depicting American history, often showing great battles and military events. Second, there are lots of headshots of military officials, suggesting that we have numerous personnel sites to consider. Such findings can be useful in other work. For example, when georeferencing locations mentioned in the .mil archive, I saw lots of locations in northern France, in addition to Afghanistan and Iraq; being able to combine this montage with keyword-in-context and geocoding allowed me to garner a sense of the archive's contents.

The .cn (China) top-level domain presents a very different and much larger data set than the previous ones. The appearance of this montage also includes splashes of color, to be sure, but there are a *lot* of white backgrounds (patches of white appear prominently). I will discuss this feature more in the next section, but in general, this montage suggests we are not seeing digital photographs. In fact, zooming in and out demonstrates that yes indeed, we are looking at multitudes of consumer products, clip art, and other sorts of characters. Perhaps with more tranches of web data, we will be able to see the spread of digital photography and the relative lesser prominence of e-commerce on a national domain.

Building on the patterns we can see in domain-oriented montages, we can use montages to learn basic elements about GeoCities neighborhoods. Consider the child-focused EnchantedForest neighborhood. It's a useful test case because, due to its nature, it had the most circumscribed and vigilant community moderators: the sites had to be by or for children to find a home there. In figure 7.2, I include the montage of the EnchantedForest as well as a cutaway zoom to see what it shows in particular. Looking at this montage we certainly can tell that, compared to the current web, there are *fewer* pictures of people. There are more examples of clip art and more

Figure 7.2. The EnchantedForest, with cutaway zoom.

low-resolution images and navigational elements. Even from afar, without the help of any further digital aids, we can see a bit less color variation: a lot of white backgrounds, for example, and consistent colors. When zooming in we find a lot of children's characters, cartoons, color, pictures of pets, pictures of babies, and so forth. The montage indicates that GeoCities' neighborhood model generally *worked* (my other research in this field demonstrates this, both through image and textual analysis). The Enchanted-Forest was for children, Heartland was for families, and so forth.

Montages have their advantages and disadvantages. They are easy to create once you have the images, they can be operated with a minimum of technical skill, and they encapsulate the bridge between distant and close

reading that generates so much fruitful engagement. Furthermore, as an introduction to web archives, they play a significant role; for instance, I can imagine them being an essential part of finding aids, if copyright issues can be resolved. But montages have substantial disadvantages as computational tools. In some sense, they are a simple visualization, a simple tiling of images, and other tools can tell us more by drawing on the images themselves. Nonetheless, montages give a blast of everyday life on the web, with lots of the smiling people, folks enjoying life, the products that everyday people purchase and sell, and the hand-drawn pictures that make the Web so interesting. In short, montages offer a visual compendium of humanity.

There are, however, other more analytically useful ways for us to study these montages. One potentially provocative type of analysis is through prominent colors. Colors are, at least on the surface, quite hard to explore. In the .ca collection, for example, there are 2,485,238 distinct colors, representing nearly every part of the color spectrum. Yet we can extract colors and normalize them; that is, we can take colors that are close to each other and chart them.¹⁶ For example, there are hundreds of different shades of red, but we might want to cluster them into simply three or four major shades of red and see how frequently they occur. The results of color clustering are promising, allowing me to discern differences between various top-level domains and GeoCities neighborhoods.

The color schemes of the .ca domain are dominated by relatively demure shades and are, to some degree, generally earth tones (the same can be seen in .com as well). In .ca, we have white (or MintCream), two shades of gray, black, firebrick red, blue, and so forth; meanwhile, in .com, we find two shades of gray, a snow white, indigo, brown, light blue, and orange.

Now compare these color schemes to a sample of images from GeoCities (in this case the EnchantedForest neighborhood), which are far more vibrant: a similar graph shows high frequencies of shades of blue, mint, beige, and sienna. Gone is the domination of grays. Instead, we find lots of blues, some green, and far more colors evenly represented. There are similarly vibrant colors in other GeoCities neighborhoods, too, making these image sets more similar—to a degree—to another top-level domain, the Chinese-based .cn, which has high frequencies of smoke, black, goldenrod, and blue. Indeed, the .cn domain is a bit more vibrant than the first two top-level domains I have looked at. And while differences may not be dramatic, coupled with analysis of montaged images, we begin to see the broad contours of an archives color profile. Still, color can be difficult to work with for several reasons, not least of which are my own accessibility

issues due to colorblindness. But even those who can see colors perceive them differently, something that needs to be kept in mind when visualizing colors (hence I largely rely on labels rather than the colors themselves).

Finally, we can begin to plot these images and arrange them based on various characteristics, such as hue, saturation, or brightness, to reveal patterns within the web archive, including the overall proportion of digital photographs, clip art, and the color palette. There are other ways that we can cluster images, from *k-means* to latent dirichlet allocation, but they are very computationally intensive. Therefore, I have selected the method discussed below because it effectively clusters images based on their color and brightness characteristics and does so on a personal computer.

Consider this visualization from the GeoCities EnchantedForest that plots median image brightness on the *x*-axis and median image saturation on the *y*-axis (see figure 7.3). On the *x*-axis, starting from the zero point in the middle and radiating out, we measure brightness. A pure black image would exist in the middle of the visualization, and a pure white image would be at the extremity of the sphere. On the *y*-axis, we measure saturation or the intensity of a color, ranging from washed-out faint colors (on the far left) toward very vibrant, pure colors. In practical terms this means that on the far left in the middle, we see washed-out pastel colors, and to the far right we see pure yellow or blue instances of clip art (an apple, or a Big Bird with a yellow background). Combine the two, and we'll see pure white or blue located at the far right of the sphere—far from the starting point of blackness in the middle. In short, remember: the darker an image, the closer it is to the *middle of the sphere*, and its positioning on the left or right of that center point depends on its saturation. As a result, similar colors will cluster together. While print does not do the pictures justice, examining clusters and zooming in to see the points at greater detail, we can see:

- Point A: Pastel-colored images, mostly clip art of recurring baby-focused images. Bright images with less saturation.
- Point B: Relatively dark images with middle levels of color purity, mostly digital photography. In all data sets explored, the size of the cluster at point B was a good way to gauge the sheer number of digital photographs.
- Point C: White-backgrounded images, typically an object centered, surrounded by white. Extremely bright images with high saturation.
- Point D: Intensely colored red, blue, and green images, generally clip art, mostly of smiling green bears and other children-focused

Figure 7.3. A visualization of images from the EnchantedForest GeoCities.com neighborhood, focusing on hue and saturation.

items. Lower saturation, darkness. This area was more pronounced in the EnchantedForest than other areas, due to the child-focused nature of these pictures.

The utility of this method can be seen when comparing EnchantedForest image clusters to other sections. From these images, using the similar points of comparison, we can see that the largest concentration in the Canada TLD is in the realm of digital photography. In the Chinese domain, we find a strong concentration of white-background images with products in the middle (notice the upper right), as well as a strong concentration of

digital photography. This sort of method would be a useful way to longitudinally explore web archives. We also see in the modern web some very long images, forming background elements, which stretch across the image arrays. By thinking about these globes in terms of hemispheres, one can quickly see the basic contours of a top-level domain, which could, perhaps, be abstracted to flesh out archival finding aids.

Recurring Images: Community through Image-Sharing and Posting

Some images appear multiple times in web-archived collections, especially in early collections before the widespread advent of personalized digital photography. People mostly crafted these websites with limited computing knowledge using clip art shared within their communities or provided by the hosting companies. Returning to our GeoCities example, the image type and the number of times it appears can tell us about the cohesiveness of the neighborhood on a number of levels. For example, we can see whether users followed community guidelines: we should expect to see child-focused pictures in EnchantedForest, motorcycles and cars in MotorCity, and so forth. We can also detect whether they borrowed among neighborhoods, which might indicate levels of unity between sites. If they were simply using GeoCities as a site of convenience, we should see limited borrowing; if it was a community, we should see more extensive ones.

To explore image-borrowing, I built a visualization interface using *Mathematica*'s rapid prototyping functionality. My goal was to view images in declining frequency using the following three fields—relative frequency order (i.e., were they the most popular image or the least popular image?), absolute frequency, and image preview—and to then examine how they were distributed across the web archive. The latter was important because if a site had the same image over and over again, we would want to see if an image was widely distributed or just concentrated in a tight field.

Finding duplicate images is difficult. The easiest way to detect duplicates is to calculate the hash function of an image, which means taking an image and reducing it to a number that can be a unique identifier for that image: “28806a10de038b0bf32420971a641a603c69e721d7753364bd04ef71211e7157,” for example. These hashes can be used to identify images, find duplicates, and compare images *far* more quickly than working with the images visually—especially when working with millions of images. The unfortunate side effect, however, is that relatively minor image

edits will mean that the image is given a different hash (changing the background color, for example, results in a completely different hash). While there are workarounds—notably perceptual hashing, largely driven by a need to detect copyrighted material that has been slightly altered, such as the addition of a watermark or the changing color of a background—they are intensive when working with millions of images.¹⁷ The widespread nature of clip art in the GeoCities era makes this less of a concern than calculating the hash function of images on the Wide Web Scrape archive, where there is more widespread image modification.

Some examples help illustrate the efficacy of this approach. Within the Athens GeoCities neighborhood, for example, the most popular image was a GeoCities animated GIF advertising the service; it appeared 465 times. At right, the visualization notes that the image was widely distributed throughout the neighborhood, appearing hundreds of times on many *different* websites.

In the EnchantedForest neighborhood, we used this method to find popular clip art that users borrowed from each other. One notable one was an animated GIF of a popular children’s cartoon character, Tigger, bouncing up and down on his tail. It appeared forty-eight times and was, again, widely distributed throughout the EnchantedForest. As the eleventh-most-popular image in the neighborhood, it offers clear evidence not only of image-borrowing and reuse of images fitting within the neighborhood theme, but a clustered set of interests among users or “residents.”

Examples abound from elsewhere and help us get the sense about community in neighborhoods. Some have more recurring images that are community-specific: bouncing Tigger, SafeSurf, smiling faces, a Hello Kitty animated GIF in EnchantedForest; a flip-book animated GIF, an email scroll, a spinning globe in Athens; checkered race flags in MotorCity. Others show less consistency: FashionAvenue, the fashion-focused neighborhood, had little distinctiveness, with recurring images being largely generic or GeoCities-related. Others showed overlap: Heartland, the family-focused neighborhood, had an exceptionally large number of waving American flags, as did MotorCity.

These affinities bear themselves out in link structures. When I generated a list of incoming and outgoing links between neighborhoods, I saw that MotorCity’s closest neighbor was Heartland (with 1,058 links from Heartland to MotorCity and 657 from MotorCity to Heartland). Images corroborate this. Neighborhoods that linked more to each other also seem to borrow images from each other more, showing patterns of visiting and

reuse, and perhaps reflecting the importance of the automobile in American culture.

Within the Wide Web Scrape, we can see widespread borrowing and use of clip art and recurring images both through the above Mathematica viewer and by exploring the distribution of unique image IDs. Take the case of images in the .ca top-level domain compared to those in .cn. In .ca, the recurring images are generally navigational elements (e.g., a bilingual “email us” button, an RSS feed icon, a printer icon), followed by recurring images hosted by universities (the University of Ottawa’s brand appears prominently). In the sample, then, we see that 2 images appeared more than 300 times, only 20 images appeared more than 150 times, and most images appeared only once. With the .cn top-level domain, a different story emerges: 11 images appeared more than 1,000 times, and 430 images appeared more than 500 times. This distribution shows widespread sharing across the top-level domain, and the duplicated images demonstrate widespread image-borrowing.

At a glance, then, these distant-reading techniques tell us quite a bit about GeoCities neighborhoods and TLDs. Some of these groupings have more recurring images, widely distributed across the corpus, whereas others have more individual items. By tracking the borrowing of images, we can use the data to complement broader studies of community and other indicators of web archive content. As we move into an era of distant reading, these sorts of inquiries will become increasingly common in automatically generated finding aids for archived web content.

Faces: The Human in Humanities

I mentioned at the beginning of this chapter that working with large montages of images had a sobering effect because it gave one the ability to look at wide swaths of human activity on the web. I want to return to this in my final section, when we look at the faces that make up web archives themselves. While previous inquiries have had a computational bent, examining faces involves a qualitative turn in analysis. Faces are a powerful way into the humanity that makes up the web. One famous digital humanities project is the *Invisible Australians* project, which used some 12,000 faces of Australians excluded from their country under the White Australia policy. Using facial detection, portraits were cropped and mapped on the website so users could see an individual’s record by clicking on the person’s face. But more importantly, as creator Tim Sheratt noted, the project “helps us

understand the scale of the White Australia policy and how it impacted on the lives of individuals and communities. These were not isolated cases; these were thousands of ordinary people caught up in the workings of a vast bureaucratic system. The shift of context wrought by these digital manipulations allows us to see, and to feel, something quite different.”¹⁸ I do not purport to think that GeoCities holds anything near the significance of widespread exclusion and racism. But I do think that looking at the faces of the Web has rhetorical power around scale, and on an aggregate scale it helps us get a sense of just how “human” a collection might be. A distant view of GeoCities faces at a low resolution reveals some of the possibilities—and pitfalls—inherent in using facial recognition technology on archived sources.

Facial detection is not perfect, but its implementation is relatively robust. We can, for instance, see true positives, or successfully extracted faces. We also see false positives, or the occasional erroneous face (e.g., some cropped text as well as some concentric circles and shapes have been falsely identified as faces), but for the most part, these are either faces or cartoon faces. Yet these people did not knowingly give their consent to have their likenesses included in a web archive (inclusion in these web archives was largely a function of not having a robots.txt file, and most users do not have access to that file). And as Jentery Sayers discusses elsewhere in this volume, these techniques raise “questions about normalizing bodies . . . and treating them as data, including the relevance of computer vision to privacy and social justice issues.”¹⁹

The Association of Internet Researchers may provide some guidance here, asking us to consider both scale and expectations of privacy,²⁰ but in the near future, a conversation about ethics will be necessary: we will soon be able to leverage facial recognition algorithms already in place in corporations such as Facebook and Google, and begin to think about longitudinally tracking the appearance of a single individual across web archives. This strikes me as research that needs to be carried out with an ethical conversation, perhaps governed by granting councils or institutional review boards.

Indeed, the power to see the past in these ways will challenge current ethical processes. Formally, websites and other content are considered published: so, in countries such as France and Britain, they all fall under legal deposit requirements. In Canada, which does not have a full-fledged web-archiving process, the *Library and Archives Canada Act* gives the government the legal power to collect this information without permission

(and even to legally demand passwords).²¹ The legal power to collect material, however, does not make it ethical—especially when we consider that most of these collections have been created without the knowledge or even implicit consent of the creators. As noted, the legal basis of the Internet Archive collecting a website is respecting the robots.txt file. If the code block:

```
User-agent: ia_archiver
Disallow: /
```

appears within the robots.txt file, then the Internet Archive will not archive your website—and will retroactively remove your content. To do so, however, users need to both know that the file exists *and* have access to edit it. Institutional ethics will soon develop, but in the meantime, the power of visual analysis suggests that historians need to begin participating in “live web” researcher conversations around ethical use of born-digital material. Distantly reading material occludes some of the concerns around consent and privacy. Reading closely might also mitigate these concerns if a site is part of large networks, encourages visitors, has a well-travelled guest-book, and is clearly intended for the public, thereby allowing us to treat it more as a “publication” than a rarely accessed, privately oriented personal web page. Until a broader discussion around the appropriate uses of these sources can be carried out, we need to approach these sources cautiously, even as we recognize the potential of extracting faces from digitized archival collections and making them part of the finding aid record.

Conclusion

Images can tell us a lot about the shape and contour of large arrays of web information. Distantly reading images, as digital humanists do with text, allows us to begin to think of these web archives as more than just the sum of their parts. They help us get at research questions such as: Can we see community in images? What can the shape of top-level domains tell us? How can we quickly find the information that we, as historians, need? Drawing our attention to such questions, I argue that we can gain new insights into the nature of community and intercommunity interactions.

Web archives are still largely new ground for historians, and we are currently unprepared to engage with the quantity of digital sources that will fundamentally transform our profession. However, an understanding of these data sets is necessary as historians adapt and seek to understand the

1990s and 2000s. By beginning to investigate how we can fruitfully explore web archives, we can improve our understanding of where they have come from, various methods to both textually and visually explore them, and, finally, their implications.

What does it mean to see the past in this way? First, it opens up new frontiers around the computational access of images. Previous research has demonstrated that as sources are digitized, they are used more; those that are not digitized are used far less.²² Historians have become increasingly adept in the use of textual analysis, from keyword searching to Google Books explorations to Spotlight searches on their desktop computers.²³ Images, however, have eluded this transformational process: for most periods, they need to be scanned or photographed, cropped, and deposited into a local repository to be useful. As other chapters in this book demonstrate, this is changing with the application of computer vision techniques on digitized primary sources. Thus, as we move into large collections of born-digital images—both scanned by large organizations such as the Internet Archive and Google, and increasingly, the veritable flood coming from web archives—a conversation around these methods is important.

The methods employed in exploring the GeoCities Archive and the Wide Web Scrape to make several interventions into this impending shift toward digital sources. First, they demonstrate that images can be accessed and explored on a large scale, similar to other forms of digital sources. As such, images are no longer just complements to textual analysis, to be used for illustrative purposes in an article or monograph, but can rather become a primary frame of analysis. Second, they illustrate that we can learn something about a community through images: the degree to which images were borrowed, the colors that they used, the consistency found within, or how they look when arrayed before the user on a screen. With data sets stretching into the realm of over nine million individual images, as in the case of GeoCities, this sort of work will become increasingly necessary.

The acceleration of data-set preservation will soon raise opportunities for historians to manipulate and explore ever-larger swaths of the past. As of 2013, the popular photo-sharing service Flickr was seeing three and a half million photos uploaded *every day*; many of these photographs are georeferenced and contain robust metadata, allowing researchers to match location and time to various themes.²⁴ Furthermore, as Turkel and Imhotep have noted in their contribution to *Seeing the Past*, YouTube sees a hundred hours of video uploaded every minute. Google Streetview has been collecting panoramic images of streets throughout the world since 2007, allowing

users to see changes over the last eight years. This all lets us imagine a near future where historians can *see the past* in new ways, but we should not get carried away: the entire historical record is not being preserved, and these increasingly archived images still represent a tiny glimpse into the past. Moreover, as with all instances of having too much information, we will be seeing a past that is mediated through digital delivery systems. Nevertheless, they present a much larger visual recording of the past than has ever been preserved before, bringing with it opportunities in the realm of visual culture, built culture, art, human expression, and beyond.

As far as computer vision techniques go, the methods discussed throughout this chapter are hardly at the cutting edge of computer science, but they are understandable and relatively easily explained and taught. We can control inputs, explore outputs, and quickly see how changing inputs leads to different findings. As historical sources increasingly become mediated through the digital, this is an avenue of research that matters to historians—we think we can understand textual analysis (although many do not—whenever something appears at number 1 and something else at number 100 in an ordered-result list, it is important to grasp *why* that is) but we now also need to think about computer vision.

This chapter presents a new historical approach to making sense of millions of web-archived images, borrowing methodologies from computer science and new media scholars. These data sets are still in their infancy. Yet as we move into the web era, historians will need a capacity to meaningfully deal with images. Keyword searches and databases have revolutionized how we approach text, and soon, computer vision techniques will allow us to use treasure troves of digitized images. Returning to our metaphor of the historian staring up at a mountain of millions of digitized images at the start of the article, you might have imagined a look of horror. I hope you can now instead envision a smiling historian—at the promise today, and recognition of the opportunity ahead.

These are all techniques that you can learn on your own. While this is not a handbook chapter *per se*, some important resources to get started can be found at Lev Manovich's aforementioned Software Studies Initiative at <http://lab.softwarestudies.com/>. Most of the other examples generated in this chapter came from *Mathematica*, a commercial platform that has robust documentation at <http://reference.wolfram.com/language/guide/ComputerVision.html>. An open-source equivalent, available with OpenCV at <http://opencv.org/>, runs on your computer's command line; although this coding can seem daunting, simple lessons at the *Programming Historian*

(<http://programminghistorian.org/>) can help you reproduce the methods employed in this chapter on your own collections of digitized images.

Notes

1. This has been discussed in Ian Milligan, “Mining the ‘Internet Graveyard’: Rethinking the Historians’ Toolkit,” *Journal of the Canadian Historical Association* 23, no. 2 (2012): 21; “Web Archives: A New Class of Primary Source for Historians,” Institute for Historical Research Discussion, London, England, June 11, 2013, <http://historyspot.org.uk/podcasts/digital-history/web-archives-new-class-primary-source-historians>

2. The concept of abundance was discussed in Roy Rosenzweig, “Scarcity or Abundance? Preserving the Past in a Digital Era,” *The American Historical Review* 108, no. 3 (June 1, 2003): 735–62.

3. Milligan, “Mining the ‘Internet Graveyard.’”

4. Franco Moretti, *Graphs, Maps, Trees: Abstract Models for Literary History* (New York: Verso, 2005), 3–4.

5. In this, we see connections to the argument made in Rosenzweig, “Scarcity or Abundance?”

6. Ian Milligan, “Welcome to the Web: The Online Community of GeoCities during the Early Years of the World Wide Web,” in *The Web as History*, eds. Niels Brügger and Ralph Schroeder (London: UCL Press, 2017): 137–58.

7. “GeoCities Neighborhood Watch Program,” April 13, 1997, available online, http://web.archive.org/web/19970413015812/http://www8.geocities.com/homestead/neighbor_watch.html, accessed July 25, 2013.

8. “GeoCities Neighborhood Directory,” December 19, 1996, available online, <http://web.archive.org/web/19961219233921/http://www.geocities.com/homestead/homedir.html>, accessed July 22, 2013.

9. See for example Constance Elise Porter, “A Typology of Virtual Communities: A Multi-Disciplinary Foundation for Future Research,” *Journal of Computer-Mediated Communication* 10, no. 1 (November 2004), <http://jcmc.indiana.edu/vol10/issue1/porter.html>; Howard Rheingold; “Rethinking Virtual Communities,” in *Cybercultures: Critical Concepts in Media and Cultural Studies*, ed. David Bell, vol. 3, 4 vols. (London: Routledge, 2006), 3–66; Howard Rheingold, *The Virtual Community: Homesteading on the Electronic Frontier* (Cambridge, MA: MIT Press, 2000), <http://www.rheingold.com/vc/book/intro.html>; Nancy K. Baym, “The Emergence of On-Line Community,” in *Cybersociety 2.0: Revisiting Computer-Mediated Communication and Community*, ed. Steven G. Jones (Thousand Oaks, CA: SAGE Publications, 1998), 35–68; Lori Kendall, “Community and the Internet,” in *The Handbook of Internet Studies*, eds. Mia Consalvo and Charles Ess (Blackwell, 2011), 310–25.

10. Milligan, “Welcome to the Web.”

11. Andrew N. Jackson, “Formats over Time: Exploring UK Web History,” *arXiv:1210.1714 [cs]*, October 5, 2012, <http://arxiv.org/abs/1210.1714>

12. For this, see <http://lab.softwarestudies.com/2010/11/science-and-popular-science-magazines.html>. You can find the tutorial at <https://docs.google.com/document/d/1PqSZmKwQwSIFrbmVi-evbStTbt7PrtxNgC3W1oY5C4/edit>

13. “Software Studies: Media Visualizations—Illustrations,” <http://lab.softwar-estudies.com/p/media-visualizations-illustrations.html>, accessed October 6, 2014.
14. Daniel R. Montello et al., “Testing the First Law of Cognitive Geography on Point-Display Spatializations,” in *Spatial Information Theory. Foundations of Geographic Information Science*, eds. Walter Kuhn, Michael F. Worboys, and Sabine Timpf, Lecture Notes in Computer Science 2825 (Berlin: Springer, 2003), 316–31, http://link.springer.com/chapter/10.1007/978-3-540-39923-0_21. My thanks to Shawn Graham for the reference.
15. While I have generated data for all the TLDs, I provide examples here from ones that provided strikingly different data than the .ca TLD.
16. Drawing on <http://blog.wolfram.com/2009/02/12/flag-analysis-with-mathematica/>
17. For more on this technique, see <http://blog.iconfinder.com/detecting-duplicate-images-using-python/>
18. See <http://wraggelabs.com/shed/presentations/digisam/#/text>
19. See Sayers chapter.
20. Annette Markham and Elizabeth Buchanan, “Ethical Decision-Making and Internet Research: Recommendations from the AOIR Ethics Working Committee (Version 2.0),” September 2012, <http://aoir.org/reports/ethics.pdf>
21. “Legal Deposit,” <http://netpreserve.org/legal-deposit>, accessed November 11, 2014.
22. Ian Milligan, “Illusionary Order: Online Databases, Optical Character Recognition, and Canadian History, 1997–2010,” *The Canadian Historical Review* 94, no. 4 (2013): 540–69.
23. The growing use of digital technologies is discussed in Jennifer Rutner and Roger C. Schonfeld, “Supporting the Changing Research Practices of Historians” (Ithaka S+R, December 10, 2012), <http://www.sr.ithaka.org/sites/all/modules/contrib/pubdlcnt/pubdlcnt.php?file=http://www.sr.ithaka.org/sites/default/files/reports/supporting-the-changing-research-practices-of-historians.pdf&nid=532>
24. See many of the possibilities at http://lab.softwar-estudies.com/p/research_14.html. “The Man Behind Flickr on Making the Service ‘Awesome Again,’” *The Verge*, <http://www.theverge.com/2013/3/20/4121574/flickr-chief-markus-spiering-talks-photos-and-marissa-mayer>, accessed November 11, 2014.

CHAPTER 8

Building Augmented Reality Freedom Stories

A Critical Reflection

Andrew Roth and Caitlin Fisher

Can augmented reality (AR) in the classroom create a richer connection to history for students? Can AR contribute to a better understanding of social justice? Can collaborative projects between historians and designers lead to new innovations in storytelling? In this chapter, Andrew Roth and Caitlin Fisher take us through their experiences designing and developing classroom-based AR that explores the experiences of escaped slaves who journeyed on the Underground Railroad to settle in Canada. Unlike the authors of other chapters in this volume, Roth and Fisher write from the perspective of nonhistorians working to help humanists see the past. Their chapter provides insights into the process of collaborative design and how historians and developers can work together to create meaningful and innovative expressions of history.

The Underground Railroad was one of the most important social justice movements of nineteenth-century North America. People from many different walks of life, religious denominations, and ethnic backgrounds united in a vast, secret system to assist refugees from American slavery in their quest for freedom. The Canadian provinces of the time were a main terminus of the Underground Railroad, but even in their new homes African American refugees faced and overcame many challenges and continuing racist discrimination. In 2009 the Augmented Reality Lab and the Harriet Tubman Institute at York University began a multiyear collabo-

ration to use new digital storytelling techniques to bring this history to life and explore the role of African Canadian refugees in helping to build Canada as we know it today.¹ This paper explores the collaboration between historians and AR developers to bring together different scholarly and artistic forms and new methods of communication and technology to give public audiences in Canada a more complete understanding of the Underground Railroad. How might AR enhance our capacities to reach new audiences and enrich the public understanding of history by helping them see the past?²

This chapter discusses our experiences developing two collaborative projects, *Breaking the Chains* and *Augmented Reality Freedom Stories: Using AR to Understand, Explore and Communicate New Narratives about Canada's Role in the Underground Railroad*. These projects combined historical research with artistic practices and technical innovation in the field of AR. This chapter does not attempt to outline all of our research-creation explorations and projects during the collaboration, but instead maps a trajectory of the project as it shifted between technologies and evolved over multiple iterations, resulting in a mobile AR application. We hope this discussion of our tools and methods, and of the lessons we learned, will be useful to others interested in using AR in their own research, teaching, or dissemination. While our work was successful in several ways, including the creation of a free mobile application that has been downloaded by users around the world, the project was anything but straightforward. Our earliest work on *Breaking the Chains* involved experimenting with a variety of products designed to share the research of historians at the Harriet Tubman Institute via a web-based AR experience. This early work served as the foundation for the more-sophisticated AR application for iPhone, iPad, and Android devices, *AR Freedom Stories*, which used flash cards to trigger historical AR experiences.

Our first collaboration, *Breaking the Chains*, included classroom testing, educator training sessions with participants from across North America, conference presentations, and a comprehensive online module with links to AR stories and the biographies of over twenty-four individuals involved in the Underground Railroad. Our team, in collaboration with descendant communities, developed lesson plans and AR experiences available free to educators, students, and the general public.³

Early in the project we were presented with some fairly rigid technical constraints in response to the intended primary-school audiences: we wanted the experience to be free, and we wanted to make use of technolo-

gies widely available in Ontario public schools, such as aging computers with limited graphical capabilities. Though this meant we would seldom be working with cutting-edge technology, we knew we could still explore many possible uses of AR as a visual communication tool. We opted for a well-established web-based development kit for what we had initially assumed would be the finished project.

For this first stage of our collaboration, we aimed to develop a workflow and tool kit that would both showcase the historians' original content about the Underground Railroad and give other users the means to create their own web-based AR experiences. To do so, we researched and imagined software architectures that could be repurposed by other fine arts and digital humanities researchers. Along the way, we experimented with various open-source, off-the-shelf, immersive, large- and small-scale hardware and software solutions. We used the Apple iPad to explore and test a map that comes alive in your hands. We tried playing cards that whispered secrets to the users, and we wrote new code and used smartphones to superimpose poetic footage on the landscape. We walked through each experience at different scales, from room-sized to hand-held, exploring multiple ways to harness emerging media for the digital humanities (e.g., digital traces, including audio, video, stills, and 3D imagery) to enhance these experiences.

Our team of historians (many of them undergraduate researchers, who worked to gather primary sources), our historical research objectives, and our methods are well described and publically available on the Tubman Institute's *Breaking the Chains* website.⁴ So here we intend to focus on the methods we used to work in a collaborative, cross-disciplinary manner to use AR in the service of history. Our work moved us through a series of stages, including developing flexible and responsive design strategies, experimenting with AR technologies to critically inform the role of new technologies in historical research, authoring new expressive tools, facilitating a rapidly iterative workflow for designing both the content and tools, reevaluating the tools and products with an eye to dissemination, and finally, testing and integrating the results. The following section explores each of these stages in greater detail.

Developing Flexible and Responsive Design Strategies

Some scholars argue that AR is a new expressive medium for storytelling, like radio or television.⁵ As such, the key to exploiting AR's potential

is creating room to experiment. Producing small experiences quickly and frequently (e.g., decks of AR playing cards with videos) helped us to understand which strategies were immediately applicable, but flexible in the face of changing technologies. And yet the point of entry of any user or author comes from an understanding of these previous media forms. For example, changing the project from a magic mirror to a magic lens experience over the course of our projects largely changed the ideal viewing situation, and in consultation with educators, our metaphor changed from a pop-up book to flashcards. This was a strong indicator to us of the capacity of old media forms, such as books and flashcards, in influencing both the user and the creator as we worked to understand the new medium.⁶

We also experimented with duration. It was long suspected, and quickly confirmed through our early tinkering, that laborious video editing to craft a story was lost on the AR user, who, fatigued or simply curious and eager to move on, would tire of holding still for the duration of a several-minute-long video. Working with shorter video clips meant we needed to supplement the information in each scene through other means.

We approached this project with an understanding of what might be possible based on some earlier work. In 2009 the Augmented Reality Lab at York University had produced a series of AR vignettes for the Toronto Museum Project using 2D images, 3D models, and video. And in 2010, Helen Papagiannis (then a graduate student), in collaboration with the lab, staged an exhibit at the !dea Gallery in the Ontario Science Centre. The exhibit, *The Amazing Cinemagician*, included a take-away flyer with a black-and-white FLARToolkit marker on it. Holding the flyer up to a webcam displayed a variety of AR pop-up tunnel books, designed by Caitlin Fisher, Andrew Roth, and Helen Papagiannis and inspired by works such as Martin Engelbrecht's *Garden Scene* (c. 1740). Users observed a short animation clip by peering down the length of the tunnel book. By combining stills composited in Photoshop layers with video and sound, we effectively began to develop a technique for AR collage.

We found the use of photomontage, collage, and bricolage full of fascinating possibilities. We were interested in the creative practice of working with AR collage—the capacity to add and rearrange bits and pieces as these were discovered or made available to us, and to create new meanings through association. Collage not only provided a means to organize and display materials, but also ways to find and create pattern and to organize archival material spatially in meaningful ways: we thought of collage as argument.

Collages also allowed us to theorize our AR constructions as a three-

dimensional archive, suggestive of the viability of new visual methodologies in historical research. For example, working with AR collage-building maquettes and small scenes helped us imagine a more complete archive and, in doing so, to identify gaps in the data. Creating inexpensive models into which historical footage—photos, oral histories, videos—could be added as it was found, suggests that AR recreation is a research/creation methodology with enormous potential. This was evident in the early models for *Breaking the Chains*, which were easy and inexpensive to generate; we could build and rebuild them within a day using existing tools.

One potential weakness or risk of using collage stems from easy remix possibilities facilitating unanticipated juxtapositions. Specifically, collagist practices can help generate unexpected or undesirable meanings, or imply causality or other relationships that do not exist. This might be an attractive feature for poetic or purely artistic endeavors, but it can be problematic in documentary and historical presentations. In our work with collage, we sought to mitigate this risk by involving the historians in an ongoing evaluation and critique of the visual relationships were proposing.

Using AR Technologies to Create, Rather than Simply Disseminate Knowledge

Democratized storytelling and production tools, such as cellphone cameras, “are as much about literacy as they are about innovation or entertainment,”⁷ and the capacity to both participate in and consume interactive media forms is often referred to as a new form of literacy. Rather than building an experience that unilaterally disseminates historical interpretations, we advocate the power of knowing through making, which involves communication and learning that is “facilitated by students actively participating in the creation process of digital storytelling.”⁸ To that end, we privileged a vision of AR for history that involved the dissemination of easy-to-use expressive tools. Ironically, the more polished and sophisticated final products we published to the app store lacked this very desirable knowledge-making feature that some of our early, less-sophisticated pieces were able to incorporate.

Authoring New Expressive Tools

From the commercial ventures of Layar, Aurasma, and BuildAR, and our lab’s own SnapDragonAR, to open-source initiatives such as the Designers’

Augmented Reality Toolkit (DART) and Flash-based FLARToolkit, recent directions in AR research and technology have involved creating easy-to-use AR authoring tools, significantly reducing the barrier of entry for new content producers. Certainly none of the existing platforms is an ideal solution to the challenges of authoring AR experiences, but each can be leveraged to make more polished experiences and to attract new audiences to the content. Moreover, each tool can help the user to think, imagine, and work differently. Similarly, to facilitate participation in new literacies in the classroom, we hoped to reduce the complexity of the authoring software we created for *Breaking the Chains* and *Augmented Reality Freedom Stories*.

The open-source software we used for *Breaking the Chains* would have allowed for technologically savvy audiences to build their own content, but we recognized that this would be a very small and specialized group. Furthermore, investigations into designer-based tools report that “the extensive time required to create content and to write programs to handle all the components of an AR experience often means that there is a large amount of latency between the inception of an idea and its realization,”⁹ reinforcing the need to simplify AR authoring tools. One of our goals, therefore, was to develop a tool and dissemination channel with a low barrier of entry to target history and social studies students in primary and secondary schools. Creating constraints around the technology (such as the number of markers or limiting the use of 3D models) provides a way to manage time and focus on content design.

Identifying Tools and Establishing a Workflow

While working on the *Augmented Reality Freedom Stories* we were also evaluating and experimenting with a new generation of imperceptible marker-tracking technologies to seamlessly add digital information to physical print.¹⁰ Although these technologies offer an enhanced user experience and a way to incorporate more visual information, the technology required significant development resources and was not readily available or cross-platform at the time. Still, to take advantage of any potential outcome of these experiments, we had to remain flexible in our approach to generating media assets. Recognizing that several other researchers in the AR Lab were using planar video clips to author AR experiences—in part owing to both video’s ubiquity and our close association with York’s Department of Film—we collaborated to produce technically ambitious

software that could streamline the workflow of using moving-image clips in AR experiences.

The first software we developed in the Augmented Reality Lab was SnapDragonAR, a drag-and-drop AR authoring tool. In her early work, Helen Papagiannis was using the runtime of an early AR marker library to overlay planar videos on individual black-and-white markers, creating a photo album with moving images.¹¹ Around the same time, graduate student Geoffrey Alan Rhodes was creating a prototype for 52 Card Psycho using Macromedia Director and the Designers' Augmented Reality Toolkit (DART) plugin.¹² To map fifty-two individual videos on a consumer-grade laptop computer, Rhodes was evaluating several AR marker libraries, eventually choosing Dr. Mark Fiala's ARTag because it demonstrated a minimal rate of confusion between markers as well as "robust lighting and occlusion immunity."¹³ Partnering with two PhD students in computer science and engineering, we developed a plugin in Max/MSP to facilitate rapid authoring using a graphical programming language as a user interface (UI).¹⁴ We chose Max/MSP as a UI for its tight integration with Apple QuickTime (including networked real-time streaming) and its extensive compatibility with both USB and Firewire webcams.

Creating a UI or "front end" in Max/MSP had several benefits. First, it became possible to change videos and perform live code changes at runtime. Second, it allowed us to save software iterations that could be useful in multiple projects, and distribute them as stand-alone applications. An added bonus was that the drastically reduced production and testing time caused the number of software iterations to explode. Finally, it allowed us to change content quickly, making it possible for individual artists to focus on refining the content rather than recompiling software. An extended collaboration with Fiala resulted in a refined version of that technology, a commercially available AR authoring product called SnapDragonAR, which allows beginners to author an experience using videos and stills within a few hours. Users can then publish these experiences to a player, and freely distribute them, which is a common feature in contemporary desktop AR software such as BuildAR Pro and AR Studio.¹⁵

Pitching the Project

Using SnapDragonAR, we developed a prototype for overlaying videos on paper map locations and demonstrated it to researchers at the Harriet

Tubman Institute, showing how the pieces could be used with integrated or handheld webcams. We also explained how the software could support video see-through AR head-mounted displays such as the eMagine visor.

Researchers from both the Tubman Institute and the AR Lab were enthusiastic about SnapDragonAR's potential, especially its ease of use, which, we hoped, would enable the team to focus on content and the relationship between the stories and the technology, rather than the technology itself. But there was a significant barrier to moving forward with SnapDragon: it was a Mac-only tool, and most school labs were outfitted with older-generation PC computers. Therefore, SnapDragonAR was necessarily limited to in-house experimentation and prototype development. Even so, throughout the Breaking the Chains and Augmented Reality Freedom Stories projects, SnapDragonAR remained the primary tool for experimenting with story length and the granularity of the digital videos in small-scale tests. It also remains the AR Lab's go-to prototyping tool, allowing us to develop and storyboard in a single day what might take a month to build with tools that deliver a more polished final product.

Needing a Windows-friendly, web-based option, we quickly moved to FLARToolkit. The open-source Flash Augmented Reality (FLAR)Toolkit (including FLARManager v0.4) functions as most people familiar with early AR would expect. It detects black-and-white AR markers (variously called *fiducials*, *barcodes*, *tags*, *glyphs*, *trackers*, etc.) through a webcam image (figure 8.1). The effect is known as a magic-mirror AR experience. Software superimposes digital artifacts on a video image of the user holding up a marker.¹⁶

Our work with the FLARToolkit Actionscript library had begun years earlier with the Toronto Museum Project: Handheld City project, which aimed to virtually showcase items and stories in the Toronto Museum archive (figure 8.2). Though the Toronto Museum Project website hosts hundreds of high-resolution 2D scans of artifacts, few of them were available to us for incorporation into the design before the official launch of the museum website. Therefore, we used open-source images and models to produce six vignettes with narrative audio related to the collection as well as a collage of 3D models, music, video, and still images. Our experience with FLARToolkit, as well as the knowledge we gained about basic information architectures, helped our collaboration with Tubman.

For example, the FLARToolkit framework we initially developed lacked the ability to change assets at runtime, making the design process slower. To reduce the complexity of recompiling the project, we modified the proj-

Figure 8.1. Test cube and marker outline superimposed on FLARToolkit marker.

Figure 8.2. Toronto Museum Project: hand-held city screenshot.

ect for Tubman to include an XML reader instead of hardcoding the names of image textures, models, and videos in ActionScript; by adding the texture names to the XML file, we could easily change the project with a plain text editor or a server side script, such as PHP. Removing the ActionScript compiler from the everyday development also bypassed a significant technical obstacle for designers on the project. As was previously observed with the use of SnapDragonAR, the number of iterations exploded. While in the end, SnapDragonAR was not a final software solution for either *Breaking the Chains* or *Augmented Reality Freedom Stories* for the reasons mentioned earlier, the fact that the software was so easy and fast meant that almost all of our finished projects had been prototyped or imagined in SnapDragon at some point. It remained a vital software and learning tool in the lab during this period.

Shifting between SnapDragon—the preferred authoring tool—and software solutions designed to cater to our audience, we developed an early sensitivity toward making platform-independent assets. In the early days of working with SnapDragonAR, we learned the value of creating an AR tool using 2D videos and images exclusively. Each of the authoring environments we used (e.g., Layar, Flash, Max/MSP, Unity) had differing support for 3D models, textures, and animations, and the difficulties users had creating, or even locating, thematically suitable 3D models in an appropriate format and polygon count prompted many of our researchers to make extensive use of 2D. This was certainly the case during the first iteration of *Breaking the Chains*. Historians on the team most wanted to communicate text-based material; audio, a few photographs, and original illustrations came later, but 3D never did. While the historians we were working with intended to pair their stories with 3D models for increased visual impact, they had neither the desire nor the skills to create the custom 3D models required. And viable 3D materials were not easy to source.

Our experience with the Toronto Museum project taught us that small, hand-held scenes could be quite compelling when they made good use of AR's capacity to create the illusion of 3D spaces and worlds. So instead of lamenting the lack of 3D models in the historians' archive or focusing on developing custom 3D models, we opted to create 3D virtual objects using primitives and texture them with alpha channel-enabled textures (e.g., a simple sphere overlaid with clouds or birds). Easy to create, the resulting complex layering of textures on 3D primitives produced a compelling array of vignettes that, visually, was quite magical. Perhaps more importantly, this method for designing content independent of platform allowed for the

eventual publication of a polished application created in the Unity game engine, which used many of the same assets to texture 3D primitives.

Facilitating a Rapidly Iterative Workflow

We had imagined that the Tubman historians on the project would want to work closely with these new AR tools or, at the very least, work in close physical proximity with AR Lab researchers to create vignettes themselves, or with our close assistance, thereby allowing for rapid iteration. But most did not feel comfortable editing XML files; others felt that their time was better spent conducting research and supervising the creation of scenes. Such reticence, however, worked against some of the project's foundational goals. Aside from creating a workload imbalance, the AR lab researchers—not the historians—were getting hands-on experience in combining primary materials with the affordances of AR. As a result, AR Lab researchers, rather than the historians, were pushed to think deeply about how the pieces made meaning and what, if any, difference an arrangement made for communicating history. To address these shortcomings, we had to create specific times to solicit feedback that might have emerged organically if historians had been more active in the AR process. Some feedback we received referred to technical bugs or design issues: “Title seems like it’s a lot higher than the rest of the images. I didn’t see it until I held the barcode a lot further back.”¹⁷ More critically, several incidents also highlighted the importance to the intellectual integrity of the project of having excellent cross-disciplinary communication at multiple stages. In one scene, for example, we used found footage of a train crossing a bridge between Canada and the US; however, an historian familiar with the route pointed out that the train was headed toward the US, rather than toward freedom in Canada as the scene intended. Because such feedback was vital, some historians’ resistance to being involved with AR technology slowed our ability to rapidly iterate.

Other efforts, however, made workflow more efficient. The first stage in the design process, for example, was to create an overview of the design materials and tasks for constructing each AR scene. In response to our sample prototypes, as well as the AR Lab’s emerging understanding of appropriate granularity and story length in AR experiences of this type, the Tubman Institute research team determined that each scene would showcase a short, scripted story, generally thirty seconds or less, communicated as voice-over. A theme would then connect the visuals within individual

scenes and across the entire experience. This aspect of the collaborative workflow moved quickly.

Many of the primary research materials were being located and scanned, and interviews were being conducted, while the scenes were being constructed—not ideal, but a function of the timelines of the project. In other cases copyrights took time to clear. Scenes were loaded to an FTP server, each with its own script and identifying information for some of the assets. To stay organized, we used a spreadsheet in Google Docs we referred to as a “shot sheet.” Members of the AR Lab design team assigned their names to scenes that included the name of the subject (person or place), a delegate at the Tubman Institute responsible for the assets in that scene (especially historical accuracy), the file names of the individual assets, changes that still needed to be made, and the overall status of the scene. We also developed a system for AR Lab designers and historians to communicate with each other to clarify when the designer should use the assets uploaded to the folder to build a scene, when the historian should view the scene and offer feedback, and when the historian and designer agreed that the scene was finished. The shorthand “Ready,” “Set,” and “Go” for these three phases of completion or milestones became important to introduce new users to the workflow and quickly quantify how much work remained on the overall project.

While Tubman historians were willing to use the research to script original stories and record them, they had little commensurate enthusiasm for finding or creating images. The majority of images favored by the historians were black-and-white, text-based documents. The importance of visually compelling and accurate images—as exemplified by the train leading freedom-seekers back south—further complicated matters. AR researchers, for example, struggled to create a visually appealing color palette within and between scenes with little direct collaboration and input from the historians. Moreover, subtle changes in the color or composition of an element often produced significant changes in how the scene was received by its audience. It was only after the majority of the scenes started to develop a complex palette that we better understood how to select a few scenes and use the desaturation of new clippings alongside black-and-white photography to dramatic effect (figure 8.3).

Other elements that affected project workflow include the technical demands of sharing scenes between the historians and the team in the AR Lab. Specifically, we needed to be ever-mindful of which version of the Flash player was being used and whether the browser cache was storing

Figure 8.3. Levi Foster AR scene from the iPhone version of AR Freedom Stories.

out-of-date images. Access control (primarily to control intellectual property prior to the release of the project) was a factor we should have anticipated, but did not. Simply creating a local .htaccess file to restrict viewing of the final scenes to servers on the York campus and a list of personal IP addresses for project members would have been sufficient, but the fact that we did not have this system in place at the start of our work complicated our workflow. Finally, we realized that test users needed access to the scenes to help evaluate user experience, so we implemented a low-security (client-side), ad hoc password-protection scheme. Because such a solution relies heavily on security through obscurity, it wasn't ideal. If we had considered security requirements earlier, we could have incorporated an open-source PHP/MySQL user management system, such as UserCake, to accommodate customizable access.

Reevaluating the Tools and Products

As we were developing an AR application suitable for mobile devices in Unity, two new AR marker libraries came to our attention: QCAR beta (now Vuforia) and String. Both of these libraries included plugins for the Unity 3D game engine and could build for iOS, so to test their viability for our purposes we ported a scene from the Ontario Science Centre project to

Unity and started developing a workflow that could be scaled up to accommodate larger projects. To develop a mobile application, Unity and String, or QCAR, seemed to be the best option for the large-scale dissemination of AR Freedom Stories, but the commercial license for String and Unity combined was too expensive, and the Apple Store had not yet approved QCAR apps for distribution, prompting us to explore other options.

And once the assets were in place, we were inspired to imagine new possible versions of the project, with different viewing situations, potentially reaching new audiences, reimagining, for example, Breaking the Chains resources as life-sized virtual walkthroughs, and a complementary mobile application that would deliver locative tours and mobile AR experiences. The increasing ubiquity of mobile devices also increased the possibility of dissemination for the latter two experiences to educational settings beyond technology-equipped classrooms.

The AR Lab's experience with locative applications suggested several possible scenarios for locative storytelling. One popular technique involves *in situ* storytelling, which replays story arcs when the user with a GPS-enabled device approaches a predefined set of coordinates. Although this technique would have allowed us to present the history in or near the space where it actually occurred, we were concerned that stories for this particular project would be distributed too far apart to be practically navigable. Another technique we considered involved overlaying stories at a deliberately different location, rather than seeing the landscape and the story as coconstitutive of meaning, in order to recontextualize both the physical landscape and the story and allow both to be interpreted in a new and potentially powerful way. Beyond being theoretically interesting, in practical terms this meant that we could have clustered story nodes closer together, thus reducing the space required to tell the story. But this technique also has the potential to affect user's reception of the story in a negative way—for example, by communicating historical inaccuracies. We were also concerned about point of view in a locative media application—what are the politics of positioning the user as an objective outsider witness to this kind of story while you actually walk in the footsteps of these historical figures? These remain urgent and interesting questions for documentary work in *in situ* locative AR, and playing with points of view can be a powerful tool for work in locative media, but we felt that we could not tackle these issues within our time frame.

While we would have preferred to develop a polished, GPS-based, locative app, then, we abandoned that goal in favor of producing a ro-

bust handheld mobile experience that could be used within classrooms and workshops and that seemed both more practical and less problematic. The Vuforia library made it easier to use multiple markers with the AR experience and therefore made AR Freedom Stories on a mobile device possible.

In early iterations, the markers we used for the Augmented Reality Freedom Stories Vuforia-based app were simply images pulled from the existing FLARToolkit scenes, creating two problems we had to overcome in the final design. First, the user had no way of knowing which marker corresponded to which scene—there was often no indexical relationship with the content, making any connection to the physical marker as tenuous as any connection to a black-and-white icon. Second, not all images produced high-quality *trackables* according to Vuforia's training algorithms; in other words, the camera would not identify the images reliably, if at all, making images useless for our AR experiences. To address these difficulties, we added quotations and names to each *trackable* image and then arranged the images we liked best—even ones that did not initially track well—alongside new edges and contours (using filters, palimpsests, and geometric shapes) and text (using a jagged edged font), thereby providing the necessary contrast to create both high-quality *trackables* and images that would be intelligible to users.

On March 15, 2013, the first Augmented Reality Freedom Stories app was approved for distribution in Apple's store, and as of April 20, 2018, the app has been installed over 5,500 times, including to 292 devices outside of the US and Canada.

Testing and Integrating the Results

Although one of our key target populations—Ontario middle-school students—initially mandated an AR solution that would accommodate aging Windows-based machines, a wide range of currently available AR technologies can now be used with a wide variety of classroom equipment, from tablets to projectors to head-mounted displays. And more subjects—from art to math to biology—stand to benefit from this technology.¹⁸ Encouragingly, there is evidence to suggest that the use of AR displays is as at least as effective as conventional displays for the purpose of education or training.¹⁹ Yet our experience shows that AR currently presents several challenges to storytelling in this transitional moment, as both literacies and hardware and software change: existing “AR literacy” and understanding, users' physical comfort while using devices, user attention span, and user

expectations for AR experiences. Surprisingly, our extensive work in classrooms reveals that these challenges often have low-tech solutions.

Researchers at the Tubman Institute conducted four workshops with teachers to gather feedback about the online resources. Natasha Henry, a Tubman Institute educator and historian, demonstrated the *Breaking the Chains* resources and *Augmented Reality Freedom Stories* app to this audience and reported that the web-based AR applications online were generally well received. She attributed much of the participants' interest as a response to technology-based learning directives in their schools and to schools both supplying technology and encouraging teachers to make use of the technology students have at their disposal. Teachers said that they appreciated the new approach to presenting black history in Canada through digital storytelling.

We anticipated some resistance to the use of experimental technology that required peripherals such as a printer or a webcam, but found it was largely absent. Teachers indicated that the necessary technology (such as computers with integrated webcams) would already be in the schools. There were unanticipated issues, however. One teacher was concerned, for example, because the administrators at their school blocked access to the computer's webcams for security reasons.

Several questions consistently emerged throughout the workshops, often with regard to the use of the marker. Notably, teachers wondered if students would physically be able to hold the marker up for the duration of all the stories. Our experience with AR revealed that short stories, thirty to sixty seconds long, were well suited to the physical demands of holding markers or positioning devices, but we have never held up markers to experience thirty stories in a single lesson, and we are not sure how audiences might ultimately consume this content in a limited time frame.

While it was common for users to view multiple scenes sequentially, we never considered that viewing most or all of the scenes was likely. Yet, surprisingly to us, this was how the workshops were conducted. In anticipation of user fatigue, Natasha demonstrated exceptional creativity by affixing the paper marker to a table tennis paddle. The handle and stiff backing made it easier for teachers and students to aim the marker at the webcam and may have resulted in a more sustained engagement with the content.

This concern about the printed markers is reflected in one of the published app store reviews: "The app is slow to start—it takes longer than it should to get to the point where you can scan the postcards. The images are good, but you can't really read the text from the article clips they have

as part of the augmented reality. I do suggest printing them out on card stock if you are going to use them in the classroom.”²⁰

A few teachers and researchers at the Tubman Institute also commented on the illegibility of text within the AR work. While this was an issue for us in the AR Lab, we also had approached the vast quantity of text as being, in some instances, suggestive of a set of ideas rather than assuming that each digital artifact would be scrutinized individually. For the AR experiences, we felt that the intelligibility of a newspaper’s headline would be sufficient, given that historians developed associated lesson plans to work in tandem with the digital storytelling. During the workshops, the AR experiences were always demonstrated alongside lesson plans. As mobile devices become more powerful, it may be possible to use higher-resolution textures that retain their legibility. It will also be possible to allow users to zoom in to or access the online resources from within the app. What is interesting, here, however, is the desire to use the AR as more than just a complement to the primary documents: the AR is the gateway to accessing the historical record.

Conclusion

Augmented Reality Freedom Stories was about building tools, content, methodologies, and community, while being responsive to changing technologies and expectations over time. We aimed to understand AR technology and its potential better through making and creating new opportunities for sharing knowledge, and we sought to enhance both digital and historical literacy by nurturing new communities of both makers and users. In these areas, we were successful in providing literally thousands of educators and students with innovative access to new research on people of African descent who came to Canada in search of freedom. In the process, our work adds to a growing list of tools and workflows that can be used to create and disseminate new kinds of historical texts. It is also suggestive of how AR technologies might provide powerful new ways to build and share knowledge. However, as P. Juala notes, the humanities “as a whole [tend] not to be aware of the tools developed by DH practitioners,” placing a greater burden on AR creators to craft and deliver bespoke content themselves rather than building on the work of others.²¹ Knowing this, we worked very hard to bring our tools to the attention of diverse audiences. For example, we made the FLARToolkit architecture freely available to communities of individuals who did not readily identify as programmers,

and Andrew Roth developed video tutorials, walk-throughs, and lightning talks both to disseminate this work and to support those interested in making AR without compiling code.

We engaged the Vuforia community when we encountered technical issues while developing in Unity and Vuforia, and we freely shared the patches we developed to overcome limitations in the existing code; one such patch for loading large numbers of videos dynamically has over 22,300 page views.²² Yet despite these efforts and successes, we are compelled to ask, why do so many excellent tools and scalable techniques remain largely unused or unrecognized on completion? There are several possibilities.

One explanation is that all software requires continuous upkeep to stay viable on new hardware architectures and new operating systems. As the scale of distribution increases, so do issues of security and censorship, making it sometimes impossible for digital tools to reach their intended audience. Notably, the cost of building software also becomes relatively small compared to the time and energy that goes into supporting communities of users once the software has been released.

If digital humanists are building communities of practice along with the software they create, then it is likely that the success of any tool can be defined in its own terms, but surely one vision for the future of this kind of work is the ability to leverage the fruits of earlier projects in the service of new endeavors.

Another possible explanation for a lack of uptake in the tools “is a mismatch of expectations between the expected needs of [the] audience (market) for the tools and the community’s actual needs.”²³ Because of our small development team and the scale of the projects involved, we were able to perform informal user testing with collaborating researchers at the Tubman Institute, teachers who attended training workshops, and students taking university courses involving AR. These users are also frequently the project’s continuing evangelists. Nevertheless, we had assumed, incorrectly, that historians would want to use our tools, if only we reduced the barriers to entry sufficiently. Key lessons to share from our experiences are the need to include stakeholders early and frequently in a workflow designed for rapid iteration, to find ways to encourage hands-on use of tools by new users outside the digital humanities umbrella, and to make a strong case for what is at stake in doing so: namely, new ways to share knowledge and reach new audiences and to learn through building with next-generation expressive tools

Our experience shows that, in practice, it is all too easy in cross-

disciplinary collaborations for participants to solely focus on areas of comparative advantage: historians understandably want to focus on history. Yet the process of creating these projects taught us the important lesson that we need to be open to new methods and stretch ourselves beyond our academic comfort zones, so that we can discover new avenues for learning, teaching, and creating new ways to see the past.

Notes

1. AR Freedom Stories and Breaking the Chains were developed with the support of the Social Sciences and Humanities Research Council and the Canada Research Chair Program and under the direction of international award-winning AR artist and theorist and Canada Research Chair in Digital Culture, Dr. Caitlin Fisher, director of the Augmented Reality Lab at York University, with AR Lab Technology Manager Andrew Roth. The AR Lab worked with one of the foremost authorities on the Underground Railroad, Distinguished Research Professor and Canada Research Chair in African Diaspora, Dr. Paul Lovejoy, and Governor General award-winning historian, Dr. Carolyn Smardz Frost, author of *I've Got a Home in Glory Land: A Lost Tale of the Underground Railroad*.

2. There is extensive literature on pedagogy and historical thinking skills that we do not have space to explore, since this chapter focuses on the experience of working with and designing for historians. For a discussion of the work of Peter Seixas, Sam Weinberg, and others and how they relate to AR, see the chapters by Comepeau and MacDougall, and Kee, Poitras, and Comepeau in this book.

3. Our research has been widely disseminated at events such as the tenth anniversary of HASTAC in 2012 and “Dialogue on the Diaspora: A Black Heritage Expo” in 2011, at conferences for the Elementary Teachers’ Federation of Ontario, a Peel District School Board summer institute, and at the National Park Service National Underground Railroad Conference. Additionally the work has been featured by educational technology bloggers, and appears regularly on embedded ed-shelf lists under the category Apps for Education. C. Beyerle, (n.d.), “Augmented Reality for ED” [Smore Online Flyer], from <https://www.smores.com/u00w-augmented-reality-for-ed>, retrieved October 18, 2014; J. Corder and J. Gore, “Augmented Reality—J2 Training,” from <https://tackk.com/augmentedreality>, retrieved October 18, 2014; C. Pepe, APP of the Week—“Freedom Stories” [YouTube], from <http://youtu.be/UvOP-Wb90ME>, retrieved October 18, 2014.

4. <http://tubman.info.yorku.ca/educational-resources/breaking-the-chains/>

5. Blair MacIntyre, Jay Boulter, et al., “Augmented Reality as a New Media Experience,” *Proceedings IEEE and ACM International Symposium on Augmented Reality* (2001), doi:10.1109/ISAR.2001.970538.

6. <http://www.ontariosciencecentre.ca/sciencenow/ideagallery/cinemagician/>

7. E. Rennie and J. Hartley, “The Story So Far: Digital Storytelling, Narrative and the New Literacy,” *Image, Text and Sound 2004: The Yet Unseen: Rendering Stories*. (Melbourne: RMIT Publishing, 2004): 1.

8. B. R. Robin, "Digital Storytelling: A Powerful Technology Tool for the 21st Century Classroom," *Theory Into Practice* 47, no. 3 (2008): 220–28.
9. Blair MacIntyre, Maribeth Gandy, Jay Bolter, Steven Dow, Brendan Hannigan, "DART: The Designer's Augmented Reality Toolkit" *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003 Proceedings*, 329–30.
10. Vincent Lepetit and Pascal Fua, "Keypoint Recognition Using Randomized Trees," *Transactions on Pattern Analysis and Machine Intelligence* 28, no. 9 (2006): 1465–79; Vincent Lepetit, Peter Laguerre, and Pascal Fua, "Randomized Trees for Real-Time Keypoint Recognition," in *Conference on Computer Vision and Pattern Recognition* (San Diego, CA: 2005), 775–81. DOI: 10.1109/CVPR.2005.288.
11. Helen Papagiannis, "Augmenting Digital and Analog Memory," *Public* 34 (Fall 2006): 52–59.
12. Maribeth Coleman, *Creating Augmented Reality Authoring Tools Informed by Designer Workflow and Goals* (2012), retrieved from <http://smartech.gatech.edu/handle/1853/45845>
13. M. Fiala, "Artag Revision 1, a Fiducial Marker System Using Digital Techniques" (November 2004), retrieved from <http://nparc.cisti-icist.nrc-cnrc.gc.ca/npsi/ctrl?action=rtdoc&an=5765376>
14. Andrei M. Rothenstein and Mikhail Sizintsev, "52 Card Psycho: Implementation details," July 9, 2008, http://52cardpsycho.com/52CardPsycho_whitePaper.pdf
15. HIT Lab NZ, "BuildAR," from <http://www.buildar.co.nz>, retrieved September 27, 2014; S. Malik, G. Roth, and C. McDonald, "Robust Corner Tracking for Real-Time Augmented Reality," in *Vision Interface* (May 2002): 399–406, retrieved from <http://nparc.cisti-icist.nrc-cnrc.gc.ca/npsi/ctrl?action=rtdoc&an=9062774>
16. J. Grosjean and S. Coquillart, "The Magic Mirror: A Metaphor for Assisting the Exploration of Virtual Worlds," in *Spring Conference on Computer Graphics (SCCG)* (Bratislava, Slovakia: 1999).
17. From a personal communication with D. Vavala (2012).
18. Mark Billingham and Andreas Duenser, "Augmented Reality in the Classroom," *Computer* 45, no. 7 (2012): 56–63; Hannes Kaufmann, "Collaborative Augmented Reality in Education," keynote speech at Imagina Conference, retrieved from <http://www.ita.mx/files/avisos-desplegados/ingles-tecnico/guias-estudio-abril-2012/articulo-informatica-1.pdf> (2003); Mehmet Kesim and Yasin Ozarslan, "Augmented Reality in Education: Current Technologies and the Potential for Education," *Procedia—Social and Behavioral Sciences* 47 (2012): 297–302.
19. A. W. Stedman, K. Hill, R. S. Kalawasky, and C. A. Cook, "Old Theories, New Technologies: Comprehension and Retention Issues in Augmented Reality Systems." *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 43 (23) (1999): 1358–62.
20. Socdownload. (n.d.). "Not Bad." [iTunes App Store review], <https://itunes.apple.com/us/app/augmented-reality-freedom/id616766825?mt=8>, retrieved October 18, 2014.

21. Patrick Juola, “Killer Applications in Digital Humanities.” *Literary and Linguistic Computing* 23, no. 1 (2007): 73–83.
22. Qualcomm Vuforia Developer Portal—posts by cap10subtext. [User support forum], https://developer.vuforia.com/forum/user_posts/11690, retrieved April, 2018.
23. Juola, “Killer Applications.”

CHAPTER 9

Experiments in Alternative- and Augmented-Reality Game Design

Platforms and Collaborations

Geoffrey Rockwell and Sean Gouglas

Inspiring humanists to engage with new technologies and new ways of sharing history is a tall order, taller still when the technology is experimental. In this chapter Professors Geoffrey Rockwell and Sean Gouglas share their experiments in developing mobile experiences, bringing together students, historic sites, and businesses in their attempts at creating immersive place-based applications. Their experiences provide a snapshot of an emerging discipline, in which the trial and error of developing history games to explore a former frontier help push the boundaries of a new frontier of technology and public engagement.

Since the release of the iPhone in 2007, smartphones and their affordances have inspired innovative services and social toys that help us “augment” and simplify large segments of our daily lives. They also allow us to augment our experience of space and play using our location, as evident in one trend in digital entertainment: locative games. Offering rich simulations of complex environments, alternate- and augmented-reality games (AARGs¹) invite participants to navigate physical locations, collect information, and solve problems collaboratively, providing an interesting forum in which to present and model knowledge, promote playful user engagement, and encourage meaningful learning.² Yet despite the proliferation of devices that combine imaging, ambient data, and locative sensors, designing, testing,

and refining such games remains challenging and time-consuming. Nevertheless, we believe there is significant promise in developing game design platforms for collaborative learning. In an effort to understand the opportunities and limits of developing AARGs for educational purposes, this paper explores the rewards and challenges we encountered in two AARG collaborations.

To ensure the most strategic context and incentive for potential players, we partnered with two stakeholder organizations that had domain expertise in both the proposed play spaces and the presumed players. We also wanted to compare different authoring environments, so one collaboration used an AARG platform we created; the other collaboration used SCVNGR, a commercially supported product. After an overview of AARG and the ways its technology has evolved, we turn to the locative games we developed. For each, we explain the game's objectives, advantages and disadvantages in using the particular design platform, and participant responses. We conclude with a discussion of where locative game designers might best spend their resources in future developments.

Alternate- and Augmented-Reality Games

Although AARGs seem to be a byproduct of today's technology, AARG-like game elements predate modern incarnations. An element in G. K. Chesterton's *The Tremendous Adventures of Major Brown*, published in 1905, is a case in point. Bryan Alexander (2006) notes that, like AARGs, "The Adventure and Romance Agency, Ltd. use[s] planted characters in Brown's neighbourhood, 'chance' encounters, mysterious messages, [and] hidden locations to simulate an invigorating adventure. The plot only unravels when Brown discovers the Agency, who then asks him to pay the game's bill." David Fincher's film *The Game* (1997), in which a wealthy investor is given a live-action game by his unpredictable brother, offers a modern-day interpretation of Chesterton's work. William Gibson's novel *Pattern Recognition* (2003) features a community of film buffs, scholars, housewives, and others who piece together clips of scenes to solve a larger mystery, while his *Spook Country* deals explicitly with AARG art. Also falling into this category are works that incorporate many AARGs' "This is not a game" (TINAG) philosophy,³ reflecting a wonderful tradition of narrative hoaxes; the *Codex Seraphinianus* (a purported alien encyclopedia) and Orson Welles's radio adaption of *War of the Worlds* (1939) are two such works. Other ancestors of the AARG form include live-action role-

playing games (LARPing), murder mystery evenings, scavenger hunts, and Gotcha-style assassin games.

More-formal AARGs began appearing in the late 1990s, particularly when players could collaborate and consult more easily through the internet. Two of the most influential of such early efforts were *The Beast* (produced by Microsoft as a tie-in for Stephen Spielberg's movie *A.I.*) and *I Love Bees* (produced as a promotion for the launch of Microsoft's *Halo 2*). These games were remarkably successful in developing a coherent community of participants who eagerly and enthusiastically collaborated to solve puzzles and advance the story.⁴ Such product tie-ins drove (and continue to drive) AARG development, including *The Heist* for the Audi A3 and *Year Zero* for Nine Inch Nails' album of the same name.⁵ In time, such large-scale AARGs, which involved participants and events that crossed continents, began sharing the stage with smaller grassroots games. As Andrea Phillips and Adam Martin note, "Typically costing around \$500–\$1000 at the low end, and around 1–4 months part-time pre-production work for 2–4 people followed by a further 3–6 months of part-time work whilst the game is running, grassroots games are well within the grasp of pretty much any group of friends with dedication and good ideas."⁶

Advancing digital technology continues to blur the lines between physical and virtual domains, allowing users to seamlessly traverse between their physical and digital lives, and game designers have capitalized on this ever-merging space, moving beyond games that favor one space over the other and creating alternate-reality games and augmented-reality games—products that allows users to exist in both spaces simultaneously. By interacting with their physical world, players can change the game's status in the digital world.

Many of AARGs' constituent elements have appeared in other game forms (e.g., puzzles, role-playing, collaborative play), but an AARG field of play tends to be vague, its goals shrouded, its rules unclear, and its collaborative play emergent.⁷ This diversity of forms has sparked the imagination of game designers, the public, and the academy—especially considering that "the unique activities of ARG communities can enlighten producers of media on how to design for participation, provide insight into tools for empowerment, are ideal marketing hubs and also illuminate the nature of communication and networks in general"⁸—yet finding the most strategic ways to advance work in this area has been challenging, if only because game designers, players, and theorists disagree on what elements of AARG are essential.⁹ Most will include a narrative structure, collaboration, inter-

activity, multiple forms of media, and the TINAG game-play philosophy,¹⁰ which reflects a common practice of game players and designers to suspend disbelief at all times as to the game's falseness. In other words, even though the game clearly runs parallel to reality, the fact that it is a game should never be admitted—though some may dispute whether the TINAG ethic is essential to AARG.¹¹ All, however, agree that puzzles are essential. According to Andrea Phillips, five types of challenges are AARG staples: cryptography, games, achievements, social engineering, and puzzles. Together, these provide the shape and nature of a story that is, generally, revealed over time.

With this context in mind, we now turn to our experiments in developing two AARGs, demonstrating the potential of locative games for educational purposes.

Experiment 1: *The Intelliphone Challenge*

Produced in 2012 as a collaboration between the University of Alberta Humanities Computing department, in conjunction with the Computing Science department, and the city of Edmonton's Fort Edmonton Park historical site, *Intelliphone Challenge* is a locative game that invites players to learn about Edmonton history through engaging narratives, which guide visitors through Fort Edmonton Park and its facilities. Wanting to see how players engaged with and moved through a particular space, we designed the game with the fAR-Play (which stands for "for augmented-reality play") platform, which made it possible to examine user experience, with a particular emphasis on interface design.

fAR-Play is a software platform we designed to help others easily create geolocate, scavenger-hunt-style games that would use the GPS and web-connectivity features of a player's smartphone to help players engage with their physical environments in novel, educational ways.¹² The front end, which delivers game content to players, is built in a combination of HTML5, PHP, and JavaScript. The back end, which keeps track of the games and players, is created with MySQL. JSON encoded data and dynamically generated SQL queries allow these two ends to interface.

At the most basic level, fAR-Play is a tool that tracks player's proximity to arbitrary points in space defined by a game. When players sign in to an AARG created in fAR-Play, they are presented with a list of games, a brief description of each game, and the various "adventures" within each game. To progress through each adventure, players must travel to a pre-

defined “node” (that is, a physical location game developers set using map coordinates) and “capture” the virtual points of interest (vPOIs). Players complete an adventure when they have captured all of the nodes within it.

Depending on the game’s objectives, players can capture nodes in various ways: answering a particular challenge question, scanning a specially created QR code located in the physical space, or using Layar’s various augmented-reality features to access a range of extra content that is overlaid on the physical play-space or that launches another application for extra functionality. Developers can implement hints in the form of images, audio, or video files to guide players through the experience; they can also determine whether players must visit vPOIs in a particular order, or if order is inconsequential. Each captured node rewards the player with a predetermined number of points, which functions both as a score and as a currency; if players are having trouble with a challenge, for example, they can use their existing points to purchase hints. Designers can also assign “achievements” to reward players with extra points for completing certain tasks or challenges. Similarly, designers can penalize players for various behaviors; for example, if designers want to encourage players to discover answers for themselves, they can reduce a node’s value by a predetermined amount if players use hints to capture it.

To track players’ progress, the fAR-Play website can have a leaderboard that displays player profiles, which includes each registered player’s accumulated points, captured nodes, and achievements, as well as the player’s level (novice to super player) based on the player’s total score; this data lets designers identify players according to two classifications of achievements, “Percentage of Adventure Completed” and “Percentage of Game Completed.” Players can also share their adventure progress, achievements, and vPOI captures on Facebook. Collectively, these fAR-Play features allowed us to develop a game that would encourage Fort Edmonton Park visitors to explore the facility in ways they might not otherwise have done.

Intelliphone Challenge’s narrative follows the members of a fictional Edmonton family as they experience growth and development in three areas the park wants to emphasize: security, communication, and community. The park itself is divided into four sections, each representing a key era in Edmonton’s history. Using three of these sections—the 1885 Street, the 1905 Street, and the 1920 Street—we created three distinct adventures, each of which offered seven nodes for players to explore. To capture these nodes, players would need to interact with the exhibit and the park’s in-

interpreters to acquire the information they would need to answer specific questions about the facility and its history.

Running during the summer of 2012, *Intelliphone Challenge* attracted twelve players, three of whom finished the game. fAR-Play tracked and recorded players' actions during game play,¹³ revealing that the shortest time to complete the game was seventeen minutes and the longest one hour and fifty-four minutes. Players spent an average of thirty-two minutes playing the game, and on average, captured nine of the twenty-one nodes. Discounting four players who began the game and captured one or zero nodes, game metrics show that players spent an average of forty-eight minutes playing the game.

In addition to game metrics, Fort Edmonton Park staff members gathered surveys from players once they completed the game, resulting in four surveys: three from those who finished the game and one from a player who did not.

Results

One of the earliest concerns players had was the game's log-in requirement. Although we initially included a mandatory log-in to be able to collect research data, players could not understand why they needed to provide personal information to participate. Based on player feedback, we decided to implement an anonymous log-in option, allowing players to log in to the game without an account. This change increased our sign-ups on the fAR-Play platform by 300 percent within two weeks, but the amount of nonanonymous user signups remained constant over the same period.

We were able to resolve players' concerns about log-in requirements early and quickly, but their other concerns were more complicated, mostly because they involved various elements of the game's interface. Wanting to have an attractive and flexible design, we originally designed the interface to look good on a laptop without taking into account how crowded the design would look on a mobile device. We also didn't take into account how lots of menus can be confusing on mobile devices. Unfortunately, while players said they enjoyed the game's content, they also noted that "the layout was confusing," "it was too complicated," and "it should be super simple. The content was great, but the format was difficult." The menus and navigation were particularly cumbersome. Players could not figure out how to accomplish certain actions in the game, and things did not always

function as expected; at times, players became so frustrated that they exited the application and started again.

While most players' comments concerned usability and not the look of the system, we decided to redesign the interface for fAR-Play to make it as simple and intuitive as possible. This allowed us to enhance players' access and experiences with the game. To help in this effort, we used Nielsen's ten heuristics for interface usability: visibility of system status, match between the system and reality, user control and freedom, consistency and standards, error prevention, recognition rather than recall, flexibility and ease of use, aesthetic and minimalist design, well-structured and easily discriminated features, and the use of default values.¹⁴ Recognizing that users would spend most of their time playing on a mobile device rather than a computer, we also decided to focus exclusively on the mobile interface and spent most of our time addressing the following elements.

Menu System: In the game's original iteration, we used a menu bar of simple text strings that gave many different options. Rather than making it easy for the user to find and engage with their information, however, the menu presented a cluttered and dense interface that made it hard for users to discover what they were supposed to be looking at and what each of the values meant. The game's new iteration offered a large drop-down menu at the top of the screen. When users tapped on the menu, a drop-down list would appear, allowing players to swipe and scroll through the options and select the page they wish to access. This option, making use of the ever-increasing screen real estate available on smartphones, gave users something substantial to tap with their finger, minimizing the chance that they would accidentally tap the wrong menu option during game play.

We also rewrote the menu item's copy, making it easier to understand each item in the menu; this was one of the reasons we changed the original game's use of "vPOI" (virtual point of interest) to "nodes."

Buttons: Like the old menu system, buttons in the games original iteration were simple text strings that users had to click on to accomplish a task. These strings were not only hard to press, but difficult to recognize as actual buttons, since they weren't well distinguished from text or hyperlinks. Although the new buttons take up substantially more space on the screen, they are easier to tap with a finger and their function is more readily apparent.

Loading screens: Another player concern was how difficult it was to know when the system was working on a task, especially in one area of Fort Edmon-ton that received poor cellphone reception. Users would click on a link and,

not realizing the system was loading, click on the link again and unknowingly restart the process; eventually the system would time out, leaving users frustrated. To alleviate this potential problem, we implemented jQuery mobile's page-loading system so that, when users tap on a function, a message would let them know that the system was in progress. If for some reason the page could not load, users would see a "Page could not load" error. This change not only let players know when something went wrong, but it minimized their need to decipher what, exactly, the problem was—unlike the conventional Error 404 (e.g., "not found") and Error 408 (e.g., "timed out") messages.

Lists and inputs: In the original iteration, clicking a link presented users with all information in that list at once. While this setup presented much information, it also meant that users who did not immediately recognize what they were looking for had to read through all of the titles and texts to discover the option they wanted. This situation was particularly noticeable in the game's list of adventures. Therefore, to help users find and select adventure options more efficiently, we implemented a unified list display; when users tapped on a list heading, they would get more information about that particular adventure, including its description and their current progress. Like other clickable options, we also made these list headings large and easily selectable. Another problem was user input boxes, such as the log-in screen. In addition to having difficulties clicking the small text-input boxes, users said it was unclear whether adjacent fields were related. Therefore, we implemented a minimal but graphically consistent box to help users identify related text fields. We also increased the vertical size of text boxes and made them large enough to occupy a substantial width of the screen to allow for easier selection.

Building the *Intelliphone Challenge* on the fAR-Play platform helped us understand more about the ways players participate in AARGs, as well as the design elements that can help or hinder players' access to and engagement with the information these games aim to present. But to learn more about design features that can enhance instruction-oriented AARGs, we decided to compare our experiences building *Intelliphone Challenge* with those we had building a geolocative game on a commercial platform.

Experiment 2: *Return of the Magic*

A collaboration between researchers from the University of Alberta and members of the Old Strathcona Business Association (OSBA), *Return of the Magic* is a geolocative game that aimed to promote the histori-

cal neighborhood and local businesses in the Whyte Avenue area over the Christmas season. Beginning our conversations in early 2011, we struggled to identify the most appropriate format, interface, theme, and content for the game. We knew that we wanted the game to be an event, and we believed that evoking and celebrating the charm of the Whyte Avenue area during the holidays would make that possible. We also wanted to offer more than just the extrinsic rewards typically linked to location games, and we believed that digitally overlaying a fictional world onto the physical world could provide a novel and thus intrinsic motivation for people to participate. We wanted to promote the historical neighborhood and local businesses in a fun, interactive manner, and we believed that a geolocative game that offers interesting content and riddles would encourage people to explore the avenue while simultaneously giving us an opportunity to study what makes such games entertaining for players. Ultimately settling on a reinterpretation of the paper-based scavenger hunt that OSBA ran the previous holiday, we decided to create a digital scavenger hunt with a holiday magic theme that would incorporate various elements from the area's vintage Christmas-window displays.

This decision offered several advantages. First, it gave us a nicely defined time to run the game; the game ran during the holidays, between November 2011 and January 2012. Second, OSBA gave us access to all of the materials it created for its 2010 scavenger hunt, giving us place-based trivia questions players could answer by visiting particular window displays (e.g., "In the Nativity Scene display, what color are Mary's eyes?") and saving us substantial development time. Third, even though few people completed the game the previous year, we thought people who were familiar with the area's window displays might be interested in playing a game connected to them. Finally, because Whyte Avenue visitors could play the game on their mobile devices, we thought people would enjoy the novelty of seeing their immediate surroundings with a digital layer, and thus remain actively engaged in the game.

We envisioned that the story behind *Return of the Magic* would focus on the story of two rabbits, Chloe and Chance, who must help Mrs. Magpie, an addled old bird, remember what she is doing on Whyte Avenue. By looking at the display windows on Whyte Avenue and visiting local businesses, players would get clues to help the rabbits answer the riddles Mrs. Magpie had rattling around in her head. As players help Mrs. Magpie remember

more about herself, they would come to realize that she is Edmonton's Christmas Spirit come to bring the holiday magic back to the city.

To build *Return of the Magic*, we decided to use SCVNGR (Scavenger), a mostly free, geolocative gaming platform that encouraged players to discover new places, find new things to do, and share their activity with friends on social media. Building a “game” layer on top of the real world, SCVNGR allowed developers to design challenges that would lead players to specific places; by going to these places and completing the challenges linked to them (e.g., answer questions, take a photo, leave a comment, scan a QR code), players would earn points, with results appearing on a leaderboard, that could be used for prizes.

The SCVNGR platform offered several advantages. One advantage concerned metrics. With SCVNGR we could track player participation universally within the game, at individual locations, and even more specifically, for each challenge question, providing valuable research data. We could also use these metrics to determine which players would receive the prizes local merchants were donating in exchange for having their business linked to the game. The main appeal of SCVNGR, however, was that it was a robust, easy-to-use platform. Developing the interface amounted to essentially filling out a series of web forms with the content we created. During development, we could also turn on or off the entire game or individual locations, making it easy to test and debug elements in the interface that might affect the game play experience, thereby minimizing our need to focus on technical problems in our overall assessment.

Despite these advantages, the fact that SCVNGR was a completely closed system introduced several obstacles—some of which were further complicated by context surrounding the game's inception and development—leading to discrepancies between the game we envisioned and the game we could ultimately produce. In the discussion below, we explain the challenges we faced and the ways we tried to address them. We then discuss reactions from players and merchant partners before turning to lessons for improving future locative game design.

One challenge we encountered is that SCVNGR was not built with stories in mind. Originally we planned to include a segment that introduced players to the story's characters and then use the description fields to continue and build on the story throughout the game. However, we discovered that all of the fields in SCVNGR are limited to 160 characters to be SMS-compatible. We tried a work-around that would use the built-in QR reader

or paste in a hyperlink to load a web page where we could establish *Return of the Magic's* larger narrative, but SCVNGR does not allow hyperlinking to content outside of the app because, as the developers told us, they wanted the platform “to be like a casino: once you get inside, they don’t want you to leave.” These limitations forced us to fit the story, character, plot, and motivation—some of which originally had hundreds of words—into 160-character chunks of descriptions, questions, and successful player responses that we could slot into the various SCVNGR fields. We eventually began creating “start” and “end” locations to give a little more context for this story and to give it a satisfying conclusion, but because the locations could be played in any order, there was no way to ensure players were accessing these locations as intended.

Another obstacle concerned the type of problems we could offer to players. In the early stages of development, we were working with eight businesses that would have window displays, giving us concrete elements around which to build the game’s challenges. To increase visibility for the larger area, however, OSBA invited all of its members to donate prizes for game participants in exchange for having their stores linked to the game. As the original eight locations increased to seventeen, we suddenly found ourselves needing to create more content to extend the story. The greater problem was that the recently added locations did not have window displays, forcing us to come up with other ways players could find these locations—a task that SCVNGR made more difficult.

The easiest way to help players navigate to and through various locations was by giving challenges players could complete by visiting display windows along the avenue. Yet without the benefit of a display window, businesses needed to become more creative in helping players find their particular location. One merchant, for example, agreed to challenge players to a game of mini ping-pong. Another suggested photographing the store’s resident cat, appropriately named “Bunny,” while another hid a stuffed penguin in the store’s freezer for players to discover. Unfortunately, players needed more context to understand these and the other challenges that were not linked to specific window displays—something that SCVNGR’s 160-character fields did not readily accommodate. Another option was for merchants to host the image of a game character, which players would then need to find and photograph. But even here, we encountered difficulties because SCVNGR restricts the size of the images players can upload, affecting the success of game play and diminishing the ways the game could use the avenue’s visual charm to enhance the players’ experience.

In addition to the design-oriented difficulties, there were also user-oriented difficulties. One such problem was helping prospective players learn of the game's existence. SCVNGR assumes that players would first find the SCVNGR app and then serendipitously discover the content, but without any user-generated content in Edmonton, there is little reason for anyone to have the SCVNGR app in the first place, requiring us to promote the app and our scavenger hunt simultaneously. Fortunately, OSBA was already promoting the Christmas displays, so we could piggyback some of our promotion onto OSBA's, and we landed a brief spot on a local television show. In addition, we distributed postcards to businesses in the area and put up posters in all of the participating businesses' windows. There were also articles about the game in the local newspaper, in the OSBA newsletter, and on the OSBA web page; OSBA also tweeted and posted on Facebook about the game. But despite these efforts, only two players noted in their postgame surveys that they had heard about the game from print media; most players learned of the game online and through word of mouth.

Then there were the challenges of simply getting players started. Like the individuals playing *Intelliphone Challenge*, those playing *Return of the Magic* faced barriers before they could participate in the game, and these barriers may have affected who ultimately played the game. For example, to begin, users needed to own a smartphone, be savvy enough to know how to download the app from the app store or marketplace, be willing to create a SCVNGR account, and then navigate the SCVNGR menu to the game environment we had built. Although our data did not account for this information, we do wonder how many people who may have been interested in playing the game abandoned it simply because of the hurdles they faced to get started.

These were challenges we encountered before the game actually began, and the game metrics and player surveys, which we discuss in the next section, indicate the impact these and other challenges may have on players' ability to access, participate in, and learn from AARGs.

Results

To identify elements designers should consider when developing geolocative games, we used game metrics, player surveys, and merchant-participant surveys.

Game Metrics

Using the metrics that SCVNGR recorded, we were able to see how players moved through and interacted with their environment. For example, we found that, collectively, players started 541 challenges, and completed 311 of them. Looking more closely at the data, however, we can learn more about those completing the challenges.

To identify player engagement, we classified users into one of four categories, based on the number of challenges they played: “curious” (played 0–3 challenges), “casual” (played 4–7 challenges), “avid” (played 8–20 challenges), and “hard-core” (accomplished all 21 challenges). According to game metrics, of the fifty people who downloaded the game during its six-week run, there were thirty curious players, nine casual players, seven avid players, and four hard-core players, suggesting that most of the challenges may have been completed by avid and hard-core players. In contrast, more than half of the people who downloaded the game stopped after only a few challenges, suggesting that something stopped them from continuing.

The metrics also reveal where players moved during the game. The business with the most visits recorded thirty-three unique players, and the business with the fewest visitors had three players. Stores that had window displays (eight of seventeen participating businesses) generated the most challenges completed (49 percent), even though these locations provided less than a quarter of the challenges in the game. These data seem to indicate that businesses not having window displays, thereby requiring players to complete challenges indoors, may have created more barriers to player involvement.

Player Surveys

In addition to game metrics, we used player surveys to learn what motivated their participation. We contacted all registered players, giving them a survey with a series of Likert-scale questions (using a scale of 1–10), as well as multiple-choice and open-field questions that asked participants about their reasons for playing and their levels of enjoyment. Of the fifty registered players, thirteen returned their survey. Of this group, four were from our curious players category, two were the casual player category, three were from the avid player category, and four were from the hard-core player category. Although the response rates skew data toward avid and hard-core players, the surveys offer some interesting results about the

reasons people played and the ways they played. For example, to identify what motivated the respondents to play the game, we provided both a multiple-choice (which allowed players to select more than one answer) and an open-field option. Among the multiple-choice options, the highest-rated answer was “fun,” generating eight responses, followed by “interest,” which earned seven votes. Four respondents named “potential prizes” as their primary motivation, but the same respondents also cited an interest in learning more about the Whyte Avenue area. Only three players were interested in the game story’s ending, which does not necessarily negate developers’ need to build their game’s brand; it could simply suggest that some players may not find finishing the narrative a motivation in itself for game play.

The surveys also revealed that players considered the game a social activity, and an opportunity to participate in a new activity. Despite the game being played on a smartphone, which is generally thought of as a solitary gaming platform, only one person played the game alone; most played in teams, either by using one shared phone, or by playing together using phones that each had its own account. Although nine of the thirteen respondents indicated that they had never used a location-based app before playing *Return of the Magic*, most noted that the interface was easy to use. Furthermore, in response to the question asking whether they would recommend *Return of the Magic* to others, four players marked “strongly disagree” and five players marked “strongly agree.” But when asked if they would be interested in playing a similar game in the future, eight players strongly agreed with the statement, and only two players strongly disagreed, suggesting that different or additional geolocate games could find an audience.

Merchant Surveys

To identify the best ways for developers to engage community partners in geolocate projects, we distributed two surveys: one before the game ran and one after the game ran. The following discussion explains each survey and its findings. The first survey, distributed at the start of our collaboration with OSBA, was to learn more about the businesses that might be affected by our game’s development and to gauge these businesses’ expectations of location gaming in general. In fall 2011, we emailed surveys to local businesses asking for their thoughts about the responsibility businesses have to the community, the place of technology in their own business, what a

location game should offer its players, and what motivations they thought players might have to participate. The OBSA marketing manager and thirteen merchants replied. According to the surveys, all of the merchants were independent business owners. These merchants almost unanimously supported community events, and most noted that they were at least partly responsible for the well-being of the entire neighborhood. The surveys also suggested that, while merchants varied in their technological abilities, almost all of the respondents felt that technology was an inescapable part of doing business, even if they were still trying to figure out the best ways to use it and maintain a more “human touch.”

For survey questions asking about location games, the general sentiment was that the game could offer an unconventional way to promote the community, entice customers into stores they would not ordinarily visit, and increase participant enjoyment. Merchants also liked the idea that they could promote their business with minimal expense (that is, businesses simply needed to donate a product or service that would serve as a prize for players, and their business would appear in the game).

When asked about what they thought would motivate someone to play a location game, like us they tended to overvalue the draw of prizes. To their credit, however, the merchants’ highest response was that the game’s inherent fun would be the primary reason people would play, even though they could not agree on what, exactly, would make the game fun or what long-term value that fun might create. Most business owners also noted that foot traffic would be their primary measure of the game’s success, with a few notable exceptions; one merchant, for example, listed “my customers having a fun time” would be his key measure, while two others mentioned the game’s ability to generate a buzz in the community. We used these surveys to select seventeen businesses that could best accommodate the game’s setup. We also used these surveys to create and refine the game’s challenges and larger narrative in ways that could best promote merchant goals.

Once the game was over, we emailed a postgame survey to each business that served as a location in *Return of the Magic*. Adapting questions from the pregame survey for postgame assessment, this survey asked merchants to comment on the extent to which they thought the game was successful. Ten of the seventeen merchants responded. On the pregame survey, merchants said they liked the possibility of the game bringing more people to their store—even if they didn’t buy anything—but the postgame survey revealed business owners’ general disappointment with the lack of foot traffic that moved through their stores. Many were also disappointed by the lack

of promotion for the game, suggesting that it was partly to blame for the lower-than-expected customer traffic. In response to the question asking why people would be motivated to play a game like this, merchants gave different answers. Merchants listed prizes (four times), fun (three times), exploring (three times), and social engagement (two times); four merchants specifically said they “didn’t know.” Although our merchant partners were disappointed by the final results, most said they would be willing to participate in another such game in the future.

Before launching *Return of the Magic*, we had no idea how many players to expect, and when only eight people tried the game during its first week, having fifty people try the game during its entire run seemed like a moderate success—a sentiment the surveys seem to echo. And while representing only a portion of those who participated, many of the players who did respond said they had fun, discovered something new on the avenue, and would be interested in participating in more games like this in the future. Together, these results suggest that location-oriented games can be tools for learning.

Conclusion

We believe that geolocative games can be a fun way to motivate new behavior, and these two experiments reveal several considerations for designing locative games that can be successful for all participants.

First, players seem to treat locative games as social experiences, so these games should promote group interaction. Although solo players can enjoy following the hunt and solving a riddle, the best experiences seem to come from playing with friends, perhaps because players find it easier and more enjoyable to play within a space where play is foreign when they are not the only one doing it. By providing opportunities for players to work with or even compete against others, designers can take advantage of the social affordances of locative games.

Second, players rely on easy-to-use interfaces, so games should offer clean, intuitive designs. As we learned from our work on *Intelliphone Challenge*, the way players see, understand, and interact with a game’s elements can affect not only how, but if players are able to participate as intended. By streamlining and simplifying the technical elements a player must use and navigate, designers can create more intuitive interfaces that let the player focus on the game and not the technology.

Finally, players seem to enjoy locative games that offer a combination of

intrinsic and extrinsic motivation, so game creators should know their target audiences and find ways to cater to their interests and values. This work may be particularly challenging in AARG development. After all, whereas most video games create fun by carefully controlling the mechanics and conditions of play, locative games cannot control either the environment or the players' freedom in the same way; instead, they must rely on players being able to create their own fun between the moments players start and stop the game. Nevertheless, locative games can generate and sustain player interest by incorporating elements that promote discovery: little-known information, materials that offer clues for a puzzle or directions to the next discovery, feedback to help players know when and how they are successful. Here, then, is why knowing the target audience matters. For example, when designing *Return of the Magic*, we believed the primary audience would be young children, when in fact the primary players were adults. Had we known this, we could have generated content that would have engaged more of the players for longer than our game did. By considering target audiences and what might motivate them in particular contexts, designers can focus the content of the AARG.

AARG technology is still evolving, making it challenging and time-consuming to create locative games, but our experiments highlight some of the ways locative games can inspire engagement and new behaviors in their participants.

Notes

1. Whereas alternate-reality games are conventionally thought of as games that present a fictional world alongside the real, and augmented-reality games typically refer to games that add a layer of information to our understanding of the existing world, almost universally both categories rely on multimedia. This can involve cameras and screens that overlay digital information onto real-world environments or create new environments, including text, sound, still images and graphics, video, GPS data, and other elements that together construct and advance alternate yet contemporaneous narratives. Recognizing the overlap between these two forms of games—especially when such technologies are incorporated into the game-play mechanics—we use the pirate-like acronym AARGs (for alternate- and augmented-reality games) to discuss projects that combine alternate and augmented gaming elements.

2. Clark Abt was an early advocate of learning through play. In *Serious Games*, originally published in 1970 and reprinted in 1987, he argued that “serious games offer us a rich field for a risk-free, active exploration of serious intellectual and social problems. In games man can once again play the exciting and dynamic roles he always enjoyed before society became so compartmentalized” (pp. 13–14).

3. For a discussion on historical thinking concepts and how they can be taught through augmented- and alternate-reality gaming, see the chapters by Compeau and MacDougall, and Kee, Poitras, and Compeau in this book.
4. Andrea Phillips, "Methods and Mechanics," in *2006 Alternate Reality Games White Paper*, ed. Adam Martin (2006), 36–37; Jeffrey Y. Kim, Jonathan P. Allen, Elan Lee, "Alternate Reality Gaming," in *Communication of the ACM* 51 no. 2 (2008): 36–42.
5. See <http://www.argn.com/> for a more complete list.
6. Andrea Phillips and Adam Martin, "Business Models," in *2006 Alternate Reality Games White Paper*, ed. Adam Martin (2006), 50.
7. Montola 2005.
8. Christy Dena, "ARGs and Academia," in *2006 Alternate Reality Games White Paper*, ed. Adam Martin (2006), 58–78, <http://www.igda.org/arg/resources/IGDA-AlternateRealityGames-Whitepaper-2006.pdf>
9. Phillips, 36–37.
10. Bryan Alexander, "Antecedents to Alternate Reality Games," *2006 Alternate Reality Games*, 9–14, <http://www.igda.org/arg/resources/IGDA-AlternateRealityGames-Whitepaper-2006.pdf>; David Szulborski, *This Is Not a Game: A Guide to Alternate Reality Gaming* (Lulu, 2005).
11. Jane McGonigal, *Reality Is Broken: Why Games Make Us Better and How They Can Change the World*, (Penguin Book, 2011); Phillips, 36–37.
12. This project was supported by the GRAND Network of Centres of Excellence (grand-nce.org) and the University of Alberta. A number of people worked on the development of the fAR-Play platform and games, including Shannon Lucky, Joyce Yu, Calen Henry, Aiden In, Lucio Gutiérrez, and Shayne Mugford.
13. This was done with players' consent on their signup for the fAR-Play platform.
14. Jakob Nielsen, "Heuristic Evaluation," in *Usability Inspection Methods*, eds. Jakob Nielsen and R. L. Mark (New York: John Wiley & Sons, 1994).

Tecumseh Returns

A History Game in Alternate Reality, Augmented Reality, and Reality

Timothy Compeau and Robert MacDougall

Can students and public audiences be trained to see the past like historians? Timothy Compeau and Robert MacDougall share their experiences crafting and running two educational games designed to teach historical thinking skills. Using alternate-reality gaming techniques and augmented-reality print, Compeau and MacDougall helped students peel back the “authoritative” layers of history to see how historians construct history from the fragmentary remains of the past.

In September 2011, in the first week of the new school year, about two dozen students in and around the Public History program at the University of Western Ontario received an unusual email. “My name is Sophia,” it began. “I got your name off the UWO History website, and I was hoping you or your friends could help me.” Most of the students deleted the email without a second thought, but a few followed the attached link to an even more unusual video. The video showed a man, dressed in a Napoleonic-era uniform, apparently locked in a small room. The man, identified only as Captain Smith, seemed to have no memory of how he got there, or indeed of any historical events from the last two hundred years. But he, too, asked for help.

Clues in the video and on Captain Smith’s uniform led the students to another website, and then to the university library, where, in dusty old

books about the War of 1812, they found torn scraps of paper bearing cryptic notes about the Shawnee leader Tecumseh and the war. Soon the students were working together and also communicating by email and social media with a number of mysterious characters, both malevolent and benign. They found themselves in a race against time with a shadowy conspiracy—is there any other kind?—researching history and racing around southwest Ontario to museums, graveyards, and historical sites. Classes in the Public History program at Western came to a standstill as the Tecumseh mystery became all-consuming. Students stayed up all night decoding nineteenth-century poetry and debating who among them might be in league with Smith’s nefarious captors.

Of course, there was no “Sophia” or “Captain Smith” really. The students were playing *Tecumseh Lies Here*, an alternate-reality game or pervasive game designed to teach historical research methods while exploring the history of the War of 1812. We designed *Tecumseh Lies Here* as an experiment in history pedagogy and subversive commemoration, and we sprang it on the students in our program as a test. The first iteration of the game ran from September through October 2011; both the students and the game passed this test with flying colors. *Tecumseh Lies Here* was a fun, engrossing, and educational experience that exceeded almost all of our expectations. It was undoubtedly one of the most memorable teaching experiences of our careers. And we unanimously agreed we would never do it again.

This first iteration of *Tecumseh Lies Here* took over our lives for more than a month. We missed deadlines and neglected our other duties while we stayed up to the wee hours crafting new content for the game, desperately trying to stay one step ahead of our surprisingly dedicated group of players. And when it was all through, we had a fantastic story to share, but it would be difficult, if not impossible, to repeat the experience—not without months of lead time and another research grant. Therefore, the 2011 alternate-reality game became the first of several experiments with different media to see if we could design games and playful exercises that made the act of doing history—exploring historical sources, evaluating competing interpretations, and composing narratives and theories—the central focus and propelling force of the play. Rather than simply inject historical themes into preexisting game genres, as so many historical games do, we wanted to create an experience built around inquiry-oriented pedagogy, where students learned history by *doing* history. We hoped to inculcate historical thinking skills, such as grasping different perspectives and recog-

nizing the biases inherent in primary sources.¹ This chapter explores our efforts to use new techniques and technologies to design and run games that teach historical thinking in this way.

While our 2011 game was an intense and rewarding experience, we wanted to streamline it into a more repeatable version that could be used by teachers in their classrooms or as a model for museums and other heritage institutions. We created two new versions of *Tecumseh Lies Here*: an as-yet-unpublished “modular” version of the 2011 game, and an augmented-reality “un-textbook,” which we launched to coincide with the autumn 2013 bicentennial of the death of Tecumseh. Though the form of each game was quite different, the basic premise remained the same. By applying a playful approach, we wanted to help our students learn about and engage with the history of the War of 1812, while also teaching them how professional historians piece together evidence to form narratives and arguments. In each version of *Tecumseh Lies Here*, our intention was also to create a sort of subversive commemoration that challenged simplistic interpretations of the War of 1812 and questioned supposedly authoritative accounts of the past. In this chapter we will consider the strengths and weaknesses of these different iterations, and the unique affordances each type of game provides.

Alternate-Reality Games

Building an ARG is like running a role-playing game in your kitchen for two million of your closest friends.

—Sean Stewart

We spent many years throwing, essentially, rock concerts. Very large, real-time, elaborate experiences that were really cool and really fun for the people who were involved with them. . . . But you had to be there.

—Jim Stewartson²

We began development on *Tecumseh Lies Here* in 2009 with ambitious goals explained in our chapter of *Pastplay: Teaching and Learning History with Technology* (2014).³ Our intention was to explore how the fascinating new genre of alternate-reality games could be harnessed and adapted for history education. At first we wanted to create an immersive experience that could attract an international, multilingual group of players and enlist the emergent collective intelligence of crowds to solve complex historical puzzles. This vision was based on the ARGs of the early 2000s: games such

as *The Beast* (2001), *Chasing the Wish* (2003), and *World Without Oil* (2007). Perhaps most famously, the alternate-reality game *I Love Bees* (2004), designed and run by a team including Jordan Weisman, Sean Stewart, Jim Stewartson, and Jane McGonigal, exemplified the standard ARG form. Commissioned by Microsoft as a viral marketing campaign for the sci-fi video game *Halo 2*, *I Love Bees* eventually attracted thousands of players who, working together, cracked what seemed like nearly impossible puzzles while also discovering and building the game's story. An ARG is built by piecing together aspects of the narrative and the puzzles across a variety of media—from websites to email to instant messaging and even packages in the mail. As Sean Stewart explains, “we . . . present the evidence of that story, and let the players tell it to themselves.” He refers to this as “story-telling as archaeology,” an ideal concept for helping students explore how history is crafted from the clues of the past.⁴

After receiving a grant from the Social Science and Humanities Research Council of Canada, we set out to design a similar game that would teach history and historical research skills.⁵ With the War of 1812 bicentennial approaching, the topic seemed ideal, and we soon zeroed in on the mystery surrounding the life and death of the legendary Shawnee leader Tecumseh. His story and the history he represents were perfect examples of how mysterious and contested history can become.

Tecumseh was a charismatic orator and visionary leader who attempted to form a confederacy of the various tribes and nations of the indigenous peoples from the Great Lakes to the Gulf of Mexico, in a bid to stop the relentless westward expansion of American settlement in the early nineteenth century. Along with Tecumseh's message of political solidarity, his mysterious and often misunderstood brother Tenskwatawa—“The Shawnee Prophet”—spurred a religious revival among the First Peoples of the old Northwest, which called for the rejection of European ways and a return to traditional Native practices. This fusion of a political and religious ideology in the movement deeply worried the American government and frontier settlers. Tecumseh's intransigence and rejection of some important land cessions, along with increasingly plausible American suspicions of clandestine British encouragement of the northwestern tribes, were two of the most significant causes of the War of 1812.⁶

With the outbreak of hostilities between the United States and Britain in June 1812, Tecumseh forged an alliance with the British in Canada. After a series of stunning victories, Tecumseh became a legend—hailed as a savior by the Canadians and British, and respected and revered throughout

his fledgling confederacy. To Americans, Tecumseh became both a terrifying specter of Native vengeance and a romantic image of the “noble savage” fighting bravely, even if in vain, for the freedom and ancestral lands of his people. Whereas Tenskwatawa was routinely depicted as cowardly and deformed, a conniving trickster who used Native superstition to empower himself, Tecumseh was remembered as brave, handsome, and noble, a warrior who sacrificed his life for his people. The grudging admiration for Tecumseh among his enemies is plain to see in the number of American towns, statues, and schools bearing his name; white parents even named their children after their former enemy, including the parents of Union general William Tecumseh Sherman, born in 1820. The attitude of nineteenth-century white North Americans to the Shawnee chief is perhaps best summed up by the Canadian novelist and poet John Richardson, a British officer who fought alongside Tecumseh at several engagements, including the Battle of the Thames where Tecumseh was killed. “Ever merciful and magnanimous as he was ardent and courageous,” Richardson writes, “. . . his heart was formed to glow with all the nobler and more generous impulses of the warrior.” He was the epitome of the noble savage imagined by European settlers, a leader who “was a savage; but a savage such as civilization herself might not blush to acknowledge as her child.”⁷

Tecumseh was killed at the Battle of the Thames on October 5, 1813, about seventy-five miles northeast of Detroit in present-day Ontario. The exact circumstances of Tecumseh’s death have never been conclusively explained. At least two different American fighters took credit for killing him, sparking a simmering debate between old soldiers that went on for decades. More important, and more mysterious, was the fact that no one knew what became of his remains. Stories circulated that vengeful Kentuckians had mutilated his corpse and carved him into grisly trophies. Others claimed that his body was misidentified by the Americans and that Tecumseh’s followers actually recovered his body and buried him in a secret grave now lost to history or known only to a select few. Some refused to believe that Tecumseh had died at all. For the next century, amateur historians attempted to track down Tecumseh’s remains, perhaps in a misguided attempt to venerate the fallen hero or simply to solve the mystery. Grave hunters identified several sites purported to contain Tecumseh’s bones, and the remains of at least one unfortunate casualty from the battlefield were exhumed and paraded for public inspection during the 1840 presidential election campaign of William Henry Harrison, the American commander at the Battle of the Thames. This ghoulish episode was part of a long tradi-

tion of white defilement of Aboriginal burials; it is also a physical manifestation of the many attempts to appropriate the legacy of Tecumseh for political or nationalistic purposes.⁸ The use and abuse of Tecumseh's life and legend provides a compelling object lesson in the making of history. Our title, *Tecumseh Lies Here*, was a sort of dark pun on the Tecumseh mystery, and the lies and distortions that have been attached to his legacy. Nobody knows where Tecumseh lies, but lies about Tecumseh are everywhere.

With these historical controversies in mind, we began designing *Tecumseh Lies Here* in 2009 with a small team of graduate-student researchers. Our planned narrative for the game focused on three competing factions of historical enthusiasts whose quarrel over the meaning and ownership of the War of 1812 had spun out of control and now involved a bizarre kidnapping. Each faction represented a different way of thinking about the war and about history in general. There was a New Age spiritualist group who considered themselves spiritual heirs to Tenskwatawa and his struggle; a group of 1812 re-enactors with a strongly nationalist, pro-Canadian reading of the war fighting against the forces of historical revisionism or "de-enactment"; and finally, a nefarious conspiracy of professional historians altering the historical record for their own murky goals. In the end, of course, none of these organizations had any real claim to Tecumseh's legacy, and all three twisted and distorted the historical record to fit their own ideologies. We built websites for each group, laying out their manifestoes and their back stories, but also embedding each site with hidden messages, clues, and password-protected areas. These websites served as the trailheads for our fictional plot and as entry points into the true history engaged by the game.⁹ We planned to drop our players into the middle of this fictional power struggle just as things began to turn hot, and the leader of the re-enactor organization ("Captain Smith") went missing. It would be up to our players to retrace his steps, reconstruct his research, and figure out what happened to him.

As months of development quickly slipped by, we realized that our initial goal of a cross-border, multilingual experiment was reaching too high. Instead, we ran a "beta test" of *Tecumseh Lies Here* in fall 2011 for a manageably small group of twenty to thirty players from the University of Western Ontario and the surrounding community; many of our most active players came from that year's cohort of Western's Master's Program in Public History. In the first week of classes, we asked our colleagues in Public History to announce that we were looking for volunteers to test a new type of historical game and to forward the email addresses of any willing students. We

also invited other friends, colleagues, and students to join the game. About ten days later, all those willing received a strange, personalized email—the “rabbit hole” into the Tecumseh mystery and our game.

One thing that makes an alternate-reality game compelling and immersive is its lack of obvious rules, directions, or parameters. The players were left to figure out what to do from this point. Nobody gave them direct instructions. Eventually, they began to Google small hints provided in the five-minute video that led them deeper into the scattered pieces of the story. Our fictional characters had names, back stories, group affiliations, photographs, email and Twitter accounts, and even a cellphone. The unclear rules or boundaries of the game added to the sense of immersion and even some paranoia. The students playing the game began to wonder if their professors were behind it, and some began to suspect that certain classmates were actually moles working for the game designers. (They were not.)

It would be difficult to retell the entire narrative of the game. In giving talks and presentations about *Tecumseh Lies Here*, we have been forced to admit that there is a “you-had-to-be-there” quality to alternate-reality games, which is almost impossible to convey after the fact. More-experienced game designers than us have struggled to describe the form and player experience of their ARGs.¹⁰ But a few details may suffice to suggest how the ARG genre can help students explore history.

Within a few days of receiving the initial message and video, our players had unravelled enough clues to begin building the story and tracking down further clues and puzzles. They determined that the imprisoned amnesiac was a man who went by the name Captain Smith, who led a reenactor troupe based on the Independent Company of Foreigners, a notorious British army unit made up of French deserters and sent to fight in North America. They also learned that Smith had been searching for the true resting place of Tecumseh so that there could finally be a European-style monument constructed on that site.

With every puzzle our players solved, a new one would appear, and a new element of the story would begin in the form of videos, blog posts, or other media. Several larger puzzles went on throughout the game, such as continual discovery of wheels from an eighteenth-century decoder wheel that the players would need to decipher an important letter. In one of the initial puzzles, players encountered a scrap of paper with an encoded message written in Shawnee. When deciphered, the note revealed a call number to a copy of Edward Eggleston’s *Tecumseh and the Shawnee Prophet* (1878)

in the D. B. Weldon Library at the University of Western Ontario. Inside the book, at important pages, the players discovered torn sheets from Captain Smith's research diary. These led them to pages in other books also pertaining to Tecumseh, and helped them to not only assemble a map of a local museum, but to explore several key works in the historiography. The players were on their way to discovering why Captain Smith had been kidnapped and how they might free him. They had also achieved a crucial step in historical investigation, by becoming acquainted with the secondary literature. Over the next few weeks, the players worked through more puzzles and tasks that introduced them to copies of primary sources and took them to a variety of actual historical sites. One of the unplanned by-products of the game was that it introduced several newly arrived students to the London, Ontario, area and also served as a powerful team-building exercise for that year's Public History cohort.

We communicated with players via email, Twitter, and the dedicated cellphone, playing the roles of various allies and antagonists. The sheer number and complexity of individual interactions needed to run an ARG made us grateful that we had not attempted anything more ambitious. Meanwhile, the players discussed the puzzles and their progress on Twitter, which made it easy for us to follow their work, intrude on their discussions with occasional tweets and messages from in-game characters, track their progress, and to some extent, track their learning. Many of the students became deeply, powerfully engaged with the game. Several claimed to be genuinely frightened of potential encounters with the villain's henchmen. They also inadvertently alerted us when they had solved a puzzle and were rushing to the revealed destination a day before we believed they could have possibly succeeded. This left us racing to alter game elements and weaving through rush-hour traffic to get the next clues in place before the players arrived. We knew to increase the pace of the game when one player tweeted, "This is boring," and we were kept in the loop when aspects of the game failed, such as when the players erased the contents of a pivotal flash drive by guessing the wrong password too many times.

One of the game's final tasks involved opening a wooden box locked with a GPS device that could only be opened after taking the box to several physical locations significant to Tecumseh's life. One of the most satisfying moments for us was learning from the players that the box had indeed opened to reveal the final pieces of the mystery at the exact spot where Tecumseh was purportedly killed. The contents of the box led our players to the final stage of the game, which took place at Fanshawe Pioneer Vil-

lage in London, Ontario, amid several hundred costumed 1812 reenactors during a living history weekend. Using the clues and information they had discovered throughout the game, our players located Captain Smith in the flesh, freed him from his captors, and completed the game, which was followed by a reception in an authentic nineteenth-century tavern.

It took our players over a month to perform all of the tasks and solve the complex puzzles involved in the game. Rather than be put off by the game's difficulty, the players reported that the challenges actually enhanced the group dynamics. As one of our players reported, "The intricacy and range of the ARG's content appealed to many learning styles and research strategies, giving everyone in the group their moment to shine. In order to solve clues and move on to a next round, it questioned the 'acceptable' way in which we traditionally interact with historical discourse—it challenged the way we viewed history and its creators."¹¹ We were certainly pleased by that sort of testimony and by the immensely enthusiastic reaction of our players to the game. By happy coincidence, the Public History program's group project that year was the creation of a GPS smartphone tour application for the "Tecumseh Parkway," which guides tourists from Windsor to London exploring historic War of 1812 sites along the way. Much of the material our players collected for *Tecumseh Lies Here* was put to use in creating this app.¹² By all reports and by almost every measure, we considered our "beta test" of *Tecumseh Lies Here* to have been a smashing success. Why then did we not want to repeat or replicate the game?

Our budget was modest. After our initial grant had been spent on graduate-student researchers, little money was left to actually run the game. But the most limiting factor for us was not a lack of money, but of time. All of the game designers, actors, and facilitators were also busy with teaching, research, and writing (not to mention life), and the time we could devote to this demanding enterprise was limited. An alternate-reality game depends on an engaged group of players, otherwise the plot will not advance or even really take shape. It is not a genre for those seeking to be passively entertained. We were lucky that our players were very keen, because without them the game simply could not have happened. But when the players are indeed engaged, it means that the game runners, or "puppet masters," have to keep an eye on events around the clock. The game happens when it happens. At more than one point, this meant rushing out at midnight to record a new video so that the players would wake up with new material and another challenge, or scrambling at 6 a.m. when the locking mechanism on the GPS box broke. We were

constantly on the phone with each other debating the difficulty of our puzzles and what twists should come next. We had to stay one step ahead of the players, but also ensure that the game's objectives were achievable. And when our players took the narrative to unexpected places, we had to rush to alter our plans. Added to that were the dozens of tweets, emails, text messages, and blog posts we needed to generate in the personas of the various characters, which are needed to create the immersive quality of an ARG. It was exhausting.

So while we enjoyed the work and were very happy with the outcome, we did question the reproducibility or viability of this form of educational game. In the end, our team included seven people, many of whom volunteered their time, working very hard to produce a game for about twenty active players, with maybe another dozen occasionally chiming in. Those are thin margins. We questioned whether the work involved in designing and running an ARG would be feasible for most history educators. Could *Tecumseh Lies Here* be "scaled up" to reach larger audiences or be reused again and again? Would other university professors, K–12 history teachers, or museum staff want or be able to reproduce this effort? Would it make sense for them to do so? "Playing in the 'real world,'" we wrote prior to running the game, "means accommodating real-world constraints on budget and time. A pedagogical idea that cannot be employed in actual educational institutions, by individual teachers and professors, by small museums and heritage sites, by people on the front lines of history education, is unlikely to take root."¹³

Finally, we found it hard to preserve or reproduce the gaming experience after the fact. One of the most difficult aspects of the game was explaining to people in the aftermath what had happened and what it was all about. We had many people ask us if they could play the game, or experience a demonstration, and we could not simply provide either. We could show blog posts or strings of tweets, but explaining the original context or what had been happening in the game at that point usually involved convoluted and tedious explication. We met with our players and collected their thoughts and even snippets of video they took, and some joined us in giving talks about the game. But they had as much difficulty as we did in explaining what had been so compelling about the experience. In the end we often found ourselves repeating a line that had frustrated us when reading about other memorable ARGs: "You had to be there," prompting us to ask whether it was time to put *Tecumseh Lies Here* to rest.

Subversive Commemoration

The new warriors have mobilized a formidable, long-term project which, if successful, could change the country beyond recognition. . . . Appealing to atavistic and violent conceptions of blood and soil; proudly flourishing the age-old symbols of empires and cherishing as heroes their often violent partisans; . . . the campaign aims to supplant any vestige of that modest, imperfect but promising experiment called Canada.

—Ian McKay and Jamie Swift, *Warrior Nation*¹⁴

And yet we found that we could not escape *Tecumseh Lies Here*. We had enjoyed running the game immensely and kept looking for ways to capture the same intensity in our other teaching. We remained firm believers in playful historical thinking and in teaching history through play. And we found ourselves seeing traces of Tecumseh everywhere—reminders of his life, his struggle, and the ways his history has so often been appropriated and misused. Like our players, we had become “paranoid in a good way.” These feelings only intensified with the arrival of the War of 1812’s bicentennial in 2012.

Canada’s federal government spent some \$28 million commemorating the two-hundredth anniversary of the war, with television commercials, parades, commemorative coins, dozens of museum exhibits, hundreds of historical reenactments, and a new national monument in Ottawa. As pleased as we were to see this money being spent on public history—and this by a government that removed almost all federal funding for provincial and community archives, made deep cuts to the library and archives of Canada, and eliminated Canada’s long-form census—we were uneasy with the tenor of the celebrations and disappointed by the quality of some history on display.¹⁵

It was not only that the government’s presentation of the war was simple and uncritical, though of course it was. The theme of the official celebrations was the idea that, in 1812, a diverse population of English, French, and Native peoples had united as Canadians to defeat the American invaders. “We stood side by side and won the fight for Canada,” said the commercials and the posters.¹⁶ This narrative is debatable—to raise just one objection, in 1812 there was no nation known as “Canada” to defend—but that in itself was not especially surprising or alarming. Public history must often be painted in broad strokes. We were more troubled by the way the enthusiastic celebration of 1812 fit with a broader project by the Ste-

phen Harper government and the Conservative Party of Canada to make military narratives and symbols central to Canadian identity, in what critics have seen as an effort to “rebrand” Canada as a “warrior nation.”¹⁷

Most problematic of all was the way the “Fight for Canada” narrative distorted Native history and the story of Tecumseh in particular. To secure Native aid in 1812, the British promised Tecumseh and his followers a large swath of territory, from the Ohio River to the Great Lakes, in which to create an independent Native state. Tecumseh’s warriors played a pivotal role in the northern and western theaters of the war; the British could probably not have held Upper Canada without them. But it is a dubious claim to say that the Native confederacy fought “for Canada” or even to say that they stood side by side. Tecumseh himself had never been to the territories known as Upper and Lower Canada before the war. He spent only a few weeks in what is now Canada, while retreating from Detroit, and fought only one battle there: the one in which he lost his life. Even before his death, Tecumseh and his followers felt betrayed by the British, who surrendered the forts and territory Tecumseh’s warriors had captured at Michilimackinack and Detroit. And the British promise of an autonomous Native territory died, if not with Tecumseh in 1813, then with the Treaty of Ghent that ended the war. No provision was made at Ghent for Native lands, and peace between Britain and the United States allowed white settlers from both nations to push westward. After Tecumseh’s death, British North Americans, and later Canadians, came to lionize him as a martyr for the British Empire, but that perception elides Native goals in the War of 1812 and how dearly the conflict cost Tecumseh’s people in the end.¹⁸

The 1812 bicentennial also coincided with the eruption of the grass-roots protest movement known as Idle No More. Idle No More was triggered by a wide range of issues affecting indigenous people in Canada, but at its heart was an argument about history: whether the relationship between Canada and its Native peoples is that of a government and its individual citizens, as the government of Canada maintains, or that of parties to a treaty, as Native leaders such as Attawapiskat Chief Theresa Spence insist. Idle No More’s central demand was for the Canadian government to “live the spirit and intent of the treaty relationship.”¹⁹ In this context, the question of whether Native warriors who fought in 1812 were “Canadian” or “fighting for Canada” is not just semantic or academic. It is contemporary and real.

Thus by late 2012, the prediction we made in a grant application four years before—that the 1812 bicentennial would promote a banal sort of

nationalism—seemed even more true than we had expected. And our call for a “subversive commemoration” that honored the real complexity of the war felt more urgent than ever. Perhaps Tecumseh and Tenskwatawa were not done with us yet. We set our minds to what we saw as the two big problems of the alternate-reality game: the effort required to run a game for even a modest audience, and the transient nature of that experience. We talked about “scaling up” *Tecumseh Lies Here*, by which we meant both reaching a larger audience and making game elements that remained usable and educational after the initial period of play. How this was to be done, however, was not clear.

The D&D Module

What lies ahead will require the use of all your skill, put a strain on your imagination, bring your creativity to the fore, test your patience, and exhaust your free time. Being a DM is no matter to be taken lightly!

—Gary Gygax, *Dungeon Master’s Guide*²⁰

Our first idea for scaling up *Tecumseh Lies Here* was to design a kind of kit that would allow others—history teachers, or museum or library staffers—to run their own version of the alternate-reality game for their students or patrons. We called this plan “the D&D module,” after the adventure scenarios available for *Dungeons and Dragons* and other tabletop role-playing games. Running such a scenario is necessarily a mix of beforehand preparation and in-the-moment improvisation. The prewritten scenarios for such games—historically called “modules,” though this is an idiosyncratic use of that word—provide as much information and guidance as possible, but it is always the job of the individual game runner to present that material to his or her players and to respond and adapt to player input in a way that shapes the narrative and makes the game come alive. We planned to write a detailed guide to running *Tecumseh Lies Here* and produce and package the materials needed to do so: historical background, the fictional plot, web content, paper handouts, reproductions of primary sources, and other clues. We would then make these kits available to any history educator interested in running their own iteration of *Tecumseh Lies Here*. Running such a game would be demanding, as we well knew, but we hoped our kit would do as much of the work as possible, while offering

advice on customizing the game for different audiences and responding to player choices and actions.

We began working on the kit in earnest in late 2012, and by the end of the year, we had written a substantial first draft of a guide to running the game, with a rough plot and many new puzzles and clues. This draft was primarily written by freelance writer and game designer Bill Templeton. Bill's face is known to every player of *Tecumseh Lies Here*—he has played Captain Smith in each iteration of the game—but his many contributions behind the scenes have not been adequately acknowledged to date. Bill's guide to running *Tecumseh Lies Here* is a remarkable piece of work, and we regret that we have never been able to complete and release it. But writing this draft made manifest many challenges of the “D&D module” approach.

An alternate-reality game is a complex undertaking with many contingencies and moving parts—considerably more complex than most tabletop role-playing scenarios. Trying to write a document that anticipated all possibilities and situations (or even most) was perhaps a fool's errand. It seemed necessary to include a fairly detailed history of Tecumseh and 1812, the fictional framing narrative for *Tecumseh Lies Here*, thorough advice on running a complex open-ended game such as this, and at least some explanation of our pedagogical choices. The draft version of our guide soon ballooned to 30,000 words, yet it still felt incomplete. The place-based nature of *Tecumseh Lies Here* created another significant challenge. One of our favorite features of the original game was that it took players to real historical sites and made use of their physical surroundings. How could this aspect be reproduced in a game that would be run by different people in many different cities and towns?

Finally, in trying to write a generic version of the game, one that could be played by different kinds of audiences in many different settings, we were forced to acknowledge just how much the first *Tecumseh Lies Here* had relied on assumptions we could make about its original players. All of the players in the 2011 version of the game were university students or recent graduates. Most were history majors. Many, though not all, identified as gamers. In other words, they were all well-educated young adults, computer literate and social-media savvy, with ready access to the internet, smartphones, and automobiles, and with no particular curfew or bedtime. We could expect our players to be handy with the internet and familiar with basic research tools. We could scatter clues across southwestern Ontario and trust them to get to those clues and to find their way home. We didn't need parental permission to involve them in the game, and we

remained happily, even wilfully, innocent of liability waivers and the like. Reproducing *Tecumseh Lies Here* for, say, an audience of middle-school students, or teenagers, or seniors, would involve very different parameters.

Around this time, we compared notes with a team of faculty and students from the University of Maryland that had created *Arcane Gallery of Gadgetry*, another alternate-reality game for history and science education.²¹ Like us, they had run a version of their game in 2011, and in 2012 were trying to scale it up for a larger public launch. We realized we shared many of the same challenges, although in at least one way, we were revising our games in opposite directions. The first iteration of *Arcane Gallery of Gadgetry* had been played in school by groups of eighth-graders, and the Maryland team was redesigning their ARG for adult and university-age players. We ran our first game for university students but wanted to design something that could be used in middle schools and high schools, too. Talking to the *Arcane Gallery* team made clear to us some of the real constraints involved in designing extracurricular activities for children. We had known that teachers could not send twelve-year-olds to a graveyard at midnight, but we had not really realized that middle-school teachers could not ask their students to use the internet without supervision or expect them to get to a public library on their own steam.

For all of these reasons, the D&D module began to feel unworkable. Many of the things that make a pervasive ARG memorable and affecting are hard to reproduce in a generic, modular way: the way game events respond directly to player actions; puzzles tailored to specific sites, or even to specific players (one of the most delirious moments in the original *Tecumseh Lies Here* came when a video clip of the sinister Mothmen called out several of the players by name); the paranoid feeling that the game is everywhere. To some extent, the “problems” we were trying to solve by scaling up *Tecumseh Lies Here*—the fact that it was a transient experience and one that could only be enjoyed by a select few—were the very things that had made the original game powerful and fun. In our previous chapter, we asked if the intensity of the classic alternate-reality game experience was in fact predicated on its exclusivity and irreproducibility.²² By the winter of 2012, we had come to believe it was.

Threshold Concepts for History

Feeling we had reached an impasse, we returned to our first principles. From the beginning of this project, we had always insisted that the deep

lessons of any educational game do not come from its ostensible subject matter, but from the decisions players make and the actions they perform. The key to good educational game design, we believed, was to design activities where the actions performed in play were the very skills and lessons you wanted to teach. For all its complexity, the 2011 version of *Tecumseh Lies Here* had grown from a simple idea: we wanted a game that was actually played by doing historical research. We also thought more about our audience. We realized that the first iteration of our game had been designed for a very specific audience, although we had not always admitted that to ourselves. So we stopped trying to design the new version of our game for a generic audience of “anyone” or “the public.” Instead, we chose a specific audience and context for play. We set out to design a version of *Tecumseh Lies Here* that could be played in school by middle-school or high-school students. Much of the interest in *Tecumseh Lies Here* had come from history teachers at that level, who wanted to teach historical thinking skills in their classes, yet had difficulty finding materials that were engaging to students and supported a critical, inquiry-based pedagogy.

We consulted history teachers and the literature on history pedagogy at the middle-school level. Though we were determined not to dumb down our material, we had to rethink some of the assumptions and expectations that came from our experience teaching history at the university level. The literature told us that many teachers are interested in an inquiry-based history pedagogy, but do not always have the tools at hand to construct those exercises. Our conversations with middle- and high-school teachers bore this out. An inquiry-based pedagogy is one that presents history as a discipline driven by questions, one that exposes students to ways of thinking critically about history, one that requires them to confront the nature of historical evidence, and teaches them to develop and defend evidence-based interpretations of the past. All that said, the specific praxis of historical research—the navigation of finding aids and call numbers and archival databases that made up so much of our 2011 game—plays very little part in high-school and especially middle-school history. Instead of trying to force middle-school history classes to conform to our definition of doing serious history—for instance, the often-unexamined assumption among professional historians that close reading of primary documents is the *ne plus ultra* of historical work—we worked to boil historical thinking down into a few essential threshold concepts, and then build our game up from those.

“Threshold concepts” are a powerful idea from the study of teaching and education. The term refers to “core concepts that, once understood,

transform perception of a given subject.”²³ Threshold concepts are often difficult or counterintuitive. They create bottlenecks in learning that trouble a large number of students year after year. Once grasped and integrated, however, they are transformative, leading students into new ways of learning and knowing. Since the idea of threshold knowledge was articulated by Jan Meyer and Ray Land in the early 2000s, it has been explored and developed in a variety of fields. Every discipline, it seems, has its own threshold concepts, though these are not always obvious to expert practitioners. Another aspect of threshold learning is that it is often irreversible. Once a threshold concept has been learned, it can be difficult for students—and more importantly, for teachers—to retrace their own steps and remember what a subject looked like in the days when that threshold concept escaped them.

There have by now been several efforts to define a set of threshold concepts for learning history, including Sam Wineburg’s “historical thinking,” Charles Anderson and Kate Day’s “ways of thinking and practicing,” and Peter Seixas and Tom Morton’s “big six historical thinking concepts.”²⁴ The exact typologies vary, but there is general agreement on most of the big ideas. We made an effort to identify just three threshold concepts for understanding history and to express them as clearly as possible for a middle-school audience. Our three concepts were as follows.

First, we said, *historians are detectives*. By this we meant that historical knowledge is always pieced together from incomplete evidence. As Bruce Van Sledright argues, “lay people seldom appreciate the idea that historical narratives are constructed from evidence that has been questioned, pieced together, and interpreted.”²⁵ Like detectives, historians study past events by sifting through clues or sources in the present. Those sources can be misleading, contradictory, or downright baffling. We wanted students to consider for themselves what actually constitutes a reliable source—what clues they should accept or dismiss, and why. Embedded in this concept is a fairly sophisticated epistemological lesson that how we know, or why we think we know, is just as important as what we know.

Our second threshold concept was the idea that *history is a conversation*. Just as primary sources disagree, so do historians. We wanted to show that different people, and groups of people, can have very different interpretations of the same events. Our hope was for students to explore different versions of the same past and see for themselves how perspectives can differ, weighing whether some versions of history are more accurate or

trustworthy than others. Most importantly, we wanted students to be able to explain why.

Our final concept was the idea that *history itself has a history*. How and why people remember a given event changes over time. As a result, the stories we tell about the past say something about us and our own present times, too. This is the essential historiographical idea. It is rarely encountered in elementary-school history, yet for many professional practitioners of history, it is the key concept that makes doing history interesting and worthwhile.

Each of these key concepts offers a new way of interpreting history and seeing the past. None of these concepts were our own discovery or invention. They align closely with the “historical thinking concepts” or “ways of thinking and practicing” described by Seixas, Wineburg, Anderson and Day, and others. If there was something novel in our project, it was our determination to bring these concepts to elementary-school students and to do so in the form of an engaging mystery game. But how was this to be done?

From Alternate Reality to Augmented Reality

The medium that tantalizes us so has gone by a number of names: computer simulation, artificial reality, virtual environments, augmented reality, cyberspace, and so on. More terms are likely to be invented as the technology’s future unfolds. But the enigmatic term “virtual reality” has dominated the discourse. It has defined the technology’s future by giving it a goal. . . . Virtual reality is not a technology; it is a destination.

—Frank Biocca et al., *Communication in the Age of Virtual Reality*²⁶

After consulting with high-school and middle-school teachers, we began to zero in on a way to design a game that would be accessible to younger students while also being reproducible, scalable, and reusable after our work was done. We eventually developed the most recent iteration of *Tecumseh Lies Here*—the “un-textbook” we shared and played through with several Ontario middle-school history classes in September and October 2013—by considering the differences between traditional public school and university pedagogy.

Perhaps the biggest difference between our approach to teaching his-

tory at the university level and the traditional public-school classroom is the centrality of the textbook. “In education circles,” writes Bruce Lesh, “textbooks are often presented as the origin of all that is wrong with history education.”²⁷ Textbooks are accused of promoting stereotypes, of presenting monolithic master narratives, of being politically correct, of not being politically correct enough, of being too expensive, and of being dull. All of these criticisms are fair. Yet to us, the most significant problem with history textbooks is the way they can work against real historical thinking. Textbooks make invisible the construction of history. They conceal the sources, questions, and arguments that are the real stuff of historical thinking. But it does little good for academic historians to berate high-school teachers for the use of textbooks (written, by and large, by academic historians). There are many economic, institutional, and practical reasons why teachers cannot simply abandon their textbooks. There may even be pedagogical reasons, too.

While we were considering the centrality of the textbook in middle-school history classes, we were also exploring the emerging medium of augmented-reality print, specifically an off-the-shelf software platform called Layar. The affordances of augmented-reality print immediately struck us as an ideal way to bring the inquiry-oriented approach of the original *Tecumseh Lies Here* to a traditional history text, while at the same time encouraging students to question the learning materials they were accustomed to. The history textbook itself became a subject of our game. We asked ourselves: Could we design a game that deconstructed the history textbook, one that made its own construction visible? Better yet, could we design a textbook that deconstructed itself, that was itself a game?

Layar’s main focus is advertising, though the company has expressed interest in using their platform in education and other fields. Content creators pay to upload images of print pages to Layar’s website, and then augment the print with extra content: video, audio, animations, links, and the like. When a user with a smartphone or other mobile device scans a printed page or advertisement using the free Layar app, the app recognizes the page and augments its image on the screen of the user’s device. The results can be quite impressive, with still images on the page seemingly coming to life, moving, speaking, and interacting with the user. Though primarily used to lure readers of print advertising into engaging with advertiser websites, we found the affordances of augmented print ideal for opening up or deconstructing printed content, overlaying competing narratives onto our textbook, and even revealing the sources that go into crafting a particular

interpretation. Augmented print let us deconstruct or explode the smooth, authoritative façade of the traditional textbook page.²⁸

Augmented print was and remains a fairly new medium, and few of the students we worked with in 2013 had ever experienced this technology, giving it considerable novelty at the time—though, as expected, it is becoming more familiar. The technology gave us the ability to let students uncover material on their own in a fairly unstructured way. We decided to produce a kind of “un-textbook”—a printed paper workbook that could be used and reused in a public-school classroom without any special technology at all, but that could also be augmented with additional content, puzzles, and activities. We embedded a loose narrative into the textbook and tied to an augmented reality scavenger hunt that we would hold on the 200th anniversary of the Battle of the Thames. But the book was designed to be reusable by anyone with a mobile device and the Layar app well after that anniversary had passed.

The Un-Textbook

The notion that students must first be given facts and then at some distant time in the future will “think” about them is both a cover-up and a perversion of pedagogy. . . . One does not collect facts he does not need, hang on to them, and then stumble across the propitious moment to use them. One is first perplexed by a problem and *then* makes use of facts to achieve a solution.

—Charles Sellers, “Is History on the Way Out of the Schools and Do Historians Care?”²⁹

We decided to make our un-textbook look like the rough research notes of Captain Smith, the obsessed researcher and reenactor from the first iteration of *Tecumseh Lies Here*. The booklet has an intentionally messy, scrapbook aesthetic to replicate a passionate amateur’s notes. This helped make a virtue of the fact that the game was made with very little money, and by people without technical training in layout or graphic design. But the aesthetic also reflected our ideas about textbooks. The production values of a glossy, well-produced textbook can actually serve to conceal the work that goes into the book and distance readers from the messy business of historical detection and reconstruction. Our workbook made the construction of history visible on every page.³⁰ In the pages of the un-textbook, Captain Smith took on the persona of a helpful guide who questioned what the stu-

dents were being taught and challenged them to determine for themselves what really happened in the War of 1812 and what it meant for them today. Using the Layar augmented-reality platform, we crafted augmented-reality “buttons” on each page, which were revealed when the user looked through their devices’ camera. Certain features of paintings or highlighted text seemed to float on the page, inviting students to press them and see where they lead. Some buttons took students to our own websites, stocked with primary sources, and others linked to external sites that explored various aspects of the War of 1812. Our favorite elements were the audio and video recordings embedded on the page. Bill Templeton reprised his role as Captain Smith and talked the users through the key ideas of the book. Seeing still images come to life on paper was a powerful part of the book’s appeal.

Instead of being a linear retelling of the history of Tecumseh and the War of 1812, we designed the *Tecumseh Lies Here* un-textbook to explore our three core lessons or threshold concepts. The book begins with the lesson that history is built from pieces of evidence, just as a detective reconstructs a mystery using clues. It explains that events such as the Battle of the Thames have to be reconstructed from fragmentary evidence. Most of our information about the battle comes to us from the testimonies of participants, but these can be contradictory or misleading. We try to convey the idea that some sources are more reliable than others, but leave it up to the students to decide which sources are credible. Our first lesson leads students through a deconstruction of the famous painting, based on a sketch by historian Benson Lossing, of Tecumseh in a British general’s uniform and a jaunty hat. Though this has long been *the* depiction of Tecumseh, we wanted the students to explore elements of the portrait and compare physical descriptions of Tecumseh provided by eyewitnesses, to recognize the unreliable elements of the famous portrait. (See figure 10.1.) Another exercise in this section challenged students to read through the earliest descriptions of the Battle of the Thames from newspapers and letters, and rank them according to which ones they believed were more or less reliable, discuss why they would doubt the veracity of some sources, and construct their own narrative from these contradictory pieces of evidence.

After students had built their own idea of what happened at the Battle of the Thames, we turned our attention to historical debates, and to the idea that historians often disagree with one another over what really happened and what those events really mean. This is the second lesson of the

workbook: that history is contested. Just as students produced many different narratives based on their reading of the primary sources, we showed that historians can also come to strikingly diverse conclusions about the meaning and significance of past events. In this section, the un-textbook presents three competing versions of the War of 1812, from the American, Canadian, and First Peoples' perspectives, augmenting this section with "voices from the past" and links to nationalist and mythic interpretations of other key events, such as the Battle of New Orleans and the burning of Washington, DC. We wanted to impress on students that how people understand events is often contingent on their present situation. For instance, we challenged the recurring Canadian claim that "we burned down the White House" by showing that there were no Canadians involved and that there was really no concept of a Canadian nation at the time. We then asked students to come up with explanations as to why so many Canadians believe this myth. A section on women during the War of 1812 resisted the "here is an obligatory section on women" approach and asked students a historiographical question instead: *Why* are women usually underrepresented in traditional histories? How would it change our understanding of history to put women's point of view at center? We concluded the section with links to televised historical debates on the question of who really won the war, and question the merit of these debates considering the devastating consequences of the war for the First Peoples.

The final lesson of our un-textbook is one that elementary school students are rarely exposed to: that history itself has a history and that the way we remember events changes over time. We wanted students to understand the basic ideas behind the changing nature of social memory, and that the stories we tell about the past, and the way we tell them, say a lot about us in the present. Students using the un-textbook compared and discussed memorials and monuments of the War of 1812 and Tecumseh, and considered the purposes and messages of these commemorations. Why, the book asks, are so many things in the United States named after Tecumseh, if he was their enemy in the war? Why do Canadians remember Tecumseh as wearing a British uniform or describe him as fighting "for Canada"? The augmentations on the pages helped the readers trace the legacy of Tecumseh in Canada, the United States, and Europe to see how often his name has been appropriated, by everyone from the United States Navy to East German communists to lawn-mower manufacturers. We hoped that after all this work, students would reflect on the distortions caused by such nationalistic or frivolous appropriation of an important historical figure, and what that

Tecumseh's Portrait ???

Tecumseh was a Shawnee leader who organized a large confederacy of Native Americans to fight against the United States before and during the War of 1812. He died (most people think) at the Battle of the Thames in October 1813, but his body was never found. Canadians, Americans, and the First Nations or Native peoples of North America all honor Tecumseh's memory, but every group tells his story in a different way. That's one reason this book is called *Tecumseh Lies Here*. Nobody knows where Tecumseh's body really lies, but lies or myths about Tecumseh are common. Who was the real Tecumseh? And how can we know?

We don't even know what Tecumseh looked like! This is the most famous portrait of Tecumseh, but it was drawn in 1868, decades after Tecumseh died. The man who drew this portrait was an American historian named Benson Lossing. He said that he based his image of Tecumseh on an earlier sketch drawn in 1808 by a man named Pierre Le Dan. But Lossing admits he made some changes.

How would you describe Tecumseh in each of these portraits? Which parts of these portraits do you think are realistic? Which elements seem like the artist imagined them? Why do you think an artist might have chosen to portray Tecumseh in this way?

2

Figure 10.1. A page of the un-textbook explaining the problems associated with alleged portraits of Tecumseh. Students explore the mystery more deeply by holding their smartphones or tablets over the page to reveal links, video, and other content.

does to a wider understanding of Tecumseh and the War of 1812. The book concludes with a discussion of how different aboriginal artists were commemorating Tecumseh and the contributions of the First Nations to the War of 1812, and invites students or readers to try to compose a fitting commemoration of their own.

Finally, embedded in all these lessons but hidden in the background, was a mystery game element in which Captain Smith quietly challenges what the students are being taught in class (figure 10.2). Secret buttons on several pages lead to a video of Captain Smith whispering, “your teacher isn’t telling you everything.” “Someone is being left out of the history,” Smith goes on to warn. By finding hidden links, students could reveal missing pages of the textbook, accessed online. This was the rabbit hole of the game within our textbook: a set of hidden lessons behind and around the textbook that commented on and even subverted the lessons the students were ostensibly learning in class. We asked our partner teachers not to tell students about this material at first but to let them discover it on their own and make of it what they would. This hidden curriculum explored the historical legacy of General Henry Procter, who has long been blamed for the British defeat at the Battle of the Thames, and was intentionally left out of the un-textbook as a lesson on the power of historical writing to shape popular understandings of the past.³¹ Students could uncover clues hidden in the book to reveal this extra content, and revisit the history of the general, putting him “on trial” for his alleged cowardice and determining for themselves whether he deserves his ignoble historical legacy or ought to be rehabilitated. In this game feature, students are again presented with conflicting facts and competing interpretations, and get to challenge past assumptions and even revise the historical record for themselves.

The *Tecumseh Lies Here* un-textbook was delivered to four Ontario classrooms in September 2013. We met with all four teachers, and in two cases with the students themselves, showing them how to use the book and then letting them explore it, both in class and on their own. After the students had been working with the book for a few weeks, we invited all four classes to join us in the game’s finale, at a major bicentennial celebration held on the site of Battle of the Thames Historic Park outside of Thamesville, Ontario. There, students had to scour the re-enactor encampment for augmented-reality signs that held the final clues for their game, encrypted in an eighteenth-century code that required a replica of a Thomas Jefferson–designed cipher wheel to decipher.³²

Preparing the un-textbook was a fairly labor-intensive task. The final

Figure 10.2. Captain Smith, portrayed by Bill Templeton, comes to life on the page with the augmented-reality Layar app, and explains how to use his research notes. Note the mysterious call to action at the top of the page, which leads students down the rabbit hole to hidden content.

product could perhaps have been more polished, and we had hoped to include more original video and audio augmentation than was ultimately used, but we certainly feel we succeeded in creating an in-class augmented-reality experience that was engaging, educational, and cost-effective. The booklets certainly proved to be very popular. Our partner teachers were extremely enthusiastic about the project and reported that their students enjoyed the experience. They were excited at the prospect of using their mobile devices in class, and most had not experienced augmented-reality print. But they also enjoyed and took to this approach to learning history. One of our partner teachers informed us that her students were “eating this up, engaged and thinking very critically.” Overall, we and our partner teachers were pleased with the project and felt that the learning objectives had been met. We distributed many more booklets at the Battle of the Thames bicentennial and encouraged people to play along. In future projects we plan to incorporate some of the evaluation techniques described in the chapter by Kee, Poitras, and Compeau in this book.

Back to Reality

Art must take reality by surprise.

—Francoise Sagan³³

Our experiments in alternate and augmented reality proved to be deeply rewarding, but they were never without their share of unanticipated challenges, mistakes, and difficulties. In both versions of the game, puzzles that we thought would be fairly straightforward tripped up some players for days, leading to late-night debates over whether to drop hints and clues to the players, or let them fail. Yet later in the game, when the players found their stride, riddles that we believed would take days were solved in a matter of hours, leaving us scrambling. The scavenger hunt we designed for our 2013 finale at the Battle of the Thames bicentennial was really too long for most of our middle-school participants to complete during what was already a jam-packed day of activities. Creating puzzles and challenges is a vital but tricky part of both alternate- and augmented-reality games, and can really only be mastered after considerable practice.

Another concern about the un-textbook was that its augmented features would only be accessible to students with access to smartphones and mobile devices, and therefore had the potential to exclude students without such devices. Our partner teachers, perhaps exasperated at the ubiquity of

distracting devices in their classrooms, assured us that this would not be a problem, estimating that at least three-quarters of their students had devices capable of handling the Layar app. When we actually ran the game, we found that the teachers had overestimated the number of their students with mobile devices and also their facility with those devices. We were glad we had designed the un-textbook so that many of its features could be used without an electronic device.

In an alternate- or augmented-reality game, much depends on the players. These are open-ended activities by design; therefore, if players are uninterested or unenthusiastic, the game will simply grind to a halt. We were lucky in that we had active and enthusiastic players in both games, and we ourselves were keen to keep things interesting. In its un-textbook form, *Tecumseh Lies Here* needed enthusiastic teachers to buy into the project, and once again we were fortunate. We did find that teacher buy-in, and the amount of time we were able to spend with teachers prior to the game's launch, was crucial in how well the in-class experiences worked. Some of our classes had extremely enthusiastic teachers with a strong understanding of the goals and methods of our project. Their classes, predictably, had the most success with, and got the most out of, our game. Our other teacher partners were interested but less grounded in the historical and pedagogical theory behind *Tecumseh Lies Here*. In the busy first weeks of September, we did not always have as much time to prepare with them as we would have liked, and their classes did have more difficulties making progress and remaining engaged in the game. All this simply goes to show what anyone should have suspected: that classroom teachers are the vital link in any educational initiative or innovation. In any future iterations of this or other games, we will know to devote more time and effort to engaging, planning, and training with our teacher partners.

Finally, in the alternate-reality version of *Tecumseh Lies Here* we had adequate funding and completed all the preparations and most of the artwork in the history department at the University of Western Ontario.³⁴ With a smaller team and budget, the second, augmented-reality iteration of the game relied on other parties for funding and for the augmented interactive print.³⁵ Some promised funding did not materialize, forcing us to cast about for other ways to pay our bills. And while Layar initially gave us a very generous rate on our subscription, relying on a commercial service did prove to be a liability to the project. After our un-textbook was created, Layar was acquired by a different software company, which altered their rates and terms of service. We do not have the funding to continue our

subscription at the new higher rate, and so the augmented portions of our un-textbook are now inactive. We are looking for other ways to revive the un-textbook, including constructing our own augmented print application, but the future of *Tecumseh Lies Here* remains unknown.

There is work to be done, and many more experiments lay ahead, but we believe all these experiments show promise. The sprawling, immersive narrative of the 2011 alternate-reality game was an effective and hugely engaging way to teach historical research. It combined elements of the mystery and detective genres with tried and tested historical methods to produce a unique learning experience where the act of doing real history and the mechanics of the game are inseparable. Yet while enjoyable and effective, we felt it would be difficult or impractical to scale up or repeat that game. Advances in off-the-shelf augmented-reality platforms such as Layar let us imagine and design a more repeatable experience that could engage many more students or players, and our experience with the 2013 *Tecumseh Lies Here* un-textbook shows the willingness and ability of both students and teachers to embrace new methods and complex historical concepts. Yet this iteration of *Tecumseh Lies Here* also had to contend with issues of time and budget and reliance on commercial services. Whether *Tecumseh Lies Here* was presented in alternate reality or augmented reality, it could never escape reality.

Ultimately, the historical concepts we have tried to teach are not reliant on the new techniques and technology explored in this chapter. Still, alternate-reality games and augmented-reality applications do have unique affordances that help convey ideas in engaging and experiential ways. Museums and historic sites around the world are adopting augmented reality to help enrich their collections and exhibits. The best such activities, we believe, will be those that include a game-like, playful element, and encourage free thinking and investigation. They will also be those that recognize how learning objectives, game design, and pedagogy are interlinked. Our own best successes came when we aligned these three elements and when we designed our games' challenges and activities—both in the alternate reality game and in the untextbook—to mirror our learning objectives or threshold concepts. These echoes and resonances were what made *Tecumseh Lies Here* more than a fun series of games; it was effective pedagogy. Yet getting all three elements right and getting them to work together—the game, the content, and the pedagogy—is not easy. This is, perhaps, what makes educational game design so hard. But the rewards, we feel, are worth it.

Notes

1. Peter Seixas, *Benchmarks of Historical Thinking: A Framework for Assessment in Canada* (Vancouver: Centre for the Study of Historical Consciousness, University of British Columbia, 2006), 2.

2. Sean Stewart, *Alternate Reality Games*, 2008, <http://www.seanstewart.org/interactive/args>; Jim Stewartson, quoted in “Events, Not ARGs: Interview with the Founders of 4th Wall,” *Variety*, May 4, 2009.

3. Compeau and MacDougall, “Tecumseh Lies Here: Goals and Challenges for a Pervasive History Game in Progress,” in *Pastplay: Teaching and Learning History with Technology*, ed. Kevin Kee (Ann Arbor: University of Michigan Press, 2014), 87–108.

4. <http://www.seanstewart.org/interactive/args/>

5. Image, Text, Sound, & Technology (ITST) Grant from the Social Sciences and Humanities Research Council of Canada (SSHRC). Robert MacDougall was principal investigator, with Kevin Kee and William Turkel listed as coinvestigators, Shawn Graham as a collaborator, and Timothy Compeau as project manager. *Tecumseh Lies Here* was also supported by the Ontario Augmented Reality Network.

6. For scholarly histories of Tecumseh see R. David Edmonds, *Tecumseh and the Quest for Indian Leadership* (Boston: Little Brown, 1984); John Sugden, *Tecumseh: A Life* (New York: Henry Holt & Co., 1998).

7. John Richardson, *Richardson's War of 1812* (Toronto: Historical Publishing Company, 1902), 154; Richard White, *The Middle Ground: Indians, Empires, and Republics in the Great Lakes Region, 1650–1815* (New York: Cambridge University Press, 1991, 2011), 518.

8. For a comprehensive history of the debate and mystery of Tecumseh's death and final resting place see Guy St. Denis, *Tecumseh's Bones* (Montreal: McGill-Queen's University Press, 2005).

9. <https://independentcompanyofforeigners.wordpress.com/>; http://intangibleharmonics.blogspot.ca/2011_07_01_archive.html

10. This idea was also explored by filmmaker Spencer McCall in his documentary *The Institute* (2013), which followed the San Francisco-based ARG *The Jeune Institute*.

11. Adriana Ayers, “Tecumseh Lies Here,” *Active History*, May 2, 2012, <http://activehistory.ca/2012/05/tecumseh-lies-here/>

12. <http://activehistory.ca/2012/05/the-development-of-the-route-1812/>

13. Compeau and MacDougall, “Tecumseh Lies Here,” 105–6.

14. Ian McKay and Jamie Swift, *Warrior Nation: Rebranding Canada in an Age of Anxiety* (Toronto: Between the Lines, 2012), 297.

15. <http://www.cbc.ca/news/politics/long-form-census-cancellation-taking-toll-on-statscan-data-1.1176466>; <http://www.cbc.ca/news/canada/ottawa/federal-libraries-archives-shutting-down-1.1139085>; <http://www.theglobeandmail.com/news/politics/war-of-1812-extravaganza-failed-to-excite-canadians-poll-shows/article8906910/>

16. Official website of the Canadian government's War of 1812 bicentennial: <http://1812.gc.ca/eng/1305654894724/1305655293741>

17. For two opinions on this recent shift see Ian McKay, *Warrior Nation*; and Noah Richler, *What We Talk About When We Talk About War* (Fredericton, NB: Goose Lane Editions, 2012).
18. Alan Taylor, *The Civil War of 1812: American Citizens, British Subjects, Irish Rebels, and Indian Allies* (New York: Vintage Books, 2010), Treaty of Ghent and the British failure to support Native buffer state, 413–14.
19. <http://www.idlenomore.ca/>
20. Gary Gygax, *Advanced Dungeons and Dragons: Dungeon Master's Guide* (Lake Geneva, WI: TSR Games, 1979), 86.
21. <http://www.arcane-gallery-of-gadgets.org/>
22. MacDougall and Compeau, "Tecumseh Lies Here," 102–3.
23. J. H. F. Meyer and R. Land, "Threshold Concepts and Troublesome Knowledge—Linkages to Ways of Thinking and Practising," in *Improving Student Learning—Ten Years On*, ed. C. Rust (Oxford: Oxford Centre for Staff and Learning Development, 2003); *Overcoming Barriers To Student Understanding: Threshold Concepts and Troublesome Knowledge*, eds. Jan Meyer and Ray Land (London: Routledge, 2006).
24. Sam Wineburg, *Historical Thinking and Other Unnatural Acts: Charting the History of Teaching the Past* (Philadelphia: Temple University Press, 2001); Peter Seixas, *Benchmarks of Historical Thinking: A Framework for Assessment in Canada* (Vancouver: Centre for the Study of Historical Consciousness, University of British Columbia, 2006); Charles Anderson and Kate Day, "Subject Overview Report: History. Report from the Enhancing Teaching-Learning Environments in Undergraduate Courses Project." University of Edinburgh (2005), <http://www.etl.tla.ed.ac.uk/docs/HistorySR.pdf>
25. Bruce Van Sledright, "Narratives of Nation-State, Historical Knowledge, and School History Education," *Review of Research in Education* 32, no. 1 (2008): 109–46.
26. Frank Biocca, Taeyong Kim, and Mark R. Levy, "The Vision of Virtual Reality," in *Communication in the Age of Virtual Reality*, eds. Frank Biocca and Mark R. Levy (Hillsdale, NJ: Lawrence Erlbaum Associates Inc., 1995), 4.
27. Bruce Lesh, "Why Won't You Just Tell Us the Answer?" *Teaching Historical Thinking in Grades 7–12* (Portland, ME: Stenhouse Publishers, 2011), 22.
28. See Lendol Calder on "uncoverage" pedagogy for the history survey class: <http://www.journalofamericanhistory.org/textbooks/2006/calder/>
29. Charles G. Sellers, "Is History on the Way Out of the Schools and Do Historians Care?" *Social Education* 33 (May 1969), 511.
30. The layout and book design was led by Oliver Charbonneau, a PhD candidate in history at Western, using Swift Publisher 3.
31. For an overview of "the much-maligned career of Henry Procter" see Sandy Antal, *Wampum Denied: Procter's War of 1812* (Montreal: McGill-Queen's University Press, 1997).
32. <http://www.monticello.org/site/research-and-collections/wheel-cipher>
33. Blair Fuller and Robert B. Silvers, "Francoise Sagan, The Art of Fiction, No. 15," *The Paris Review* 14 (Autumn, 1956), <http://www.theparisreview.org/inter-views/4912/the-art-of-fiction-no-15-francoise-sagan>

34. Along with Robert MacDougall and Timothy Compeau, the 2011 *Tecumseh Lies Here* team consisted of researchers Kristen Way and Anna Zuschlag; writer and film editor Oliver Charbonneau; actors and writers of additional material: Bill Templeton, Megan Baxter, and Nicolas Virtue; technology, Devon Elliot; and additional artistic work, Westley Cote.

35. The 2013 *Tecumseh Lies Here* team consisted of Robert MacDougall and Timothy Compeau. Layout and writing: Oliver Charbonneau; acting: Bill Templeton; filming: Megan Baxter.

CHAPTER I I

History All Around Us

Toward Best Practices for Augmented Reality for History

Kevin Kee, Eric Poitras, and Timothy Compeau

Augmented-reality applications are being used around the world on city streets and in museums to help people see the past like never before. In this chapter, historians Kevin Kee and Timothy Compeau team up with psychologist Eric Poitras to take a closer look at location-based AR apps and consider ways that this technology can be used to teach historical thinking skills. Just as importantly, they also explore methods of determining the effectiveness of these techniques for improving education and public history.

Over the past decade many educators have watched in horror as smartphones have entered their lecture halls and classrooms, throwing the ecology of the learning space into disorder like a destructive invasive species. Some professors and teachers have banned the devices, while others have vainly called out students for clandestinely tweeting and texting below their desks. Smartphones, many are convinced, are disruptive and detract from the overall effectiveness of the learning environment.¹ Museums and heritage professionals, on the other hand, have been more receptive to the potential uses of smart devices to engage students and the visiting public. Many have explored how mobile phones equipped with audiovisual capabilities can facilitate place-based learning and enrich exhibits with expanded information, varied sensory experiences, and interactivity.

What can history educators learn from museum and heritage professionals about incorporating smart technologies into history learning? In this chapter we suggest that the most promising educational applications are those that combine augmented-reality (AR) technology with playful exercises to promote the inculcation of historical thinking skills. Furthermore, by expanding our concept of the learning space and by moving beyond our classrooms and campuses into the places and spaces where history took place, we can use AR to reveal the history all around us in ways that can be immersive, challenging, and new to our students.

We begin by placing our work on AR within the research on interactive media in general. Two decades ago, “interactive media” meant “virtual environments”; we transported ourselves into computer-created worlds. Today, “interactive media” have been extended, allowing us to take computing into the real world. Therefore, for those of us interested in using interactive media to support teaching and learning history, AR provides opportunities to better see the past in various ways. This chapter examines one such way: contextualization, a technique that allows users to imagine themselves within the original context of the artifact, place, or event.

In the second part of this chapter we explain how we attempted to use AR to better see the past through our design and development of *Niagara 1812* and *Queenston 1812*, two of the earliest location-based-history iPhone applications. These apps took users on tours of the villages of Niagara-on-the-Lake and Queenston, Ontario, sites of some of the fiercest fighting in the War of 1812. *Niagara 1812* and *Queenston 1812* not only guided users to points of historical interest and significance, but challenged users to solve a historical mystery. Although most history AR applications have yet to embrace the concept of gamification, we suggest that incorporating play is the most effective way to deepen users’ connection with, and understanding of, the past.

In the third part of our chapter we evaluate and test more recently developed historical AR applications to identify the best practices of and uses for such technology in the classroom. We conclude with discussions on how we might develop AR technology—ideally with more involvement of humanists—to improve public history education.

Engagement and Historical Thinking in Interactive Media Environments

For educators interested in communicating history in new and compelling ways, AR offers intriguing possibilities. Interested in the design, develop-

ment, and testing of interactive media environments for history, we initially focused on best practices for history simulations and serious games in virtual worlds. This research was inspired, in part, by the work of Janet Murray,² who suggested that computer-supported media possess specific attributes or “aesthetics.” Arguing that we should explore new ways to leverage those aesthetics, Murray believed that we were on the cusp of an exciting digital media revolution. In many ways she was correct, except that our understanding of digital media and “computing space” changed fundamentally in ways she did not anticipate.

Murray focused on the holodeck—the hologram-filled virtual-reality facility on starships in the popular television and movie series *Star Trek*—reflecting an interest in leaving behind physical space to explore computer-generated “virtual environments.” In the last few years, however, we have moved into a new paradigm, one in which physical space is augmented with electronic and visual information. Instead of putting people in an artificial world, we can now augment the physical world by embedding it with digital data, networking, communication abilities, and enhanced properties.³ The image of a user in virtual space, exploring an immersive simulated environment, has been replaced by a person checking her email at the airport, in a car, or in any other existing space. Yet while chapters in this book illustrate various ways that AR technologies can enhance our understanding of the past, our specific interest is in AR applications that tell a story within a physical space, drawing on well-established conventions in the discipline of history. In the following discussion, we examine one component that can make such storytelling more effective in AR: contextualization.

To be on the same ground where events took place and to both see and experience specific objects and landmarks in ways that they may have been experienced in the past can help our students and visitors to heritage sites better understand the behavior of historical actors. To use an example drawn from the Battle of Queenston Heights (a pivotal battle in the War of 1812 that has become a legendary moment for Canadian nationalists), a history teacher can explain why the American strategy failed, but to be on site and see the swift-flowing river and the steep and imposing hill brings the conditions to life. AR can then flesh out the historical detail, revealing how the terrain appeared in the past so users can better grasp why the invaders were so eager to control the heights, and why the British general Brock was so desperate to retake them. This is what we mean by contextualization with AR: it reveals what people might otherwise miss—especially in landscapes that have been altered since the historic events in question.

In doing so, contextualization challenges the familiar and promotes inquisitiveness toward seemingly mundane features of a landscape: What is beneath that mound? Is it natural or manmade? How does the lay of the land influence how we live in it? Is there meaning behind the locations of houses, factories, or institutions? While recognizing that we can never be 100 percent successful in achieving the exact context for all sorts of practical and epistemological reasons, playfully combining technology with history and the user's own imagination can help us better understand the role the physical environment played in history. Indeed, it is this feature of AR technology that so attracts public historians.

In many ways, historians and antiquarians have always been interested in AR, long before we had a name for it. Plaques on a building, bronze and concrete monuments on a battlefield, signs and historical sites, and other markers alerting visitors and passers-by that something significant happened at that specific place are all, in a sense, original forms of AR. It is as if there is an innate human urge to annotate our environments and reveal the hidden stories and meanings. As our colleague Robert MacDougall has noted, these are ways historians have tried to make the past augment the present.⁴ Yet while historians and history buffs may have a vested interest in looking beneath the surface or in understanding all that shaped an event in the past, what about those who may see history as little more than a school subject? Or what about those who may not see the relevance of history in a technology-driven world that increasingly promotes the new and the modern? Can we use technology to increase user engagement, foster critical thinking, support the appropriation of knowledge and practices valued by historians, and create deeper connections to an understanding of the past? Just as importantly, how would we be sure that we have achieved these goals?

To identify the best AR applications for teaching history, we focused on design practices that engaged students long enough to learn the content and methods historians and history educators are trying to teach. To measure engagement, we used Reinhard Pekrun's "control-value theory of achievement emotions."⁵ Connecting engagement and achievement, Pekrun's model suggests that the more people feel "in control," the more they value the task, increasing enjoyment and decreasing boredom and frustration. Gamifying historical education applications, we hope, would allow users to have more control of and engagement with information about the past, thereby encouraging students to "do history."

Doing history requires students to move beyond the "facts" of history

to understand the skills of historical practice—generating, corroborating, representing, and assessing interpretations of the past. Therefore, to measure students’ appropriation of the vocabulary, concepts, and methods required to do history, we drew upon Peter Seixas’s “Benchmarks of Historical Thinking,” which offer “structural historical concepts” that can “guide and shape the practice of history.” According to Seixas, students should be able to determine what constitutes historical significance, effectively use—including asking good questions of—primary source evidence, identify continuity and change, analyze cause and consequence, appreciate historical perspectives, and understand the moral dimension of history interpretations.⁶ By creating an app that gave users both the opportunities and the tools to learn and practice such benchmarks, we aimed to promote AR apps that were not only entertaining but educational. Keeping these goals in the forefront of our work, we began to explore how these best practices might be manifested in AR applications for history.

Niagara 1812, Queenston 1812, and Other Applications

In 2008 we began development of the *Queenston 1812* and *Niagara 1812* apps, which were intended as aids for visitors exploring the villages of Queenston and Niagara-on-the-Lake, the latter of which hosts more than two million visitors a year.⁷ Many more tourists were expected to visit the region for the Niagara region’s commemorations of the bicentennial of the War of 1812, providing us with an excellent topic and opportunity to share the work.

While learning by doing is by now a truism in education, it is just as common in app development. Created and developed through a partnership between government, university, and private-sector businesses, the 1812 apps were part of a larger strategy to provide opportunities for students and young entrepreneurs to grow the interactive media development sector in the Niagara Region.⁸ The development team consisted of a producer; a project manager; a game designer; a technical director; software architects and software developers; a writer, researcher, artist, and content coordinator; as well as web designers and multimedia developers—many of whom had little previous experience in game development. All of us were building iPhone applications for the first time.

Working with Apple’s newly released software development kit (SDK) was just one source of the challenges we faced, and if the early bird catches the worm, it also entices the best cat—and we were eaten several times. As

recent history attests, the release of the iPhone and its subsequent updates transformed the smartphone landscape around the world, but revolutions are rarely neat and tidy. The *1812* team members eagerly worked with Apple's platform, but were occasionally frustrated by the rapid changes that came with its evolution. Beginning development a few months after the first iPhone was sold in Canada, the team suffered through frequent revisions to the SDK and regular updates to the iPhone operating system (iOS). The developers would spend months hand-coding functionality in Objective C (the iOS programming language), only to discover that a new release of Apple's SDK had solved our problem, or they would hard-code functionality only to find it that it was inoperable with the release of a new operating system.

As the programmers contended with the fluid state of the SDK and iOS, the rest of the team was grappling with how best to facilitate location-based historical experiences. There are essentially two options for giving users a place-based learning experience that connected real-world environments to the past, and our app designers decided to provide both. The first option, which gives users the power to choose where they want to go and what they want to see in a self-directed tour, is made possible through "roam mode." Some visitors would be interested in restaurants, hotels, and wineries, while others might seek out historic houses, forts, and churches. Roam mode helped developers incorporate both tourist and historic information into both *Queenston 1812* and *Niagara 1812*.

The second method of place-based exploration is a guided tour in which the app makes the decisions about where users will go and what they will see, taking users along a predesignated route. The design team, however, recognized that neither a linear tour nor even the roam mode concept was particularly unique, and that in time, many similar applications would leverage the GPS capabilities of the iPhone to provide similar experiences (an expectation borne out by the subsequent release of a host of travel applications, many from recognized guidebook companies such as Fodor's).⁹ Therefore, the team decided that the *1812* apps would offer something more than a conventional, linear tour: a fun, game-like experience that, team members hoped, would both challenge and delight visitors in unexpected ways. In short, we aspired to gamify location-based history.

A well-known study conducted by Roy Rosenzweig and David Thelen in the late 1990s found that although Americans were deeply engaged with history, academic study of the past held little attraction for them; they wanted personal experiences of history. Rosenzweig and Thelen, therefore,

called for a new “participatory culture” that could allow people to explore, consider, and see the history for themselves,¹⁰ and we embraced this idea with our use of “quest mode.”

In contrast to roam mode, which we used to build *1812*’s first place-based learning experiences, we built its second experience in “quest mode,” allowing the apps to lead users around a space, as in a guided tour that offers basic information on historic places and people. But here is where the *1812* apps were unique. Rather than giving discrete, disconnected information, the apps told a story in which users must accomplish specific goals to acquire more details, further the plot, and complete the tour.¹¹ In essence, *Niagara* and *Queenston 1812* encouraged active learning engagement.¹²

In combining narrative and game elements, we hoped to encourage visitors to think more critically about the history they were encountering—a key component of the best practices of historical applications development. After all, it is one thing to show visitors an historic image overlaid on an image of the site today and explain why historians think it is significant. It is entirely different to have users move through an actual, real-world environment with facsimiles of historical sources, accompanied by some ambiguous, competing narratives, with the users having the tools and encouragement to question what they are being told and arrive at their own answers. The story and quest features of the apps aimed to offer the latter, modeling effective teaching and simultaneously addressing the perennial complaint of historians that heritage sites have a tendency to distort or even commodify the past.¹³ Visitors could now engage with history and explore questions and uncertainties not easily embossed on a bronze plaque or sculpture.

The quest-mode feature in both *Niagara 1812* and *Queenston 1812* gave users the option of exploring the sites with place-based games: *Return of the Fenian Shadow* and *The Bomber’s Plot*, respectively. Both programs took the visitor on a tour through the respective village in an attempt to solve a long-standing mystery. The tours were led by ISAAC, the “Investigative Semi-Autonomous Artificial Consciousness,” a play on the name of war hero Major General Sir Isaac Brock, who was killed leading his men in a desperate charge against the American invaders in 1812.

In the case of *The Bomber’s Plot*, users were enlisted to help solve the real-life mystery concerning who bombed the monument honoring Isaac Brock. For context, users were told that British supporters dubbed Brock the “savior of Upper Canada,” but by the 1830s, reformers and rebels viewed the monument as a symbol of oppression. In 1837 and 1838, Cana-

dians who longed for an American-style republic rose up in a short-lived rebellion, aided by private American militia who invaded Upper Canada in an attempt to liberate the British colony. The rebellions were decisively quashed and the invaders repelled, but the revolutionary spirit lingered, resulting in an 1840 bombing of the Brock monument. One of the prime suspects, a Canadian rebel named Benjamin Lett, was acquitted due to lack of evidence, and the true identity of the bomber has never been confirmed. The application users were thus given the quest to evaluate possible suspects and determine who was responsible.

Along the way, the app introduced users to primary-source evidence, evoking one of Peter Seixas's "benchmarks of historical thinking." At the Queenston Baptist Church, for example, ISAAC provided users with a fiery declaration by the American colonel James Morreau, which directly implicated Morreau as the bomber. After reading the declaration, the users of the app solved a puzzle that highlighted differences between two of Morreau's signatures, one on the declaration and one on an authenticated document. On noticing the differences between the documents, users came to realize that the declaration was, in fact, a fabrication, reinforcing the historian's need to evaluate the credibility of source evidence.

En route to solving the bomber mystery, users also solved puzzles pertaining to the real world around them. For instance, when presented with an image of an historic map, users traced their fingers across the iPhone screen to reveal hidden messages that not only advanced the game's plot, but mimicked the early-modern cryptographic technique of using a candle's heat to reveal secret messages written in lemon juice. In another case, visitors compared present-day Fort Niagara, which stood watch imposingly from the banks of the United States 200 years ago, to an image a local historian "discovered" during research to determine how the fort had changed over the centuries.

In this gamified experience, those using the app ceased to be passive tourists at historic sites and became active players. Rather than merely listen to historical facts and narratives, users had to move through the tour area, search the historical images, and study their real-world environments for clues; for those who needed assistance (or for those who did not enjoy solving puzzles), several hints, and eventually the answer, were provided.

Some of our mystery was based firmly in history; other aspects were imaginary, but the whole process was an example of "inquiry-oriented history pedagogy" facilitated through a game.¹⁴ And though playing a game may seem like a departure from serious history education, using technol-

ogy to encourage playful historical thinking is—as we have suggested elsewhere—an effective way to inculcate a complex, analytical, and questioning engagement with the past.¹⁵ It is in this blurring of playing with, and playing within, history that users come to actually *do* history.

AR Platforms

In 2008, when our team first began working on *Niagara 1812* to see how AR might encourage historical thinking and inspire deeper engagement with history, there were few examples from which to learn. Today, however, an increasing number of exciting examples of AR are being developed for public history. In this section we explore two applications that are especially noteworthy: the GeoStoryteller platform, and the Cultural Heritage Experiences through Socio-personal interactions and Storytelling (CHESS) platform, developed at the Acropolis Museum in Athens.

Offering narrative-based tours, *German Traces NYC* was the first application to use the GeoStoryteller platform. Designed by Anthony Cociolo and Debbie Rabina of New York's Pratt Institute of Information and Library Science, as well as Brigitte Doellgast of the Goethe Institute New York, GeoStoryteller is an open-source platform that allows museum professionals and others parties to create immersive place-based GPS guided tours and narratives, complete with images, text, and sound. The program itself is streamlined and straightforward, permitting designers familiar with the PHP programming language to easily create their own experiences, which are then run on Layar's AR browser. As of 2014 there were fourteen different tours available in cities in the United States and Canada.¹⁶

German Traces NYC guides visitors through the pertinent sites of New York using a multimedia narrative of videos, text, and historic images to tell the story of German immigrants to the city. Users can personalize their experience by choosing from among forty-eight different personal stories. For instance, users can explore the city while hearing about W. F. Mangels, the subject of a Gilded Age rags-to-riches tale, who arrived to New York in 1883 at the age of sixteen and founded a thriving business designing and building carousels and other amusement rides for Coney Island and other parks. Directed to the subtle but still visible traces of Mangel's past life in the streetscapes around them, users are then invited to participate in a trivia game and post answers to social media sites.

Similar to GeoStoryteller, the goal of the CHESS platform is to achieve a personalized experience with the past. Combining artifacts, narratives,

AR, and pervasive game elements, CHESS guides visitors through the exhibits with story elements geared to the individual user. Visitors to the Acropolis Museum in Athens, for example, might be guided through galleries using images of horses found on artifacts as an entry point. Whereas children may need to help a lost horse get back to its friends in Athens, adults might find images of horses highlighting specific artifacts to better understand salient points of Athenian society.

The magical effect of AR can also bring artifacts to life, restoring vibrant paint to the white sculptures, or animating the powers of mythic creatures such as the glowing eyes of Medusa, allowing artifacts to become living elements of the tour. The statue of Medusa, for example, may become an adversary that children must overcome. Relying on their tablets for protection, children would see their screens turn to stone and shatter, thus dramatically illustrating the ancient Greeks' belief in the powers of the Gorgon. For adult users, the same statue could be placed in her original context as a guardian on the Parthenon, allowing users to appreciate her intended setting and function.¹⁷

These versatile platforms highlight several of the emerging best practices for historical AR development. They both utilize a narrative-based approach that threads the history, places, and artifacts together in a coherent story to engage the user. They also rely on location-based learning experiences. GeoStoryteller is particularly effective in this area in that it places users on the actual locations where the history happened, thereby creating a deeply immersive experience. Taking place within a museum, CHESS is not yet able to exploit location-based learning in the same ways GeoStoryteller can, but its ability to use elements of gamification to breathe life into static objects and to create a tangible link between artifacts and the past can help users develop a deeper connection to historical artifacts.

Testing the Effectiveness of AR for History

The question remains, however—do these applications actually achieve the goal of deepening and enriching an understanding of history? Research on the effectiveness of AR in fostering historical thinking skills is at an early stage, and most empirical studies seem to evaluate the application's usability—that is, the user's overall satisfaction with the technology, which for our purposes represents a visitor's experience with a guided historical tour that uses a location-based AR application. To gather and record observations on users' experience, researchers rely on, among other methods,

users' self-reported evaluations of experiences with the application and tour, posttour interviews, researcher field observations, and the AR application's log file records. The resulting data allow application designers to revise the interface and system guides to improve user experience.

GeoStoryteller designers who were interested in evaluating the effectiveness of their efforts, for instance, recorded exit interviews for the first few dozen users. The resulting transcripts were then analyzed independently by two different evaluators using a rubric that placed visitor experiences on a sliding negative-positive scale. The transcripts revealed that many of the participants were already familiar with the tour route; therefore, a recurring source of engagement was learning more about the historical significance of a building or a feature of the landscape hidden in plain sight. The evaluators also learned that users were occasionally frustrated by some of GeoStoryteller's AR features; even so, these narrative features managed to resonate with users, underscoring the power of place-based learning.¹⁸

Like the GeoStoryteller team, CHESS designers also conducted empirical studies to gauge the effectiveness of their AR applications. A 2013 study used both self-reporting techniques—such as post-tour interviews, after-action surveys, and questionnaires—and observational methods in which the researchers shadowed the users as they toured the exhibits with the application; one observer took notes while the other recorded the tour on video. As one would expect, individual users varied greatly in terms of their curiosity and patience. Some users hurriedly skipped over large segments of the tour while others meticulously explored every part to avoid missing any information. Observing how users navigated through the application and tour allowed designers to come up with significant improvements, such as focusing on a central theme and allowing visitors to choose how much information they wanted on that theme and on related topics. They also determined that a “hurry” button and a timeline that indicated where the visitor was in relation to the whole tour significantly improved the experience for less-patient or time-constrained users. Further research also identified limitations of the application; ambient light, for example, could cause the AR technology to malfunction, and some users became confused when they missed place-based markers.¹⁹

While such usability testing is vital to the design process, it is insufficient for helping us answer whether an AR application meets intended instructional objectives: Is this application actually improving the learning experience? Does the application facilitate knowledge and understanding

of the history any more effectively than traditional methods? Does the application meet the benchmarks for historical thinking that we have set? We must have verifiable data to either prove our success or help us improve. So while testing the functionality of the applications is critical, accurately gauging the learning outcomes should be of equal concern. In other words, technological development, crafting engaging historical narratives, and delivering effective lessons that instill historical thinking skills must work hand in hand, and only a variety of methods will help us understand the full effect of these efforts.

The challenge is that historians and humanists are, in general, not well equipped to determine the effectiveness of the AR applications they have developed. In this case Kee and Compeau partnered with the educational psychologist Eric Poitras to fill in the gaps of knowledge and expertise. While the evaluation techniques most AR developers are using focus mainly on practical issues of usability, some of their methods incidentally catch evidence of learning-in-action, which can help us evaluate the instructional outcomes of location-based AR applications. One method in particular concerns recorded and transcribed conversations between tour guides and visitors. Dialogue analysis can reveal user thought processes and how users feel at specific locations throughout a tour, and psychologists can help us quantify, track, and interpret meaningful elements in user conversations and interviews. By identifying topics users found interesting, explications they found effective, and information they may have misunderstood because of problems either in the historical narrative or in the AR design itself, we can determine whether our learners are, in fact, contextualizing the history and grasping the ideas that we are trying to teach.

Self-reporting measures, such as brief questionnaires administered at different locations, can also provide essential insights into the teaching effectiveness of the AR applications, especially when cross-referenced with the user dialogue analysis described above. For instance, a user could be asked to respond to a question such as, “After visiting this tour location, I start looking forward to the next location.” By selecting an option on a scale ranging from 1 to 5, where 1 is “Strongly Disagree” and 5 is “Strongly Agree,” users can help us quantify their emotional states (drawing on Reihard Pekrun’s theory, summarized above²⁰) in terms of categories (positive vs. negative), arousal (intensity), and object (the features in the environment that produced the emotions). Or, to evaluate an aspect of the application, such as basic factual information mentioned during the tour, the user could be asked a simple true-or-false question (e.g., “*Benjamin Lett was an*

avid follower of the politics of William Lyon Mackenzie”) to gauge how essential information is being presented in the application; if users are consistently missing this basic information, we know something is wrong. Granted, such inquiries must be done carefully, since the questionnaires themselves can affect user experience. Nevertheless, such methods can help us identify when users are bored, confused, frustrated, or irritated by any part of the application, potentially affecting what they will learn.

The inherent challenge in evaluating applications, in terms of facilitating these types of experiences, lies in the fact that what visitors and students think or feel may not be immediately apparent; furthermore, user experience varies greatly depending on factors such as tour location, application interface features, visitor personalities, and tour-guide charisma. Therefore, we might consider more novel methodologies to test the effectiveness of AR applications. For instance, mobile eye-tracking tools, which examine areas where users fixate or shift their focus on, say, a primary-source document, would allow researchers to monitor what aspects users look at in both the application interface and tour location. Such data, highlighting interface features and contextual aspects that attract a user’s attention, could help researchers identify ways to promote historical thinking skills.

A slightly more invasive but potentially more illuminating mode of evaluation involves the use of galvanic skin response (GSR) sensors. Electrical conductance varies, depending on the amount of moisture on the skin, and since sweat is controlled by the sympathetic nervous system, GSR sensors that measure shifts in the electrical conductance as users move through the tour may help researchers identify the users’ levels of enjoyment, boredom, and frustration. If GSR readings that indicate enjoyment consistently spike at certain points in the tour, we may be able to isolate elements that users find more engaging.

Some of these methodologies are unobtrusive and could be readily employed in the context of the tour; others, however, may interfere with the user’s experience, or they may not be suitable for the outdoors. Eye-tracking equipment, for instance, relies on battery life, which is limited; interpreters may also find it difficult to distinguish a variety of other visual distractions that have nothing to do with the tour. Similarly, GSR sensors may be affected by the users’ physical activity, or even by the weather. For these reasons, virtual tours—that is, tours run with the app in a simulated environment using 360-degree panoramic touchscreens of the real locations—may complement empirical studies conducted in the actual tour setting.²¹

A pilot study to investigate the benefits of virtual tours was carried out using the McCord Museum's MTL Urban Museum AR application, which uses GPS-based markers to deliver historic images of buildings and streetscapes around the McGill University campus in Montreal. The study sought to evaluate users' ability to contextualize the past by noting historical changes or consistencies in buildings, streets, monuments, and modes of transportation. To do so, the study tracked users' eye-gaze movements and recorded user discussions to see how users interacted with the app, what aspects of the historical tour users noticed, and which elements users most commonly missed. Although most learners were able to independently recognize differences between the past and present depictions of the tour location, they needed more help identifying specific changes, such as noticing the impact of evolving transportation technology or changes in architecture. These preliminary findings suggest that AR apps that offer a variety of data sources are critical for helping users understand historical change.²²

In addition to evaluating how well users were grasping historical thinking skills through the application and tour, researchers are studying users' self-report and GSR data to better understand how users' engagement with the application itself might encourage or hinder their efforts to contextualize and learn about the past. Preliminary results in this area suggest that users enjoyed the tour, but also identified areas where their interest diminished and therefore their willingness to engage with and learn from the tour suffered. These findings have important implications for how we might revise the tour locations that users found to be of least interest.

Conclusion

In the coming years, AR technology is set to become more sophisticated and ubiquitous, and at the same time more accessible. Just as an explosion of user-generated content helped us reimagine the internet, the same will happen with AR. Additionally, wearable computers such as AR glasses, watches, and other technology will no doubt revolutionize AR and its ability to transform spaces into interactive playgrounds without the need for handheld screens.²³ We are only beginning to scrape the surface of this new medium, and future developers will perhaps look back on these early attempts as quaint. The potential for developing new methods of storytelling and sharing with AR is immense and we cannot foresee the most transformative advances.

Historians and public humanists who can keep up with these changes and contribute to their development will help set the standard for applications that do more than just use history as a theme or skin for an entertaining game; they will be able to build and promote tools that can express our understanding of the human condition, share our ideas, create new experiences, and cultivate empowered citizens.²⁴

To achieve these goals, however, we must continually identify and share our best practices for humanities AR design. That means we must design our applications in ways that foster deeper engagement with history and encourage learners to think critically about the history they are taught. The *Niagara 1812* and *Queenston 1812* projects, along with other applications of AR described in this book, demonstrate how AR can inculcate users in the practices valued by historians in challenging and playful ways. By continuing to collaborate with and draw on the work being done in other disciplines, we will become more efficient in our efforts to engage others in effective, transformative history education. We are not limited by the technology, only by our imaginations.

Notes

1. For just one of many examples that fill the media, see globalnews.ca/news/1547181/smartphones-in-the-classroom-technological-tool-or-total-distraction/, accessed Oct. 15, 2014.

2. Janet Murray, *Hamlet on the Holodeck: The Future of Narrative in Cyberspace* (Cambridge, MA: MIT Press, 1998).

3. Wendy Mackay, "Augmenting Reality: A New Paradigm for Interacting with Computers," *La Recherche* (March 1996).

4. Rob MacDougall, "AR Out of the Box," paper presented at OARN 2012, Toronto, Ontario. www.oarn.net/events/conference/videos/

5. Rheinhard Pekrun, "The Control-Value Theory of Achievement Emotions: Assumptions, Corollaries, and Implications for Educational Research and Practice," *Educational Psychology Review* 18 (2006): 315–41.

6. Peter Seixas, *Benchmarks of Historical Thinking: A Framework for Assessment in Canada* (Vancouver: Centre for the Study of Historical Consciousness, University of British Columbia, 2006), 2. Cf. Compeau and MacDougall in this book.

7. The iPhone application was accessed at www.ihistorytours.com, or directly from the iTunes store.

8. *Niagara 1812* was produced by a corporation created by Kee, with funding from the Ontario Media Development Corporation and the Ontario Trillium Foundation.

9. www.fodors.com/mobile-apps/

10. Roy Rosenzweig and David Thelen, *The Presence of the Past: Popular Uses of History in American Life* (New York: Columbia University Press, 1998), 190–200.

11. S. W. McQuiggan, J. P. Rowe, and J. C. Lester, "Story-based Learning: The Impact of Narrative on Learning Experiences and Outcomes," in *Intelligent Tutoring Systems 2008, Learning Notes in Computer Science 5091*, ed. B. E. A. Woolf (Berlin: Springer-Verlag, 2008), 530–39.
12. S. W. McQuiggan, J. P. Rowe, and J. C. Lester, "Story-Based Learning: The Impact of Narrative on Learning Experiences and Outcomes," *Intelligent Tutoring Systems 2008*, 530–39.
13. David Lowenthal, *The Heritage Crusade and the Spoils of History* (New York: Cambridge University Press, 1998), see especially chapter 4, "Heritage Assailed," 88–104.
14. Timothy Compeau and Robert MacDougall, "Tecumseh Lies Here: Goals and Challenges for a Pervasive History Game in Progress," in Kee, *Pastplay: Teaching and Learning History with Technology* (Ann Arbor: University of Michigan Press, 2014), 89.
15. Kevin Kee, ed., *Pastplay*; Jane McGonigal, *Reality Is Broken: Why Games Make Us Better and How They Can Change the World* (New York: Penguin Press, 2011).
16. www.geostoryteller.org/about.php
17. A. Katifori, M. Karvounis, et al. "CHESS: Personalized Storytelling Experiences in Museums," in *The Seventh International Conference on Interactive Digital Storytelling (ICIDS 2014)*, LNCS 8832, ed. A. Mitchell (Cham, Switzerland: Springer International Publishing, 2014), 232–35. Writers compose the narratives using the PAROS authoring tool. See Y. E. Ioannidis, M. Vayanou, T. Georgiou, K. Iatropoulou, et al., "Profiling Attitudes for Personalized Information Provision," *IEEE Data Engineering Bulletin* 34, no. 2 (2011): 35–40. CHESS was also successfully employed at the Cité de l'Espace in Toulouse, France. See M. Roussou, A. Katifori, et al., "A Life of Their Own: Museum Visitor Personas Penetrating the Design Lifecycle of a Mobile Experience," in *CHI 2013 Extended Abstracts on Human Factors in Computing Systems—CHI-EA '13* (New York: ACM Press, 2013), 547–52.
18. Anthony Cocciolo and Debbie Rabino, "Does Place Affect User Engagement and Understanding? Mobile Learner Perceptions on the Streets of New York," *Journal of Documentation* 69, no. 1 (2013): 98–120.
19. L. Pujol, A. Katifori, et al., "From Personalization to Adaptivity: Creating Immersive Visits through Interactive Digital Storytelling at the Acropolis Museum," in *Museums as Intelligent Environments Workshop (MasIE), Proceedings of the 9th International Conference on Intelligent Environments*, eds. J. A. Botia D. Charitos (Athens, Greece: IOS Press, 2013), 541–54. Also see the evaluation for Cité de l'Espace: S. J. Rennick-Egglestone, M. Roussou, et al., "Indoors and Outdoors: Designing Mobile Experiences for Cite de l'Espace," in *NODEM (Network of Design and Digital Heritage) 2013 Conference: Beyond Control—The Collaborative Museum and Its Challenges* (Stockholm: Interactive Institute Swedish ICT, 2013), 89–97.
20. Rheinhard Pekrun, "The Control-Value Theory of Achievement Emotions: Assumptions, Corollaries, and Implications for Educational Research and Practice," *Educational Psychology Review* 18 (2006): 315–41.
21. This has been tested by our colleagues at the Learning Environments Across Disciplines research partnership, <http://www.leadpartnership.ca/>
22. J. M. Harley, E. G. Poitras, A. Jarrell, L. Pipe, E. Gonzalez, M. Duffy, S. P.

- Lajoie, K. Li, M. Morton, H. Tissera, and K. Kee, "Augmented Reality with Mobile Learning Technologies: Comparing Emotions and Learning Outcomes from Outdoor and Lab-Based Studies," paper presented at Educational Media, Montreal, QC, Canada, 2015. J. M. Harley, E. G. Poitras, A. Jarrell, S. P. Lajoie, M. Duffy, D. Cataldo, and K. Kee, "Augmented Reality with Mobile Technologies: Can It Lead to Learning and Engagement with Historical Settings?" Poster presented at the annual meeting of the American Educational Research Association, Chicago, 2015.
23. www.theverge.com/2014/10/5/6912979/microsoft-roomalive-research-projector-system
24. Alan B. Craig, *Understanding Augmented Reality: Concepts and Applications* (Waltham, MA: Elsevier, 2013), 1.

Hearing the Past

Shawn Graham, Stuart Eve, Colleen Morgan,
and Alexis Pantos

In this final chapter, a group of archaeologists take us into the world of aural augmented reality, explaining how lost soundscapes can be rebuilt and the profound impact they can have on a visitor's connection with the past. This has been a book about seeing the past, but as the authors remind us, we can use more than sight to explore the past. Hearing past soundscapes can provide whole new ways of experiencing, understanding, and feeling history.

This volume is about seeing the past. But “to see” does not necessarily imply vision, for we frequently “see” things that do not exist. In this sense, to see something can also mean to understand it: “I see your point” or “I see what you’re saying.” How, then, should we “see” the past? We cannot see the past; we can only see the present. And even when we look at something “from” the past, it still lives in the here-and-now. Augmented reality (AR) can help us better mediate differences between the past and present, yet even though AR does not require vision, the majority of AR apps currently available privilege the visual, overlaying reconstructions or text on an image of the present through a keyhole—the viewport offered by our small screens. But here, too, the clumsiness of our interfaces, the clunky visual overlays, create a cognitive load, a “break in presence” that interrupts what we are seeing with awkward details, preventing us from seeing the past and understanding it in any meaningful way.¹ This is why we talk of the historical imagination, or the archeological eye.

We assume that the senses neatly cleave, allowing us to prioritize one

sense over another. With our contemporary focus on the visual, we tend to prioritize sight over other senses, but in this chapter, we suggest that “hearing” the past is a more effective and *affective* way of providing immersive AR.^{2,3} We argue from cognitive and perceptual grounds that audio—spoken word, soundscapes, acoustic horizons and spaces, and spatialized audio—should be a serious area of inquiry for historians exploring the possibilities of new media to create effective immersive AR. To do so, we explore some of the phenomenology of archaeological landscapes and the idea of an “embodied GIS” as a platform for delivering an acoustic AR.⁴ Drawing on Phil Turner’s work on “presence” in an artificial environment, we explore “breaks” in presence that occur in augmented, mixed, and virtual environments.⁵ The key idea is that presence is created via a series of relationships between humans and objects, forming *affordances*; when these relationships are broken, presence and immersion are lost. Considering that the sense of hearing depends on attention, we argue that audio AR is particularly effective in *maintaining* what Turner calls “affective” and “cognitive/perceptual” intentionality. In short, the past can be “heard” more easily than it can be “seen.”⁶ However, hearing (and the active cognition that hearing requires) is an area that has not been studied to the same degree or in the same depth as the visual.⁷ We first explore the ways in which hearing and listening affect us before turning to three case studies that offer possible routes forward for an augmented historical audio reality.

“Eh? Can You Speak Up?” The Sense of Hearing

Hearing—and understanding—can be considered a tactile, haptic experience. Sound waves actually touch us. They move the tiny hairs of the ear canal, and the tiny bones within, and the various structures of the middle and inner ear turn these into the electrochemical pulses that light up the various parts of our brain. Sound is a kind of tele-haptic:

[T]he initial stage of hearing operates as a mechanical process. Later the mechanical energy of sound converts to hydraulic energy as the fluids play a larger vibratory role. Thus at its source, touch operates with and causes sound, and it is only through touch-at-a-distance that we have sound at all. The famous story of Edison’s ears bleeding from his aural experiments makes visceral this tele-touch, which is not always a gentle stroke, no matter how pleasant the sounds, voice or music we might encounter.⁸

Bishop goes on to argue that while touch and vision are senses that can only know the surface, sound gives us access to that which is hidden. Sound waves permeate, transgress, and transcend surfaces; they cause surfaces to vibrate, to amplify, and to muffle. In so doing,

Sound provides the means to access invisible, unseeable, untouchable interiors. If we consider the import of vision to the general sensorium [what we think of as the five senses] and metaphorization of knowledge, then the general figurative language of ‘insight’ runs counter to surface vs. deep understanding of the world. Sound, it would seem, not vision or touch, would lead us to the more desired deep understanding of an object or text.⁹

To demonstrate, Bishop points to Karlheinz Stockhausen’s *Helicopter String Quartet* as a piece in which sound and touch blur (and “slur”) into a kind of synesthesia, which defies the “assumed neatness of the sensorium.”¹⁰ Chris Godsen, in an introductory piece to an issue of *World Archaeology* on the senses in the past, argues that our Western “sensorium” influences and conditions how we understand material culture.¹¹ He advocates unlearning and unpacking the privileged position of sight, what others have called “ocularcentrism.”¹²

The effect of structured sound (let us call it “music”) on movement is another interesting area where the haptic qualities of sound may be perceived. Interestingly, there are aspects of music that seem to translate into movement “primitives.” A recent study explored the relationship of musical structure (dynamic, harmony, timbre) to guided walking (mobile tours).¹³ The authors note that a focus on structure in music sits between the thematic (where the emotional content of the music is manipulated) and the sonic (which is associated with spatial cues). Thus, they wondered what aspects of structure would be perceived by their nonmusically-trained study subjects (Western university undergraduates at an Anglophone university) and how the subjects would translate these into music. The subjects listened to four distinct compositions designed to emphasize one aspect of musical structure as they moved around an open field. The subjects were observed and afterward interviewed as to why they stopped, moved to particular areas, or moved in certain ways at particular moments during the music.

The authors found that participants interpreted silence in the music as a signal to stop, *crescendi* (a rising intensity in the music) as a sign to move forward, and *diminuendo* (a lessening intensity of the music) as an indication to end movement altogether. Meanwhile, musical punctuation prompted

listeners to try to understand the significance of the particular spot they were standing on, while timbre “colored” different areas; in other words “harmonic resolution” signaled “arrival.”¹⁴ As our case studies demonstrate, such interplay of silence and crescendo can be a powerful affective tool for conveying the density or paucity of historical information in an area.

Intentional hearing—that is to say, listening—also affects us in that it requires *attention*. In the crowded foyer of a cinema, for example, it can be quite difficult to make out what the person opposite is saying. We have to *pay attention*; the act is tiring, especially if we try to read lips, attempting to match visual cues with auditory cues. Similarly, in the quiet of a classroom, with their backs turned, teachers can hear the surreptitious whisper that, while much quieter than in the cinema foyer, speaks volumes. Hearing, unlike sight, requires *active* attention that divides our ability to make semantic or emotional sense of what is being said, or even to remember quite *what* was said, when the original audio signal is poor.¹⁵ What is more, our brain is processing the spatial organization of the sound (near sounds, far sounds, sounds that move from left to right)—that is, *how* it is being said, not just *what* is being said.¹⁶

In brief, sound requires cognition to make sense; there is nothing “natural” about understanding the sounds that reach our ears, and this act of attentiveness can help elide other breaks in presence, making sound an integral component for understanding the past in the context of our world today.

Culture and Soundscape

“As a little red-headed Metis kid,” Zoe Todd writes, “it never occurred to me that the city could sound different to anyone else.”¹⁷ Todd recently wrote a moving piece in *Spacing* entitled “Creating Citizen Spaces through Indigenous Soundscapes,” in which she describes, among other things, the profound effect of a flash mob occupying the West Edmonton Mall’s replica of *Santa Maria*, Columbus’ flagship: “The sounds of Indigenous music, language and drumming soaring high up into the mall’s glass ceiling was a revelation: decolonization of our cities is not merely a physical endeavor, but also an aural one.”¹⁸

Soundscapes affect us profoundly, and as Todd demonstrates, they can be used to radically reprogram, repatriate, decolonize, and contest spaces. Work on the cognitive basis of memory has shown that, rather than being like a filing cabinet from which we retrieve a memory, the act of recollect-

tion actively rewrites the memory in the present, making our memories as much about our present selves as they are about the past. Thus, cognitive scientists working in the field of posttraumatic stress disorder are finding that they can reprogram the emotional content of traumatic memories by altering the contexts within which those memories are recalled, and sound plays a critical role in such work, literally rewiring our brains and our understanding of memory.¹⁹

The work of cognitive scientists, plus Tim Ingold's observations about the "meshworks" of interrelationships that create spaces and bind them in time,²⁰ prompts questions about the ways sound might help us access and work with the past: can soundscapes help us "visualize" the past, or at least bring to the surface different patterns in the meshwork? Can we reprogram collective memories of place with sound?

Such questions have been explored by a number of scholars, particularly archeologists in the field of archeoacoustics. Most work on archeoacoustics has explored the properties of enclosed spaces such as caves, theaters, and churches.²¹ In particular, Mlekuz has investigated the soundscape of church bells in an area of Slovenia. He takes Schafer's definition of the soundscape: whereas an acoustic space is the profile of the sound over a landscape, the soundscape is a sonic environment—with the emphasis on the way it is perceived and understood by the listener.²² This clear distinction between the mechanics and properties of the sound (the acoustic nature) with the effect it has on the listener (the soundscape) fits perfectly with Turner's idea of the "arc of intentionality." Where we may be able to recreate the sounds of the historical past, we may not be able to recreate how these sounds came together to create the soundscape of a person existing in that past. The soundscape is a combination of the acoustic properties of sound, space, and the individual. However, the acoustic nature of historical sounds *will* affect us as human beings and *will* evoke some kind of emotional response—even if it could be argued that this response is not "historically authentic."²³

The next question to ask, then, is if sounds, music, and voices from the past can affect us in certain ways, can we deliver those sounds using AR to enable an in situ experience?

Aural Augmented Reality (AAR)

Audio tours using a handheld device rented or borrowed from a museum that guides a visitor through an exhibition has been a staple of many museums and heritage sites since the 1950s.²⁴ Once a bulky device that had

to be curated and maintained by the museum or heritage site, audio tours are quickly taking advantage of the smartphone-enabled age and releasing their tours as downloadable apps or podcasts. This is democratizing the audio tour, allowing new and alternative tours of museums and cities to be released and followed, and potentially undermining the “truth” of the official tour. While we recognize that the humble audio tour is a form of aural AR, experienced in situ and influencing the way the user experiences a space, we argue that such tours merely serve as a narrative-led experience of a space (much as a tour guide in book form would) and do not often explore the haptic or more immersive properties AAR promises.

Some applications have taken the idea of the audio guide further, such as the SoundWalk project, which offers alternative tours of the Louvre with a *Da Vinci Code* theme, or walking tours of the Hassidic areas of Williamsburg narrated by famous actors and members of the community.²⁵ What makes the SoundWalk tours different is that they are GPS-powered and place-specific; for instance, you are told to open specific doors when they are in front of you, or to look left or right to see individual features. With high-quality narration, sound-recording, and music and sound effects, these tours also play with the notion of self, melding the listener with the narrator: “. . . Okay, for today you [the listener] are Joseph, that’s my Hebrew name; that’s my Jewish name, and that’s your name, for today we are one.”²⁶ Through its “high-resolution” aural experience, the SoundWalk tours’ acting, sound effects, music, and beguiling narrative all come together to give listeners a feeling of immersion; listeners are encouraged to get lost in the experience, following the voice in their head.

An application that also uses the immersive aspect of storytelling to good effect is the fitness app “Zombies, Run!”²⁷ which is designed to aid a fitness regime by making running training more interesting. Logging into the app, users take on the role of “Runner 5,” a soon-to-be-hero who is going to save the world from the zombie apocalypse. The app utilizes the user’s GPS location and compass to direct them on a run around their local neighborhood, but they are continually being pursued by virtual zombies. Go too slowly and the sounds of the zombies will catch up to the users, who hear their ragged breath as they chase them around the park. As part of the run, users can also collect virtual medical supplies or water bottles—available through the use of in-game voice—all of which help stave off the apocalypse. While collecting supplies gamifies what might be seen as mundane exercise, the app’s visceral sounds of a pursuer getting closer to the runner add an emotional level to the physical exertion of being out of

breath, tired, and aching, and creates an immersive experience. The runner is not just trying to better her time—the runner is escaping zombies and trying to save the world. The drama throughout is created and magnified mainly through sound.²⁸

The SoundWalk project tours and the “Zombies, Run!” app were not specifically created with an ear to exploring and experimenting with historical sounds or soundscapes. Yet the immersive narrative (audio tours) and the gamification of a journey through an alternate present (Zombies, Run!) do offer lessons.

Three Archeological/Historical Aural AR Case Studies

Historians and archeologists are currently experimenting with AAR technology not just as a means to tell a story, but to allow users to “feel” the sounds and be affected by what they are hearing. Three AAR apps presently in development can help us see ways we might move beyond the visual interface, concentrating instead on the power of sound to direct, affect, and allow alternate interpretations. Although the following case studies are examples of prototype applications and proofs-of-concept rather than fully fledged applications with many users, they nevertheless show how sound can enhance in situ historical experiences.

Bronze Age Roundhouses

As part of his research using the embodied GIS to explore a Bronze Age settlement on Bodmin Moor, Cornwall, United Kingdom, Stuart Eve used a form of aural AR to aid navigation and immersion in the landscape.²⁹ By using the Unity3D gaming engine (which can spatialize sound), Eve created a number of 3D audio sources that corresponded to the locations of the settlement’s houses. In the modern landscape the ancient houses are barely visible on the ground as circles of stones and rocks, making it hard to discern where each house is. Yet because the resulting app was geolocated, users could walk around the settlement in situ and hear the augmented sounds of the houses (indistinct voices, laughing, babies crying, etc.), which got louder or quieter with distance from each sound source.

Eve then introduced virtual models of the houses to act as audio occlusion layers, simulating the effect of the house walls and roofs in dampening the sounds coming from within—and only allowing unoccluded sound to emit from the doorways:

At first, the occlusion of the sounds by the mass of the houses was a little disconcerting, as [visually] the buildings themselves do not exist. However, the sounds came to act as auditory markers as to where the doorways of the houses are. This then became a new and unexpected way of exploring the site. Rather than just looking at the remains of the houses and attempting to discern the doorway locations from looking at the in situ stones, I was able to walk around the houses and hear when the sounds got louder—which indicated the location of the doorway.³⁰

By modelling sound sources and relating them to the archaeological evidence, Eve encourages questions about the site's use that visitors can explore in situ. For instance, if some of the houses were used for rituals (as is indicated by the archaeological evidence), what sort of sounds might these rituals make, and how would these sounds permeate the settlement? More prosaically, if animals were kept in a certain area within the settlement, how would their sound affect the inhabitants? How far could people communicate across the settlement area using calls or shouts?

Historical Friction

Inspired by the work of Ed Summers (of the Maryland Institute for Technology in the Humanities), Historical Friction is a sound-immersion app in the style of *Zombies, Run!* Summers programmed a web app called *Ici*, French for “here,” which uses the native browser's abilities to “know” where it is in space and to search for and return all of the Wikipedia articles that were geotagged within a radius of that location.³¹ In its original form, it returned a list with a brief synopsis of each article, but in its current iteration, *Ici* returns articles as points on a map, along with the article's status (e.g., stub, “needs citations,” “needs image,” and so on). Summers's intent was for the app to work as a call to action, encouraging users to expand the coverage of the area in Wikipedia.

Visually, it can be impressive to see dots on the map as an indication of the “thickness” of the digital annotations of our physical world—a thickness that suggests how difficult it can be to physically move through places dense with historical information. But we wanted to exploit this idea through the haptic nature of sound. Historical Friction makes that possible.³²

Historical Friction deliberately plays with the idea of creating a break in presence as a way to focus attention on areas that are thick and thin

with digital annotations about the history of a place. To do so, Historical Friction initially took the output from Ici and fed it through a musical generator called Musical Algorithms, generating “music” that would be an acoustic soundscape of quiet or loud, pleasant or harsh tones as one moved through space, creating a kind of cost surface, a slope. As we iterated, we switched to a text-to-speech algorithm so that, as Ici loads the pages, the text-to-speech algorithm whispers the abstracts of the Wikipedia articles, all at once, in slightly different speeds and tones. Our goal was to make it painful, to increase the noise and discords, so that users would be forced to stop still in their tracks, take off their headphones, and look at the area with new eyes: Would the sound push the user from noisy areas to quiet areas? Would users discover places they had not known about? Would the quiet places begin to fill up as people discovered them?

Voices Recognition

During the inaugural York University Heritage Jam, an annual cultural heritage “hack-fest,” a group of archeologists, artists, and coders took the Historical Friction application as inspiration and created an AAR app called Voices Recognition. According to its creators, “Voices Recognition is an app designed to augment one’s interaction with York Cemetery, its spaces and visible features, by giving a voice to the invisible features that represent the primary reason for the cemetery’s existence: accommodation of the bodies buried underground.”³³

To achieve this goal, the app geolocates each of the graves in the cemetery and attaches it to a database of online census data, burial records, and available biographies of the persons buried within the cemetery. Using GPS and compass to geolocate the user within the cemetery, the app then plays the contents of this database for every grave within 10 meters of the user. In the example application, the data is voiced by actors; however, the full application will probably use computer-generated voices, due to the number of graves in the cemetery and the sheer amount of corresponding data.³⁴ Here, the effect of linking sound and place would inevitably make the abstract data—the deceased—tactile and visceral for users. Moving among grandiose individual monuments, users would hear the whispers of single stories; in other places, they would experience a deafening cacophony of voices—especially in areas of mass, unmarked graves where voices literally shout out and clamor to be heard. Collectively, these sounds would completely invert the conventional cemetery experience.³⁵ If this app were to be made live, the designers would need to think carefully about

which material would be suitable for the intended sphere and about how to best present and distribute that material for greatest effect, yet the concept highlights the ways audio can relate to a cultural location at a much closer, personal level than visual overlay and presentation alone can produce.

Building an Aural, Haptic, AR to Hear the Past

In a guest lecture to a digital history class at Carleton University in the fall 2014 semester, Colleen Morgan recounted her experience with the Voices Recognition app when it was being tested: “Voices, in the cemetery, was certainly the most powerful AR I’ve experienced.”

Building a convincing visual AR experience that does not cause any breaks in presence is the holy grail of AR studies. It is also something that is virtually impossible to achieve when we consider everything that can generate a break in presence: the mediation of the experience through a device (head-mounted display, tablet computer, smartphone, etc.), the quality of the rendering of the virtual objects, the level of latency in software that delivers the experience to the eyes. The list is both endless and scaleless; once you “solve” one break in presence, another occurs. The goal then can never be to completely eliminate breaks in presence, but instead, to recognize them and treat them with an historian’s caution.

Indeed, we can play with them deliberately and use their inevitability to underline the broader historical points we wish to make. For example, the use of artificial crescendo and diminuendo (such as with the Historical Friction and the Voices Recognition applications) arrests users, making them stop and consider why the sounds are getting louder or quieter. Similarly, by inserting prehistoric sounds into the modern landscape, Eve creates an anachronistic environment—sounds that should never be heard in the present, generating a clear break in presence—but one that jars our cognitive intentional state and prompts us to examine what that sound might be and why it might have been placed in that particular location.

In this way, these case studies show that AAR does not always have to be a “recreation” or a fully immersive experience. Instead, much as we would treat the written word as the result of a process of bias (what is represented) and production (the quality of experience), we should treat any AR experience as the result of these processes and one more: delivery (the way in which it is delivered). Yet hearing the past requires that we pay attention to more than the effects. We must consider the *affect* if we are to prompt the kind of historical thinking we should wish to see in the world.

Notes

1. P. Turner, "The Intentional Basis of Presence," in *Proceedings of the 10th Annual International Workshop on Presence*, PRESENCE 2007, the 10th Annual International Workshop on Presence. Barcelona, Spain.
2. S. Mills, *Auditory Archaeology: Understanding Sound and Hearing in the Past* (Walnut Creek, CA: Left Coast Press, 2014).
3. Y. Hamilakis, *Archaeology and the Senses: Human Experience, Memory, and Affect* (Cambridge, UK: Cambridge University Press, 2014).
4. Stuart Eve, "Dead Men's Eyes: Embodied GIS, Mixed Reality and Landscape Archaeology," *BAR Reports*, British Series 600, 2014.
5. Turner, "Presence."
6. J. Picker, *Victorian Soundscapes* (Oxford, UK: Oxford University Press, 2003). A. Corbin, *Village Bells: The Culture of the Senses in the Nineteenth-Century French Countryside* (New York: Columbia University Press, 1998). (See also, in a more descriptive vein, work by historians such as Picker on Victorian soundscapes and Corbin on church bells in nineteenth-century France.)
7. Carryl L. Baldwin, *Auditory Cognition and Human Performance: Research and Applications* (Boca Raton, FL: Taylor & Francis, 2012), 3.
8. Ryan Bishop, "The Force of Noise, or Touching Music: The Tele-Haptics of Stockhausen's *Helicopter String Quartet*," *SubStance* 40, no. 3 (2011): 25–40, 25–26.
9. *Ibid.*, 26.
10. *Ibid.*, 28.
11. Chris Gosden, "Making Sense: Archaeology and Aesthetics," *World Archaeology* 33, no. 2, 163–67.
12. *Ibid.*, 166; Julian Thomas, "On the Ocularcentrism of Archaeology," in *Archaeology and the Politics of Vision in a Post-Modern Context*, 1–12 (Cambridge, UK: Cambridge Scholar's Press, 2009).
13. Adrian, Hazzard, Steve Benford, and Gary Burnett, "Walk This Way: Musically Guided Walking Experiences," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 605–14.
14. *Ibid.*, 609–13.
15. Baldwin, *Auditory Cognition*, 6.
16. *Ibid.*
17. Z. Todd, "Creating Citizen Spaces through Indigenous Soundscapes," *Spacing* (Oct 1, 2014), <http://spacing.ca/national/2014/10/01/creating-citizen-spaces-indigenous-soundscapes/>, accessed January 21, 2015. D. Favro and C. Johanson, "Death in Motion: Funeral Processions in the Roman Forum," *Journal of the Society of Architectural Historians* 69, no. 1 (2010), 12–37, explores similar affective qualities of sound along funerary procession routes in ancient Rome.
18. Todd, "Citizen Spaces."
19. Cf. review article, Stephen Hall, "Repairing Bad Memories," *MIT Technology Review* (June 17, 2013), on the state of research, <http://www.technologyreview.com/featuredstory/515981/repairing-bad-memories/>, accessed January 21, 2015.
20. T. Ingold, *Being Alive: Essays on Movement, Knowledge and Description* (Abingdon, UK: Routledge, 2011).
21. B. Blesser and L.-R. Salter, *Spaces Speak, Are You Listening?* (London: MIT

Press, 2007); I. Reznikoff, "Sound Resonance in Prehistoric Times: A Study of Paleolithic Painted Caves and Rocks," *Journal of the Acoustical Society of America* 123, no. 5 (2008), 3603. M. Lisa, J. H. Rindel, and A. C. Gade, "How did the Ancient Roman Theatres Sound?" *ForumAcusticum* 2005, <http://www.conforg.fr/acoustics2008/cdrom/data/fa2005-budapest/paper/343-0.pdf>, accessed Jan. 28, 2015; D. Howard and L. Moretti, *Sound and Space in Renaissance Venice* (New Haven, CT: Yale University Press, 2010); P. Fausti, R. Poampoli, and N. Nodi, "Comparing the Acoustics of Mosques and Byzantine Churches," in *19th International Symposium CIPA (The International Committee for Architectural Photogrammetry)* (2003), <http://cipa.icomos.org/index.php?id=61>, accessed January 28, 2015. For an excellent review of the increasingly extensive literature, see S. Mills, *Auditory Archaeology*.

22. D. Mlekuz, "Listening to Landscapes: Modeling Past Soundscapes in GIS," *Internet Archaeology* 16 (2004), http://intarch.ac.uk/journal/issue16/mlekuz_index.html, para. 2.2.1, accessed January 21, 2015; R. Murray Schafer, *The Soundscape: Our Sonic Environment and the Tuning of the World* (Rochester, VT: Destiny, 1977).

23. See for instance Richard Cullen Rathe, *How Early American Sounded* (Ithaca, NY: Cornell University Press, 2005); for the full sensorium approach, see Eleanor Betts, ed., *Senses of the Empire: Multisensory Approaches to Roman Culture* (Abingdon, UK: Routledge, 2017).

24. See <http://www.acoustiguide.com/> and <http://www.duvaws.com/company/profile>

25. <http://soundwalk.com/>

26. Extract from the Williamsburg Men Hassidic tour <http://www.soundwalk.com/#/TOURS/williamsburgmen/>

27. <https://www.zombiesrungame.com/>

28. B. Perron, "Sign of a Threat: The Effects of Warning Systems in Survival Horror Games," in *COSIGN 2004 Proceedings* (2004), 132–41.

29. Eve, "Dead Men's Eyes."

30. *Ibid.*, 114.

31. <http://inkdroid.org/gici/>

32. Historical Friction may be found at <https://github.com/shawngraham/historicalfriction>

33. Stuart Eve, Kerrie Hoffinan, Colleen Morgan, Alexis Pantos, and Sam Kinchin-Smith, "Voices Recognition," *The Heritage Jam* (July 12, 2014). <http://www.heritagejam.org/jam-day-entries/2014/7/12/voices-recognition-stuart-eve-kerrie-hoffinan-colleen-morgan-alexis-pantos-and-sam-kinchin-smith>, accessed January 21, 2015.

34. A video of the app in action may be viewed at <http://www.youtube.com/watch?v=wAdbynt4gyw>

35. S. Eve et al., *Voices Recognition* (2014). Of course, this is not the first experiment in AR in a cemetery. The "Voices of Oakland" project began by using a backpack computer, GPS, and a "Wizard of Oz"-type controller (a person who observed the participants, and manually triggered the audio augmentations), in 2005. In this project, the cemetery is conceived of as a type of museum, with linear narratives telling the story of particular individuals (often in the first person)

interred within. Steven Dow, J. Lee, C. Oezbek, B. MacIntyre, J. D. Bolter, M. Gandy, “Exploring Spatial Narratives and Mixed Reality Experiences in Oakland Cemetery” in ACE ‘05 Proceedings of the 2005 ACM SIGCHI International Conference on Advances in Computer Entertainment Technology (2005): 51–60. DOI: 10.1145/1178477.1178484. See also <http://www.jdbolter.net/the-voices-of-oakland.html>

Contributors

June Ahn is an associate professor at University of California, Irvine. He conducts research on the design, implementation, and evaluation of socio-technical systems for learning—or how social, cultural, and institutional factors intersect with the affordances of new technologies—to create enhanced and equitable learning opportunities for all learners. Dr. Ahn received his PhD from the University of Southern California.

Kate Bagnall is an ARC DECRA Research Fellow in the School of Humanities and Social Inquiry at the University of Wollongong, where she is working on a comparative historical study of Chinese colonial citizenship in Australia, Canada, and New Zealand. She has published on various aspects of Chinese Australian history and is coeditor, with Sophie Couchman, of *Chinese Australians: Politics, Engagement and Resistance* (Brill 2015). Much of her research has explored the lives of women, children, and families in Australia's early Chinese communities and the transnational connections and *qiaoxiang* ties of Chinese Australians before 1940. In 2014–15 she received an Australian Academy of the Humanities Traveling Fellowship to undertake fieldwork research in southern China for her project Transnational Chinese and White Australia. She is @baibi on Twitter and you can find her blog at www.chineseaustralia.org

Elizabeth Bonsignore is an assistant research scientist and lecturer at the University of Maryland's iSchool. Her work focuses on the design of technology-mediated, playful social experiences that promote new-media literacies, arts-integrated science learning, and participatory cultures for

youth. She also serves as the director of KidsTeam, an intergenerational codesign group at Maryland's Human-Computer Interaction Lab that aims to create technologies by children, for children. She is particularly interested in the role that interactive, participatory narratives (e.g., alternate-reality games) play in helping underrepresented youth engage in lifelong learning.

Timothy Compeau is an assistant professor of history at Huron University College at Western University. He studies colonial North America and the Atlantic world with a focus on the cultural history of the loyalists in the American Revolution. He was co-creator of the Social Sciences and Humanities Research Council-funded alternate- and augmented-reality game *Tecumseh Lies Here*, which explored the memory of the War of 1812. In 2016, he completed a postdoctoral research appointment in digital humanities at Brock University exploring the applications of augmented reality for public history.

Devon Elliott is a digital historian whose work incorporates digital fabrication and physical computing tools to experiment with the technological past and the use of data mining and text analysis to explore cultural history. Devon applies 3D printers, 3D scanning, depth cameras, projection, Arduino, Phidgets, Processing, Max/MSP, SketchUp, SEASR/Meandre, and Mathematica for much of his work.

Caitlin Fisher directs the Augmented Reality Lab at York University, where she held the Canada Research Chair in Digital Culture for a decade. A 2013 Fulbright chair, Fisher is the recipient of many international awards for digital storytelling, including the Electronic Literature Organization Award for Fiction and the Vinaròs Prize for augmented-reality poetry. She serves as vice president of the international Electronic Literature Organization and sits on the international board of directors for HASTAC, the Humanities, Arts, Science, Alliance and Collaboratory, dedicated to "changing the way we teach and learn." Currently she is engaged in a four-year SSHRC-funded research project exploring the potential of long-form interactive narrative in virtual and augmented reality.

Kathryn Kaczmarek Frew is a PhD candidate in the English Department at the University of Maryland. Her research focuses on immersion

in transmedia narratives, and she served as the character writer for Violet Cannon in the alternate-reality game DUST.

Sean Gouglas is an associate professor and senior director of interdisciplinary studies at the University of Alberta. He is a principal network investigator and member of the Research Management Committee for the GRAND NCE. He has published research on gaming and gaming technologies, in particular as they relate to the Canadian gaming industry. He has also published research on computer-game art and on issues of sex and sexuality in video games. He is part of the team that developed the augmented-reality game-authoring environment fAR-Play (which stands for “for augmented-reality play”).

Shawn Graham is an associate professor at Carleton University and an archeologist and digital humanist. His work surveys the ways new media are used to construct cultural heritage knowledge, from the perspectives of practicing archeologists, historians, and the wider public. He explores the rhetorics implicit in different ways of encoding digital tools, and the effects these have on consuming and creating archeological and historical data. He keeps an open lab notebook of his research and experiments in digital history and archeology at his research blog, www.electricarchaeology.ca

Wayne Graham is the technical director at CLIR, the Council on Information and Information Resources. Wayne has built computer graphic systems, including augmented and virtual reality systems. His research interests include computer graphics, augmented reality, frontier studies, architectural history as well as quantitative and digital methodologies.

Derek L. Hansen is associate professor of information technology at Brigham Young University. His research focuses on understanding and designing social technologies, tools, and games for the public good and falls within the academic area of human-computer interaction. Dr. Hansen received his PhD from the University of Michigan’s iSchool.

Carlea Holl-Jensen has published fiction in *Fairy Tale Review*, *Queers Destroy Science Fiction!*, and the anthology *Curious Specimens*, among others. She holds an MA in folklore from Indiana University, Bloomington, and an MFA in fiction from the University of Maryland, College Park. She is the coeditor of *The Golden Key*, an online journal of speculative writing.

Edward Jones-Imhotep is a historian of the social and cultural life of machines. His research focuses on the intertwined histories of nature, social order, and technological failure in modern Europe and North America. He is the author of *The Unreliable Nation: Hostile Nature and Technological Failure in the Cold War* (MIT Press, 2017). His current research project, *Reliable Humans, Trustworthy Machines*, combines traditional qualitative analysis and close reading of archival sources with advanced digital methods to reconstruct a lost history of machines by exploring the cultural history of technological failure in the modern period.

Kevin Kee is dean of the faculty of arts at the University of Ottawa. He has held positions as associate vice president for research (social sciences and humanities) and Canada Research Chair in Digital Humanities (Brock University), assistant professor and director of undergraduate programs at McGill University, and director and project director at the National Film Board of Canada. Kevin is interested in the use of computing to analyze and express culture, and history in particular. His research program lies at the intersection of history, computing, education, and game studies. He edited *Pastplay: Teaching and Learning History with Technology* (University of Michigan Press, 2014).

Jes Koepfler is chief UX researcher at a user-centered design studio in Philadelphia. She specializes in usability evaluation, user research, and managing multidisciplinary teams. She has a passion for deriving design insights from robust qualitative approaches such as contextual inquiry, participatory design, and interviews. Her portfolio includes a wide range of projects and clients in the nonprofit, education, healthcare, and financial services sectors. Jes holds a PhD in information studies from the University of Maryland.

Kari Kraus is an associate professor in the College of Information Studies and the Department of English at the University of Maryland. Her research and teaching interests focus on new media and the digital humanities, digital preservation, game studies and transmedia storytelling, and speculative design. In 2015 she and collaborators partnered with NASA on DUST, a game promoting the deep-time sciences. Her latest transmedia work is *The Tessera*, created in partnership with the University of Maryland, Brigham Young University, Tinder Transmedia, and the Computer History

Museum. *The Tessera* was featured as a 2017 Official Selection at IndieCade, the International Festival of Independent Games, which the *LA Times* describes as “the video game industry’s Sundance.”

Robert MacDougall is an associate professor in the Department of History at the University of Western Ontario and an associate director of Western’s Centre for American Studies. He studies the history of information, communication, and technology. He is the author of *The People’s Network: The Political Economy of the Telephone in the Gilded Age* (University of Pennsylvania Press, 2013) and co-creator of *Tecumseh Lies Here*, an augmented-reality game for history education that explored the War of 1812. He is a longtime player and sometime designer of tabletop board and role-playing games, and is interested in all varieties of playful historical thinking.

Ian Milligan is an associate professor of digital and Canadian history at the University of Waterloo. He is the author of *Rebel Youth: 1960s Labour Unrest, Young Workers, and New Leftists in English Canada* (UBC Press, 2014) and coauthor of *Exploring Big Historical Data: The Historian’s Macroscope* (Imperial College Press, 2015). He is PI of the Web Archives for Historical Research Group, supported by the Andrew W. Mellon Foundation, the Social Sciences and Humanities Research Council, and the Ontario Ministry of Research and Innovation. In 2016, he was awarded the Canadian Society for Digital Humanities’ Early Outstanding Career Award.

Bethany Nowviskie is executive director of the Digital Library Federation at the Council on Library and Information Technology, and a research associate professor of digital humanities at the University of Virginia. She is past president of the Association for Computers in the Humanities, former director of the University of Virginia’s Library Scholars’ Lab, and a long-time digital humanities practitioner.

Anthony Pellicone is a postdoctoral scholar with the Complex Play Lab at the University of Wisconsin–Madison’s Center for Education Research. He received his PhD from the University of Maryland, College Park. His research interests lie in the ways that we live, learn, and interact in the social spaces that surround games and play. In addition to his work on educational augmented-reality games, he is also interested in the ways that game culture is produced in emerging online play practices (such as

live-streaming), and how these cultures mediate the experience of playful learning in digital games. He has also worked as a research scientist on the CSforALL Consortium, and on the Susan Crown Exchange's Digital Learning Challenge.

Eric Poitras is an assistant professor at the University of Utah and a co-investigator on the Using Augmented Reality Applications to Foster Learning and Engagement of History project within the Learning Environments Across Disciplines Partnerships grant-funded project (led by Suzanne Lajoie). His research interests involve inspection of learning sciences, artificial intelligence, and educational data mining. His focus is the design, development, and implementation of technology-rich learning environments on computer- and mobile-based platforms.

Geoffrey Rockwell is a professor of Philosophy and Humanities Computing at the University of Alberta. He has published on videogames, textual visualization and analysis, and computing in the humanities including a book from the MIT Press, *Hermeneutica: Computer-Assisted Interpretation in the Humanities*. He is a co-developer of Voyant Tools <voyant-tools.org>, a suite of text analysis tools, for which he and his co-developer were awarded the CSDH/SCHN 2017 Outstanding Contribution Award. Rockwell also leads the TAPoR <tapor.ca> project documenting text tools and the Methods Commons <methodi.ca> documenting text analysis methods. He is currently the Director of the Kule Institute for Advanced Study.

Andrew Roth is a new-media artist and technical associate for research and learning support at Brock University's Centre for Digital Humanities. In his years at York University's Future Cinema Lab he led and collaborated in new-media installations, augmented reality experiences, the development of published apps, and the creation of tools for digital media artists. He has been a mentor for three CCA new-media artists in residence, exhibited work at the Niagara Artists Centre, and was a co-investigator of the Brock/IBM Arduino Youth and Learning Initiative.

Jentery Sayers is associate professor of English at the University of Victoria. He's the editor of *Making Things and Drawing Boundaries* (University of Minnesota Press) and *The Routledge Companion to Media Studies and Digital Humanities*.

Tim Sheratt is a historian and hacker who researches the possibilities and politics of digital cultural collections. Tim has worked across the cultural heritage sector and has been developing online resources relating to libraries, archives, museums, and history since 1993. He's currently associate professor of digital heritage at the University of Canberra. Tim's tools and experiments include important things such as the Real Face of White Australia, useful things such as QueryPic, and strange things such as the Vintage Face Depot. You can find him at timsherratt.org or as @wragge on Twitter.

William J. Turkel is professor of history at the University of Western Ontario and the author of *The Archive of Place: Unearthing the Pasts of the Chilcotin Plateau* (UBC 2007) and *Spark from the Deep: How Shocking Experiments with Strongly Electric Fish Powered Scientific Discovery* (Johns Hopkins 2013). His research combines computational history, big history, the history of science and technology, STS, physical computing, desktop fabrication, and electronics.

Index

- 3D scanning/modeling, 6, 34, 42, 43, 45
- Adobe Photoshop, 88, 122, 140, 186
- algoracry, 32–33
- alternate reality games, 70, 78, 158, 159–161, 176–177, 178–179, 182, 183, 184–185, 189, 190, 202
- archeoacoustics, 228
- archives, 4, 25, 43, 63, 65, 66; Australia, National Archives of, 5, 16–17, 18, 25, 39; Canada, 186; digitization, 17–18; records of oppressive systems, 16, 23, 26; web, 116, 118, 125, 128, 130–131, 132–133
- astrolabe, 73
- audiences, 3, 24, 51, 69, 83, 107, 142, 146, 148, 153, 154, 174, 176, 188, 191; primary-school, 138, 185, 190; public, 138
- augmented reality, (AR), 2, 58, 139–141, 208, 224; aural AR, 8, 224–233 *passim*, classroom education, 8, 138, 151–153, 191, 193–195, 196–197, 198, 200, 208; emotions/affect, 224–225, 228; games, 6, 158–174, *passim*, 202, 178, 212; GPS-based, 2, 150, 161–162, 184, 209, 212; magic mirror, 144; markers, 149–150, 152, 162, 168; print, 194–195; public history and museums, 70, 74, 138, 141, 162–163, 165–166, 208–215, 215–216, 229, 230–233; reconstruction/restoration with, 69, 71, 216; software, 141–143, 153–154; testing teaching effectiveness, 216–220; tracking, 151
- big data, 36, 40, 97
- binarization, 88–89, 90
- British Museum, 54
- Brock, Sir Isaac, 213
- “broken world thinking,” 6, 69–70, 80
- Brookhaven National Laboratory, 96
- Browne, Simone, 39
- burst documents, 110
- circuit diagrams, 97–98, 99–101; debates over, 104–107; and rationality, 102–103; “symbol language,” 104
- Cold War, 97
- collaborative design, 137, 138, 147–148, 154–155, 159
- Completely Automated Public Turing test to tell Computers and Humans Apart, (CAPTCHA), 33
- computer vision, (CV), 1–4, 12, 83, 88, 96, 111, 113, 117; definition, 32; history of, 34–35; present uses, 35–37, 38; privacy, 37, 131; social justice, 16, 21 39

- conjectural criticism, 56
critical editing, 50–51
- digital turn, 9
distant reading, 2, 40, 56, 116, 117–118, 121, 132
DUST, 70, 74–80
- experimentation, 5, 9, 33–34, 51, 56, 58, 59–66, 144
exploded view diagrams, 74
- fabrication, 41, 45–46
Facebook, 35, 78, 112, 131, 162, 169;
 Oculus Rift, 58, 59
face detection/recognition, 12, 18–20, 25, 83, 92–93, 130–131
Flickr, 93, 133
- GeoCities, 117, 118–120, 123–131
passim, 133
Google, 18, 33, 78, 131, 133, 182;
 Books, 86, 133; Cardboard, 66;
 Docs, 148; Glass, 58; Streetview, 133
- head-mounted display (HMD), 58, 65
hearing (sense), 225–227
Higginbotham, William, 95–96
historical method, 4–5, 40, 203
historical thinking concepts/skills, 177–178, 192, 197, 208, 211, 213, 214; contextualization, 209–210; technology, 203, 208–209; threshold concepts, 191–193, 196
Houston, Natalie, M., 56–57
- image processing, 87–88, 89–91, 96, 110; -mining, 116, 121; montages, 121–127
imagination gap, 3
Internet Archive, 93, 109, 117, 120, 133
- Layar, 141, 194, 196, 202, 215
Love, Heather, 56
- machine-learning, 33, 88, 92, 96, 111
magicians, 83, 84–85; trade magazines, 85, 92
Manovich, Lev, 39, 40, 121, 134
Massachusetts Institute of Technology, (MIT), 35, 95
Mathematica, 7, 57, 89, 92, 110, 128, 130, 134
Microsoft Kinect, 35, 66
Moore's Law, 108, 112
Moretti, Franco, 2, 40, 118
Morse, Samuel, 73
Murray, Janet, 209
- non-finito* product, 77–78
- Old Bailey Online, 86
Open Source Computer Vision (OpenCV), 18–19, 22, 25, 35, 39, 40, 57, 88–89
optical character recognition, (OCR), 2, 33, 39, 83, 86–87, 90, 92, 108; limitations of, 87
optical collation machines, 65
- pareidolia, 19, 20, 72
Pekrun, Reinhard, 210, 218
photogrammetry, 36, 42, 43, 49n30
playfulness, in learning history, 210
processing (programming language), 89
prototyping, 34, 43–44, 58
Public Face Database, (Pubfig), 36
Python, 35, 37, 39
- Rosner, Daniela, 71
- Sample, Mark, 40–41
seeing computers. *See* computer vision
skeuomorphism, 51, 61
smart phones, 58, 158, 161, 201, 207, 228; apps, 212–213, 229–230, 230–233; iPhone software development kit, 211–212
soundscapes, 227–228

- speculative computer vision, 6, 33–34, 40, 47
- spidering, 96, 109, 110
- stage magic. *See* magicians
- Sterling, Bruce, 38
- subversive commemoration, 186, 188
- Swinburne, Algernon Charles, 51; *Poems and Ballads*, 52–56
- Tecumseh, 179–181, 187, 196, 199
- Tenskwatawa, 179–180, 181
- text-mining, 40, 87, 96, 110
- Trouvé, Gustav, 32; electro-mobile jewelry, 34, 41–47 *passim*
- Ultimate Display, The, 57–58
- Underground Railroad, 8, 137–138
- Underwood, Ted, 22, 56–57
- unmanned aerial vehicle (UAV), 36–37
- virtual reality, 51, 59, 65, 66, 77, 79, 193, 209
- visual turn, 9
- War of 1812, 196–197, 199, 208; bicentennial, 186–187, 211
- Wayback Machine, 120
- White Australia Policy, 11–15, 23, 26, 130
- Wise, Thomas J., 54–55, 64
- Wolfram. *See* Mathematica
- YouTube, 112, 133
- Zappar, 62, 63