5

Diplomatic history after the big bang: using computational methods to explore the infinite archive

David Allen and Matthew Connelly*

When the first edition of *Explaining the History of American Foreign Relations* was published in 1991, it would have been hard to explain to readers why historians would soon find computers indispensable for doing their research.[1] Apple's first laptop, the "PowerBook 100," went on sale that year with a 20MB hard drive, and it cost over $4000 in today's dollars. No digital cameras were available on the consumer market. The World Wide Web was only just emerging, mainly to facilitate communication among scientists. Library catalog terminals were difficult to use, and scholarly article databases still lay far in the future.

Now it is hard to imagine writing a book without word processing software and an internet connection. Whereas the contributors to that first edition a quarter century ago would have had to spend many hours at the library just to find book reviews and check citations, these tasks are done today with a few mouse clicks. Historians now take thousands of photographs in a single archival visit, collecting gigabytes of digitized documents. We are organizing them as text-searchable pdfs in annotated, cloud-based databases, a practice that would have filled previous generations of scholars with wonder.

Remarkable though they are, these are still just technical improvements on time-honored historical tradecraft. The keyboard and screen take the role of the typewriter, databases substitute for filing cabinets and card catalogs, and track changes replace sticky notes. But the advent of digital media and recent advances in information technology portend much more dramatic changes in the very nature of our field, as is already happening in law, journalism, and literary studies. We will soon face an avalanche of electronic records, far too many to cope with using traditional methods.

This chapter describes how scholars might use computational techniques to cope with what William McAllister has called the "Big Bang" in historical source materials. This explosion began with the sudden release of a quarter of a million cables by Wikileaks; a "data dump" now dwarfed by the 2.7 million

https://doi.org/10.1017/CBO9781107286207.006 Published online by Cambridge University Press

electronic records available from the State Department's Central Foreign Policy Files (CFPF), spanning the years 1973–8. Even this pales in comparison with what is being created by newer forms of communication, such as the forty million emails generated by Bill Clinton's White House, and the two *billion* emails produced *per year* by Hillary Clinton's State Department.[2]

The digital archive is potentially infinite, as William Turkel notes, but archivists have already had to conduct triage with the relatively small collections of electronic records from the 1970s.[3] They lack the resources to cope with the much greater challenges to come in preservation, declassification, and curation, raising the question of whether archival integrity and proper finding aids will become a thing of the past. Official commissions have repeatedly called for the development of new technology to accelerate the processing of large collections, without resort to crude sampling methods or mass deletion.[4] The federal government has been slow to respond.

The risks in this new era of "big data" are therefore great, but so too are the opportunities. Applying computational methods to vast corpora of electronic records is not only essential to preserve meaningful access to the archival record. It may also begin to provide precise answers to previously intractable historical questions. In turn, it might prompt entirely new kinds of inquiries, once we can start to visualize the flow of information through the diplomatic bureaucracy in novel ways. All this may require fundamental changes in the practice of history. We may have to learn programming languages in addition to foreign languages, make our "data" available to other researchers, and form multi-disciplinary teams to conduct experiments.

Our field is unusually rich in documents, most of them in the public domain and many already available in digital form. Historians of American foreign relations therefore have a unique opportunity to lead the way, developing methods and standards that other scholars will learn from when Facebook and Twitter become the archives of social and cultural history. Reviewing previous attempts to apply computational methods to history reveals many potential pitfalls, whether neglecting older but still important questions and sources, raising false expectations of objectivity, or drawing younger scholars into what may turn out to be methodological dead-ends. Even so, perhaps the greatest risk of all would be not to rethink how we approach archival research when the archive itself is about to explode.

COMPUTATIONAL METHODS 1.0 AND 2.0: FROM CLIOMETRICS TO THE DIGITAL HUMANITIES

Great claims about the potential of statistics and computers in historical research are not new. As long ago as Frederick Jackson Turner there were calls for a quantitative political history, focusing on roll-call voting and election returns, which came of age in the mid-twentieth century in the work

of William Aydelotte and others.[5] In the 1950s, social science historians used punch cards and sorters to study urban, demographic, and African-American history. By the 1960s, a growing interest in quantitative and statistical history had spread across the field.[6] Economic historians started to call sources "data" and analyze them with innovative modeling techniques, dubbing themselves cliometricians.[7] Across the Atlantic, *Annales* historians used perforated tape to create data series for commodity prices over the *longue durée*, in what they called "serial history."[8] Articles proliferated in the *Journal of American History* and the *American Historical Review*, announcing the need to train graduate students in statistics and computer programming, lauding what computers would achieve.[9] Books described armies of research assistants and thousands of hours of data input, and came with supplementary volumes full of tables and mathematical formulas. The most zealous members of the movement argued that quantitative methods would eventually take over history, turn it from an art into a science, and rid the profession of ideological cant.[10]

Historians committed to traditional methods reacted in predictable fashion. American Historical Association president Carl Bridenbaugh railed against those who worshipped "at the shrine of that bitch-goddess, QUANTIFICATION."[11] Arthur M. Schlesinger, Jr., told the American Sociological Association that, although he did not deny the value of quantitative tools, nonetheless "almost all important questions are important precisely because they are *not* susceptible to quantitative answers."[12] The debate erupted into public view in 1974, with the publication of *Time on the Cross*, a computational attempt to reinterpret slavery.[13] It was not the first time that the "new" economic history had considered slavery, but Robert Fogel and Stanley Engerman did not just promise a revolution in academic methodology. They insisted that new methods would change public attitudes towards slavery itself, reshaping the contemporary politics of race.[14] Their book garnered unprecedented coverage in *Time* and *Newsweek*, but leading scholars took issue with unrepresentative data, misapplication of formulas and theory, and an agenda that belied the authors' pretensions to ideological neutrality.[15] In the aftermath, Fogel began writing of the limits rather than the promise of quantification, the movement's messianic fervor dissipated, and its adherents either changed tack or moved into other disciplines.[16]

Quantitative methods became dominant in economics and political science, but equations disappeared from mainstream historiography. Yet the cliometricians succeeded in engendering greater reflexiveness about method. As François Furet wrote as early as 1971, the laborious creation of data series made the historian, still male, "aware that he has constructed his own facts," something that Furet called "a revolution in the historiographical consciousness."[17] Looking back nearly three decades later, Joyce Appleby observed that quantitative history had at least made it impossible to deny the structural inequities in American history, even if the statistics had not spoken for themselves. From that insight, she argued, historians realized that deeper

analysis required revealing the power relations that had produced the numbers. That research was generally to be qualitative in nature.[18]

Social and cultural history took off, and a zeal for information technology was not reborn until the mid-1990s, with the coming of digital humanities. Then as now, the internet appeared to promise a solution to what was already seen as a crisis in the humanities. Twenty years later, the most important primers still declare that the new field of digital humanities not only revolutionizes the arts, or should, but "upends academic life as we know it."[19] For one of digital history's early trailblazers, Roy Rosenzweig, interactive CD-ROMs and hypertext online publications were guaranteed to bring history to a wider audience, and were thus steps on the "road to Xanadu."[20] For another advocate, Orville Vernon Burton, digital history was to be "a revolution in the history profession that will change the way history is done at every level of scholarship and teaching and throughout the libraries and databases historians use in their everyday work."[21]

Digital humanists have indeed created countless web-based projects, internet forums, and "unconferences." But one of their central preoccupations is whether, and how, all of their blogs, tools, and visualizations add up to a new kind of scholarship that can and should pass peer-review. Digital historians' proudest achievements have not been new discoveries or grand narratives – conspicuous by their absence – but easy-to-use interactive tools and free public platforms like Zotero. So does digital history have to meet the same standard as any other field of history, that is, to demonstrate that it has created new and important knowledge about the past? Or is its main role precisely to expose self-important academic pretensions and "shake things up," as Michael Frisch argues?

An even more fundamental debate concerns how we define and delimit the field. William Thomas suggests that it "is about the medium, not the method," which makes almost any history existing on the web "digital." Many practitioners celebrate collaboration, and point out that not every member of a team needs to have the same skillset. But for Daniel Cohen and William Turkel, only programming historians can do truly advanced research in digital history. All this creates a first mover problem, as Kirsten Sword asks: "Is it wise and fair to launch graduate students into their own, largely unsupported, digital projects when the 'best' work appears in large scale, collaborative ventures, and when scholarly articles and monographs remain our common professional currency?"[22]

One thing seems clear. There is not yet any agreement on how to define either "digital humanities" or "digital history." The debate will likely only be settled when digital historians produce "field-defining" work, the kind of work that commands the respect of the rest of the academy. As the leading scholar Cameron Blevins argues, though, "digital history seems to operate in a perpetual future tense," existing forever in the "sunrise of methodology."[23] By promising a revolution that has not yet come, digital history has aimed to

challenge every scholar without challenging anyone in particular. With few publications in leading journals, it is little surprise that discussions about digital history – what it is and what it can do – are happening mainly among self-professed digital historians themselves, rather than historians more generally. In literary studies, an *avant-garde* led by Franco Moretti and Matthew Jockers has already been bold and successful enough to spark a backlash in the core of the discipline.[24] Yet historians in the main evidence little more than mild curiosity about digital techniques, which they are as likely to read about in the arts section of their newspaper as in a scholarly journal.[25] For outsiders, it is all too easy to see digital history as something to teach but not practice, like world history, or even as a harmless form of public outreach.

This is unfortunate, because for historians of American foreign relations a digital turn is coming, whether we like it or not. For literary scholars, the development of huge corpora of digitized books by Google and HathiTrust has made computational analysis possible, but not imperative.[26] The advent of electronic records, on the other hand, has already brought profound changes in the core of our archives. While the growth of digitized and born-digital primary source collections is usually seen as an unmitigated good, the effects are complex, and in some ways deeply worrying.

THE ARCHIVAL EXPLOSION

Historians of the national security state have long been coping with the problem of "big data." What counts as "big," after all, is always relative to what has come before. As long ago as 1961, the advisory committee to the State Department's *Foreign Relations of the United States* (*FRUS*) series was warning of "the fantastic expansion of materials in the archives," which it called "a crisis of major proportions."[27] Twenty years later, Gerald K. Haines and J. Samuel Walker argued that one of the principal problems faced by historians of the future would be "an almost overwhelming task of sifting and winnowing an enormous amount of documentation."[28] They were right. The foreign policy bureaucracy expanded exponentially over the course of the Cold War, with the growth in size of the State Department and the intelligence community, as well as the addition of new players such as the National Security Council and USAID. The horizons of diplomacy widened too, as decolonization increased not just the number of states but also the burdens of global management, ranging from armed conflict to financial stability, from the eradication of disease to the control of world population growth. Consequently, as William McAllister puts it, our work over the past decades has embraced "more actors, more topics, more interaction, more documents, and more historiographical approaches," adding up "to a vastly larger universe of study."[29]

What is changing now is that more and more of the sources that make up this universe are either being digitized or were "born digital," in the sense that they

were originally created or archived as electronic rather than paper records. Alongside quantitative growth, this is bringing about a qualitative transformation even in the study of the more distant past. The individual researcher can now sift, search, and sort enormous collections with truly unprecedented storage and processing capabilities. The Federalist Papers have long been a test bed for statisticians who seek to perfect techniques in authorship attribution of anonymous documents. Now there is the prospect of applying this and other statistical techniques to much larger corpora, such as the fifteen thousand letters Benjamin Franklin wrote or received during his life, which Stanford's Mapping the Republic of Letters project has begun to quantify and visualize to try to understand Franklin's transnational connections.[30] Over 450 volumes of *FRUS*, dating to 1861, are already available digitally from the State Department Office of the Historian and the University of Wisconsin's Digital Collections, and new volumes are released in a variety of formats, from pdf to xml. Now, like never before, it should be a relatively simple matter to determine, among other things, changes in the relative frequency of references to this or that country or individual across the whole corpus.[31]

These kinds of analysis might be thought simple, except that there are not yet any web applications for the historian to identify anonymous authors, extract and map locations from documents, or conduct frequency analysis of digitized collections. The only tool that is typically offered to users of online archives is the search engine. And when it comes to the electronic reading rooms maintained by every federal department and agency to store documents released under the Freedom of Information Act (FOIA), these search engines can be quite primitive or even non-functional.[32] At the time of writing, the State Department's FOIA reading room alone offers access to over 100,000 records, which someone with modest coding skills can "scrape" (or copy) to create their own database.[33] More challenging to the would-be digital historian are the Remote Archives Capture (RAC) terminals, accessible only at presidential libraries. In an attempt better to manage declassification, the CIA and the National Archives and Records Administration (NARA) have digitized millions of pages of documents, and those digital copies are now starting to be released at libraries from Truman onwards.[34] Moreover, the CIA has declassified some 11 million pages of its own records, but only makes the full collection available at Archives II in College Park, through the CREST system.

These are public sources, and in theory anyone could print out and re-digitize the RAC and CREST materials. Yet much of our public record is already being scanned and sold for profit, as Roy Rosenzweig pointed out many years ago.[35] ProQuest hosts the Digital National Security Archive, home to over 700,000 pages of FOIAed documents.[36] It also owns the History Vault, an agglomeration of documents that includes many of the National Security Files of the Kennedy, Johnson, and Nixon White Houses, plus the archives of all the major American newspapers.[37] Gale/Cengage maintains the Declassified Documents Reference System (DDRS), composed of around 500,000 pages of

documents, mainly released through the Mandatory Declassification Review (MDR) requests that scholars have filed at Presidential libraries over the last forty years. DDRS might therefore tell us a great deal not just about history, but also historiographical fashion and declassification policy.[38] But until recently, the only way to explore it and all the aforementioned databases has been by intuiting what terms will yield interesting results from the omnipresent search engine.

Having access to more and more digitized documents has already made it cheaper and easier to conduct primary source research compared to traditional archival expeditions. Considering that some other countries have gone much further than the US in digitizing their national archives, such as Australia, Great Britain, Japan, and Switzerland, it is now possible to do multi-archival international research without ever leaving home. But the advent of "born digital" electronic record collections will again change the nature of our research, beginning with reducing our dependence on scanning and Optical Character Recognition (OCR). OCR, after all, consistently produces some garbled text depending on image quality and software quirks, which is one reason why search engines do not always produce even documents that contain the specified search terms. And scanning does not, by itself, yield metadata, or "data about data," such as the author, recipient, date, and subject of a document. Names and locations embedded in clean text can be extracted through what data scientists call "named-entity recognition." But the process is prone to error, since computers cannot tell whether "Paris Hilton" refers to the celebrity or a place to stay in the French capital.

True electronic records, on the other hand, usually come complete with native metadata, which allow for many more – and more rigorous – forms of analysis. Consider, for instance, the State Department's CFPF, the core collection for the study of American foreign relations in the twentieth century. In the late 1960s, the State Department began to experiment with automatically sorting airgrams. In the middle of 1973 it started to convert all telegrams to machine-readable microfilm. One impetus was a desire to generate internal data about diplomacy, such as through the "Traffic Analysis by Geography and Subject" (TAGS) system. These TAGS are now available to us as one of dozens of different fields of metadata that have been released together with full-text cables through NARA's "Access to Archival Databases" system (AAD).[39] "P-reel," or paper documents that were microfilmed for the record, are currently only obtainable at College Park, but the metadata for each is also available on AAD. Along with "subjects," "concepts," and other information, the metadata provides a history of how each document was declassified. For the hundreds of thousands of documents that remain classified or otherwise unreleased, electronic withdrawal cards are provided, albeit with more limited metadata than the declassified records. Each field adds another layer for potential analysis. And because humans filled each in at the time of record creation, declassification, or archival preservation, their inconsistencies are

revealing and interpretable. Simply knowing the sender, recipient, date, and subject makes it possible for the first time to conduct systematic analysis of the overall agenda and volume of American diplomacy. With each new installment, the CFPF will become ever more central in the study of international history since 1973, in digital history, and perhaps even in computer science, since work in the field of natural language processing (NLP) depends on a clean dataset that is rich with metadata.

In assessing what we have and what is yet to come, it is important to realize what we have lost. The CFPF at first glance appears overwhelming both in its size and its seeming completeness. But over one hundred thousand cables were corrupted in the transition to electronic records. For certain periods, most or all of the documentary record has simply been lost. This includes most of the cables from the first half of December 1975, for instance, and 92% of the telegrams from June 1976. Gone are records pertinent to the Indonesian invasion of East Timor, and the American response to the Soweto Uprising. Moreover, a large proportion of what survived intact remains unreleased. Whereas for the 1973 cables, 13% were withheld, for 1976, it was 24%. Withholding often occurs because certain collections are more likely to have national security or personally sensitive information, and NARA has not invested in technology to prioritize documents that require closer scrutiny. All of these records were simply printed out and reviewed page-by-page. And the years in which most of these cables were reviewed coincided with a dramatic decline in appropriations for declassification, from $232 million in 2001 to $48 million in 2004. Spending on declassification has not recovered, such that the inflation-adjusted budget in 2012 was just 15% of what it was in the late 1990s.[40] Consequently, archivists have felt compelled to delete millions of other documents (the exact number is impossible to determine). So large has the archival record become that NARA has lacked the staff to look at more than a small sample before deciding which records have enduring historical significance. Materials that did not make the cut include whole classes of cables concerning scientific research, cultural diplomacy, passports, and visas.[41]

Even in terms of records that survived, research has become more challenging. Paper documents were scanned or inputted into the State Archiving System's "P-reel" in the order in which they were submitted to records managers, not with other documents created concurrently or by the same person. Documents created in 1975 might not have been archived until 1980, and are therefore currently inaccessible without lengthy FOIA delays. Using a keyword search to identify (by metadata only, not full text) a "P-reel" document yields, after consultation with a ring-bound container list, a whole box of random documents that just happened to be scanned at the same time. This loss of archival integrity makes it impossible even to produce a thematic or institutional finding aid, a problem that will become all the more acute when archivists with deep knowledge of these collections retire and their institutional knowledge is lost with them. It is already a disaster for the historian, virtually

eliminating any chance of making a serendipitous discovery in neighboring files, or of gaining any greater understanding of the context in which these documents were produced.

The promise of rapid expansion of the virtual archive of American foreign relations can thus distract us from the pitfalls and dark corners awaiting the unwary researcher. So far, historians who wish to explore this archive have only been able to use a search engine. It is not unlike a flashlight, which we shine into the archive if only because we cannot think of what else to use. But computer science is beginning to produce a whole array of new techniques to explore virtual archives, the equivalent of infrared lenses and autonomous drones. It would be foolhardy for historians not at least to try to use them before we stumble much further into the darkness, and before millions more historical records are lost forever.

COMPUTATIONAL METHODS 3.0

So how should we grapple with these digital repositories, once we realize that they are disjunctive and disorganized, and that a large (but to some degree unknown) part of the original documentary record is unavailable because of corrupt files, deletions, and withholding?

Imagine beginning a book on Henry Kissinger's stint as secretary of state, and approaching the source base in the traditional manner taught to graduate students for decades: read everything, and then read around. It would take well over a lifetime for one person to read the cables for those few years, and nearly three years to read just the ones that were sent from Moscow.[42] One would need to add transcripts of Kissinger's telephone conversations, all 15,000 of them, plus the records of his meetings, the State Department papers that were not stored electronically, as well as Defense, Treasury, and intelligence records. Then combine all of that with the personal and non governmental archives that have been mined so profitably over the past decades by historians of American foreign relations. Finally, research in the archives of other countries and international organizations would be essential to correct for the intrinsic bias in a single government's records. Considering that the potential source base for more recent history is much larger, it is clear that new approaches are needed.

What if we made a virtue of necessity, and approached the archive in entirely new ways? Normally, when we go to a large physical archive, we enter with some idea of the key topics, consult the finding aids, learn the scope and content of the collections, and start ordering everything that seems relevant. We then look at the documents one by one, trying to glean insights. Sometimes they give us leads that we follow into other files, until we begin to think that we have some sense of how everything is and was connected. But we never have a very clear picture of the larger whole, since we never see more than a fraction of the full collection. This is the virtue and vice of "close reading."

Now it is possible to "read" an *entire* archive and analyze *every* available document and withdrawal card at the same time. We can use this power to perform a "first cut," determine the thematic topics that are statistically most prevalent, reveal what kinds of documents are particularly likely to be withheld or redacted, and rank all available documents according to their relevance to our research interests. All this can be done based on the features within the documents themselves, without presuming that we already know what topics are important or sensitive, or what terms might yield documents relevant to our research. We can then alternate from this kind of "distant reading" and other advanced methods to close reading of the usual kind.[43] Only now, we can have more confidence that – if we cannot *actually* see everything – we at least do not have the tunnel vision that results from only reading the results of search queries. So the old and the new are not mutually exclusive, indeed quite the opposite.

For collections such as the CFPF that have irretrievably lost what historians would consider their archival integrity, computational methods may offer the only hope of creating order from the chaos and producing anything like a proper finding aid. Lawyers have already discovered this when faced with huge corpora of documents produced through legal discovery. There is now a multi-billion dollar industry devoted to "e-discovery," albeit one that closely protects its intellectual property. Journalists, who write the first draft of history, were the first to create free public platforms based on machine learning and NLP. These systems automatically cluster documents and organize them in virtual files and folders in ways that resemble textual archives, thus helping the individual researcher to determine where and how to begin reading.[44]

Historians can make excellent use of these systems, but we should also be helping computer scientists to develop new ones better suited to our own work. This requires collaborative research, as the digital humanists argue. Our experience suggests that historians can contribute even if they lack coding skills. Advanced work using NLP and machine learning requires much more knowledge of mathematics and computer science than all but a handful of historians are likely to possess. But as part of a strong, multidisciplinary team, historians play a critical role in defining worthwhile questions to investigate, advising on the tradeoffs of various research protocols, and determining whether the results are valid and interesting or are merely an artifact of a flawed research design. If historians do not start working together with data scientists and developers to create reliable tools for our research, we will not have any say – nor perhaps any understanding – of the compromises and assumptions entailed in putting them together.

Historians and data scientists will also have to work with the professionals who will largely determine what kinds of electronic records we will be able to access in the future. This begins with records managers, who decide whether digital collections will be more (or less) "future-proof." Archivists have been talking about these challenges for much longer than historians, and we should

enter their discussions with all due humility.[45] We also have a role to play in helping them determine what records have permanent historical significance. For instance, what in the past seemed like mundane documents on communications procedures and records management may turn out to be the most important of all, since they are an indispensable means of reconstructing how a collection came together. Destroying them is the equivalent of throwing out the owner's manual. If instead archivists keep in mind the potential for data-mining, computer scientists can more easily develop tools to help them process text collections, while historians can better understand the documents themselves. In the meantime, deleting electronic records of any kind has to be a last resort.

In addition to records managers and archivists, librarians play a critical role in helping historians and data scientists negotiate access to collections owned by private vendors. Librarians decide what digital collections to acquire, and as customers they are in the best position to communicate with vendors about the needs of researchers. The greatest need is usually to have the raw data. Vendors are usually open to this idea since they understand that new analytical tools can greatly enhance the value of their collections.

If historians can join forces with data scientists, developers, archivists, and librarians, there is the prospect of creating a vibrant new field of research. Text processing has long depended on a relatively small number of datasets, which are large, machine-readable, public, and rich with metadata, such as the Enron emails and the Internet Movie Database. The State Department cables meet all of these criteria and bear on matters of great and enduring significance. As more and more government communications are released, slowly but still decades sooner than most private or corporate emails, there will be many more such datasets. And because there is a sizable scholarly community devoted to their study both now and stretching into the future (unlike, for instance, the Enron emails), it will be easier to develop and pursue a joint research agenda likely to result in original and important discoveries.[46]

It is still very early days for this kind of computational history, but historians understandably already want to know whether it will sway long-standing debates. Those who expect an algorithm to answer a question like "Who was to blame for the Cold War?" or "Would Kennedy have withdrawn from Vietnam?" will likely be disappointed. Answering these types of questions will always require deep and nuanced reading of key documents. But when we face many millions of electronic records, as we will when trying to reconstruct debates about the Iraq and Afghanistan Wars, even knowing what to read, and what it represents, may require computational methods.

Meanwhile, there are other questions that more easily lend themselves to computational analysis, questions that historians have neglected because, until now, they have been too difficult to answer. For instance, how, without large-scale analysis, could we rank how much attention policymakers devoted to different issues and different areas, considering that every historian tends to

claim that it was their subject which was at the "core" or "center" of it all? Similarly, without computational methods, how could we measure the bias in the historical record created by official secrecy, when most historians never mention what they could *not* read because of redactions or withholding?

What follows, unlike the other chapters in this volume, is not a summary of the state of the field. The field barely exists as yet. Instead, it is an introduction to new approaches that should prove particularly useful for historians of American foreign relations. Each technique has the potential to combine the idea of a cold start in a new archive with the sense of fortuitous discovery familiar to all historians. What has changed is that finding patterns and anomalies will increasingly require at least some "distant reading" of thousands or even millions of documents.

FIELDS OF RESEARCH

Counting

Counting is the simplest kind of computation, but it can help to answer some fundamental, neglected questions. In the absence of tools to turn words into data, historians have already resorted to using search engines to tabulate the number of references to this or that historical term.[47] Aside from the problem of corrupted text due to imperfect OCR, this method does not take into account how a corpus changes over time. If newspapers grow in size, for instance, the frequency (if not the proportion) of most terms will also increase. The Google Ngram Viewer would appear to solve this problem, since it displays word frequency relative to other words published in a given year in the Google Books corpus.[48] But it does not allow users to see what part of the corpus is being quantified and graphed. This is a fatal flaw for historians who want to understand the nature of their sources before building arguments on top of them.[49]

These problems can be overcome.[50] When they are, international historians will want to know how the volume of diplomatic activity changed over time, and where it tended to focus. They will likely debate whether the number of telephone conversations or cables transmitted provides a way to measure interest rather than just activity. With metadata, there will be new layers of analysis, and new questions. For instance, which embassies dealt with the most information that was classified secret, and why (Figure 1)? Did that reflect the sensitivity of those communications, or is it an indicator of which embassies tended to overclassify?

Once it becomes easier to turn words into data, new debates will then begin about how to interpret the data. The relative frequency with which diplomats use the term "human rights" in confidential communications, as one example, may or may not indicate whether human rights were a priority in foreign policy, much less whether that was predictive of how the United States would treat a

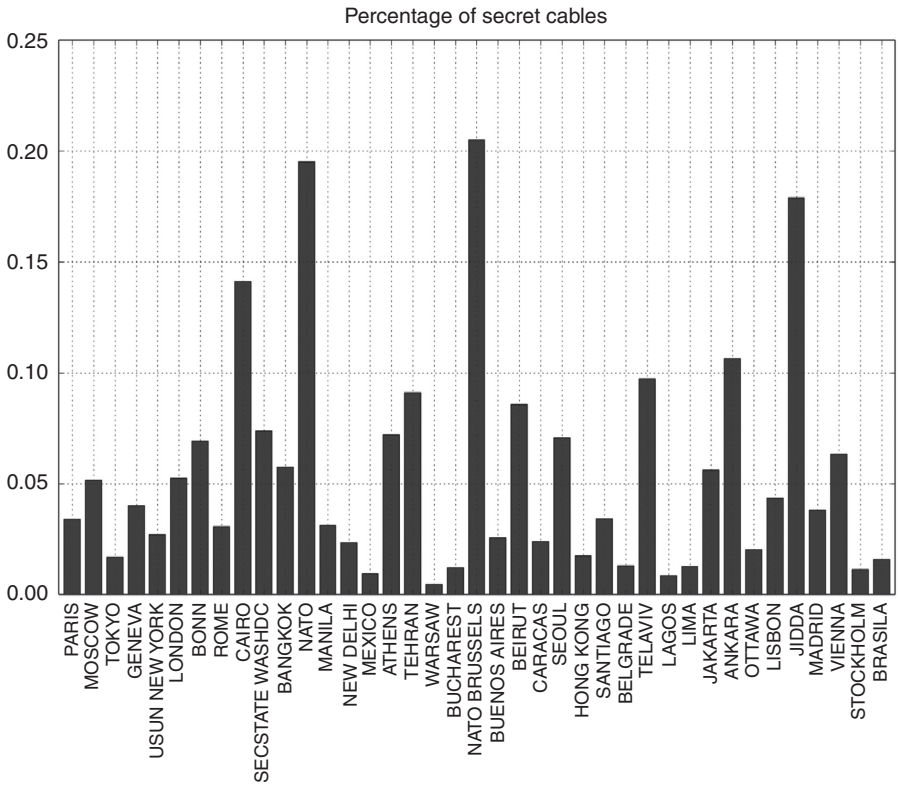

FIGURE 1. *The percentage of cables, sent by American embassies and offices around the world, originally classified secret, 1973–6. The embassies, on the x-axis, are arranged in descending order of total volume of communications, from left to right. Source: Cleaned database of fully declassified telegrams, National Archives and Records Administration, Record Group 59, Central Foreign Policy Files, 1973–8, Access to Archival Databases (AAD)*, aad.archives.gov/aad/series-description.jsp?s=4073&cat=WR43&bc=,sl. *We thank Daniel Krasner for the production of this graph.*

friendly dictator. It does show whether, when, and to what extent they believed human rights were worthy of discussion, something human rights scholars fiercely debate.[51] This kind of data is becoming more readily available, so it will become hard to ignore, and standards of evidence may begin to shift.

Traffic analysis

Counting is useful, but some historical phenomena require more sophisticated kinds of quantitative analysis. As William Sewell has written, we have always needed better ways to understand "lumpy, uneven" time, to describe how the pace of history appears to speed up or slow down.[52] The rate of

https://doi.org/10.1017/CBO9781107286207.006 Published online by Cambridge University Press

communications is a seemingly obvious measure of diplomatic activity. But when we plot graphs to reveal these "lumps," do we segment them by year, month, week, day, or hour? No one choice is more objective than another, but too small an increment will reveal many blips, rather than real bursts of activity. Conversely, if the increment is too long, these bursts will become invisible because they average out over time.

One approach that helps to resolve this conundrum is precisely to focus on this quality of "burstiness" in streams of communications.[53] The idea is to segment time into levels of activity as measured by the observed interval between cables. We demarcate the beginning and end of episodes by the escalation or de-escalation of that activity. With this model, we can then identify "bursty" time spans across the entirety of a collection, or between two embassies, and see how they relate to events. This allows us to precisely map the duration of a crisis, and to compare what we find to public assertions of what was going on. For instance, plotting the "burstiness" of communications between Vietnam and Washington between 1974 and 1976 shows how "bursts" track the dates of significant military events and how the final crisis had started long before the Ford administration admitted as much. Only a deeper dive into the documents shows that, in the end, embassy communications were mainly about refugees (Figure 2).

This is an experimental approach, but it already shows that "bursts" of activity can be a function of how the communications flow changes when a secretary of state moves through the network. When he or she goes to a foreign capital, some communications that would ordinarily have been internal to Foggy Bottom become external, leading to heightened cable traffic. How did that shift in communications affect decision-making? It will take more experimentation before this method can reveal unstudied events. We might, for instance, be able to determine whether there are particular types of cables or language within cables that are predictive of bursts of activity, or how particular kinds of metadata interact with the text. We know that intelligence agencies model and measure "chatter" to predict terrorist attacks. We ought to be able to develop our own models for diplomatic communications to model other kinds of events for historical purposes.

Topic modeling

Search engines can be a powerful way to explore document collections, but what if we do not know the exact terms we should be searching for? We may not want to presume that we already know what the "keywords" will be, or the main topics in an archive we have never analyzed before. In probabilistic topic modeling, computer scientists attempt to find the hidden thematic structure in large sets of documents. There are various kinds of models, the most common of which is Latent Dirichlet Allocation (LDA).[54] Topic modeling finds words that are likely to relate to each other statistically, and turns them into strings of

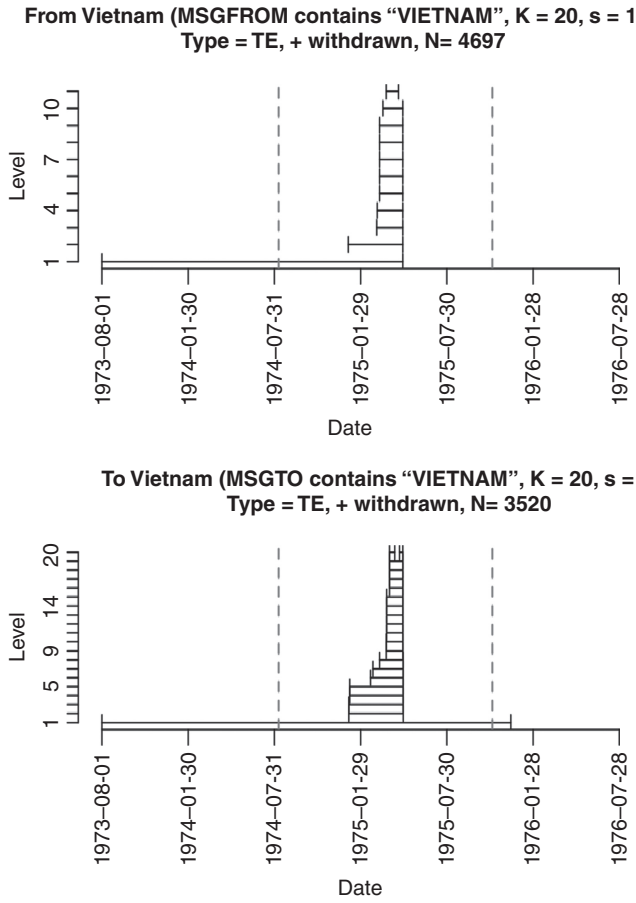

FIGURE 2. *"Burstiness" graphs showing communications and the fall of Saigon, August 1973–August 1976. These two graphs are produced by statistical work that charts traffic analysis. In both cases, a higher "burst" level represents heightened speed of communications, in and out of Vietnam. The upper limit of a "burst" is represented by "k," while "s" is a measure of graphical smoothness. N is calculated by adding the total number of fully declassified cables sent "to" or "from" Vietnam over the period to the number of still-classified cables for which we have the metadata. Dotted lines give dates for the resignation of President Richard Nixon and the firing of Henry Kissinger as secretary of state. Increased "burstiness" tracks important military events as North Vietnamese troops moved south, as well as resulting refugee crises. The precipitous stop occurs on April 30, 1975, with the evacuation of the American embassy in Saigon. Source: Cleaned AAD. We thank Shawn Simpson for the statistical work that produced these graphs.*

https://doi.org/10.1017/CBO9781107286207.006 Published online by Cambridge University Press

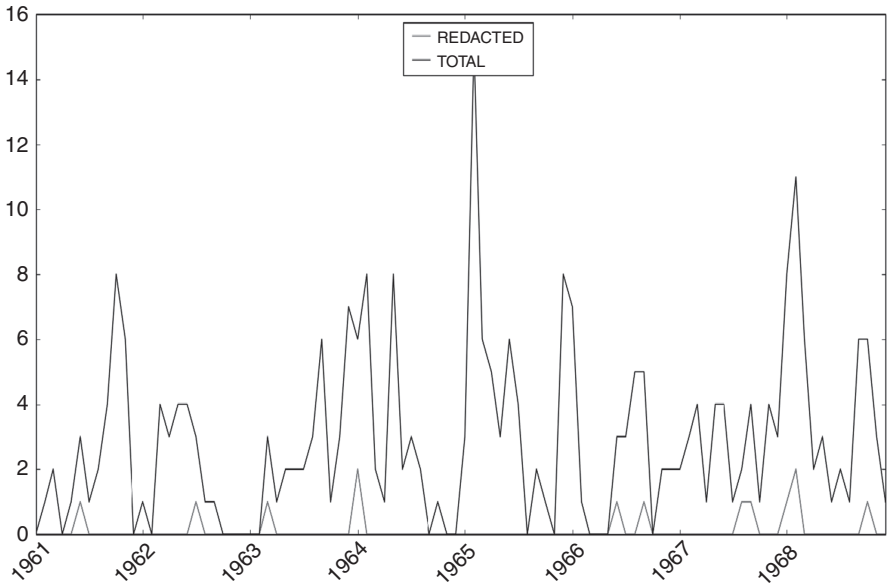

FIGURE 3. *Topic modeling* FRUS. *Topic models find an underlying, hidden structure in text. In this instance, one model of a selection of* FRUS *volumes relating to the Kennedy and Johnson administrations finds that the words listed below the graph are likely to co-occur with one another, and that the word "Vietnam" is most representative of this topic. The graph shows the prevalence of this topic over time, in terms of number of documents, as well as the number of those documents that remain redacted in published* FRUS *volumes. Source:* FRUS, *e-pub versions. We thank Thomas Nyberg for the work behind this graph.*

probabilities that make up a theme (Figure 3). To model a corpus, it assumes that a certain number of topics *generated* the documents, which reverses our usual intuition, and that each of the documents in the corpus therefore represents those topics to a greater or lesser extent. It is an "unsupervised" technique, which means that, once the parameters are specified, the algorithm is autonomous, automatically generating the combination of topics that provide the best solution.

Computers cannot actually recognize which topics are meaningful to an historian. Some, for instance, will merely represent words that frequently co-occur, such as conjunctions and modifiers. Here again, these common words, or "stopwords," must be excised with discretion and a broad historical imagination. But once the parameters are defined, one can generate a set of topics that is statistically representative of the underlying themes in a collection. When we recognize the interrelationship of terms our inclination is to label

them, something humans can do much better than machines.[55] This "science" is therefore both probabilistic and interpretive all the way down.

Topic modeling is a fairly new and quickly developing field in natural language processing, and the computer scientists doing such research are keen to work with historians to develop better ways of interacting with texts. Models have been developed that take into account historical change, networks of documents, syntax, "burstiness," and even, most interestingly, the influence of particular documents on topic distributions.[56] When used appropriately, topic modeling finds the hidden intellectual structures in our documents. Imagine, for instance, that we wanted to study how the term "national security" has changed meaning over time. Simple searching for the term in various databases will be time-consuming and often frustrating. Topic modeling has the potential not only to identify documents in the archive that are thematically related to national security, whether or not the "keywords" are actually used, but to show how it changed over time. Other applications might include showing how the language of public diplomacy differs from that of private diplomacy, and how certain topics tend to be more highly classified, or take longer to declassify.

Going "off-topic"

This kind of "distant reading" can also be used to find anomalies. Anomalies are one way that we might replicate the accidental archival discovery in the digital era. So, if we take the CFPF from the Kissinger years and apply topic modeling to all the telegrams, we find that individual embassies have specific signatures (Figure 4). The Moscow embassy talks a lot about the USSR, and very rarely about anything else, such that we can accurately predict that a cable originated from Moscow 98% of the time merely from how the words therein represent the typical topic distribution for Moscow. The London embassy, however, serves as a clearinghouse for multiple issues (Europe, NATO, trade policy, and so on), so its signature is much more diverse. If we know that particular embassies have particular signatures, we can see what happens when diplomats go "off" topic, that is, when they use words highly uncharacteristic of the embassy where they are posted.[57]

In a first run of this method, for instance, we found an unusual backchannel communication between the Soviet and American political counselors in Paris, discussing the Arab–Israeli war of 1973.[58] We also found examples of cables sent from unexpected embassies, such as a report on a Kissinger meeting with Willy Brandt sent once the Secretary had landed in Moscow, and a piece of Kremlinology pertaining to Leonid Brezhnev's power in the Politburo, relayed through the Finnish representative to the talks, held in Geneva, that led to the Helsinki Accords.[59] This method could apply not just to space, but also to time. We might, for instance, be able to estimate when a document was written, thereby helping to find documents that represent certain themes in

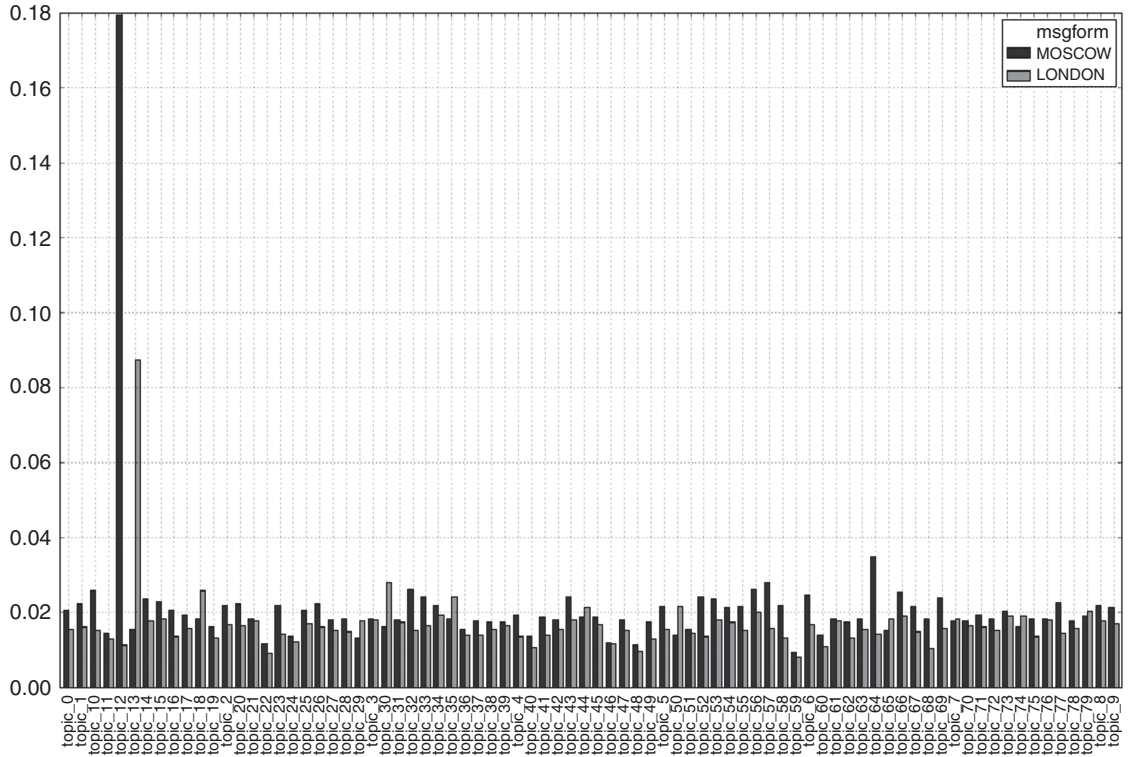

FIGURE 4. *Topic modeling of the State Department cables as a corpus reveals that individual embassies have specific topic signatures. This graph shows, from one model, the proportion of documents sent from Moscow and London that are most associated with a given topic. In this case, eighteen percent of cables sent from Moscow are predominantly about a topic that begins "Soviet, Moscow, October." Nine percent of cables from London are about a topic that begins "London, Bonn, Rome." These signatures can be used to predict from where a given cable "should" have originated. If the text of a cable does not match the usual signature of its sender, a cable is said to be "off" topic, and might be worth close reading. Source: Cleaned AAD. We thank Ian Langmore for the production of this graph.*

foreign policy before or after they become especially common. By trying to predict where a document came from, and when, we can find the unpredictable.

Authorship attribution

If topic modeling is very new, authorship attribution is very old. It dates back to the medieval scholastics, who tried to find ancient authorities for their documents. More recently, common statistical problems have involved verifying Shakespeare plays and the writers of individual Federalist papers. In their famous 1963 article on the latter, Frederick Mosteller and David L. Wallace used the frequency of words like "to," "from," and "upon" to assign contested articles to Alexander Hamilton or James Madison.[60] Since then, computers have made authorship attribution more rapid, powerful, and accurate.[61]

What we might expect from this kind of research depends on the nature of the problem. A solution will be easier if there is a small pool of potential authors, such as a presidential speech or an embassy telegram sent under an ambassador's name. Or we might have a much larger number of candidates, such as an anonymous memorandum that might have been written by one of thousands of different foreign service officers. With National Security Council or Policy Planning Staff memoranda, we might want to determine who wrote specific sections of a given document. In all of these cases, we would need authenticated examples of the writing of all of the candidates. But even if we have no idea who wrote a given text, authorship attribution techniques are reasonably accurate in determining their age and gender. All of these questions are ongoing fields of research in statistics that can help to solve longstanding questions about our sources.

Network analysis

Another well-developed field in statistics and computer science is social network analysis. The idea here is to analyze large collections of texts through the social network that is evoked in them, in ways that build on economic and sociological theory.[62] A network is a collection of people ("nodes") connected by links ("edges"). The most pertinent examples for historians of American foreign relations are the bureaucrats in government departments, and the networks of informants that embassies use to gain intelligence about the country on which they report. Network analysis provides us with ways to view structures like this across a corpus, and to see how the actions of one node have consequences for the rest of the system. Once we have extracted the network in historically representative ways, we can then begin to model it in a way that helps us to read the documents anew.

Franco Moretti has already demonstrated the power of network analysis in analyzing individual plays, using his models to question traditional

concepts in the study of literature. What happens, he asks, when you take Hamlet or Claudius out of *Hamlet*?[63] How might literary scholars define the "centrality" of a character?[64] Once historians of American foreign relations have large corpora of documents, we can imagine new approaches to classic questions, beginning with the role of particular agents within larger structures. How, for instance, did networks reflect and affect the flow of ideas and information through a bureaucracy? How did the relative connectedness or isolation of individuals and institutions help determine policy outcomes? We might study how networks developed in ways that would not be expected by the organizational charts of the State Department or the federal government, such as informal networks that formed around particular issues and pushed specific policies. We could investigate how networks of local informants to a diplomatic mission changed before and after a change of government. Modeling would not provide definitive answers, of course, but it should prompt useful questions that might not otherwise occur to us.

Mapping

Of all the digital techniques outlined here, the most established in history is that of mapping. Geographical Information Systems (GIS) have been a key part of historical geography for many years.[65] Its application has been part of a "spatial turn" in certain strands of historiography, especially environmental history.[66] For instance, GIS was used in Geoff Cunfer's innovative history of the Great Plains, in tools like the *Digital Atlas of Roman and Medieval Civilizations* and *A Vision of Britain through Time*, and in an early digital history project led by William Thomas and Edward Ayers, "The Differences Slavery Made."[67] Now, with the pioneering work of Stanford's Spatial History Lab, what Richard White calls "spatial history" seems ripe for takeoff.[68] The historiography of American foreign relations has not quite taken a "spatial turn" in the sense of adopting, as White and others have, the ideas of Henri Lefebvre and other geographers.[69] But international historians have always been interested in space in one way or another.

GIS allows for the mapping of any data point that has a corresponding location, as so many of our sources do. Once we have that kind of data, all kinds of maps can be placed on top of one another, adjusted to show change over time, and so on. Mapping brings the prospect of visualizing some of the most basic aspects of international history. One of our first efforts, as part of our "Declassification Engine" project at Columbia University, has been to begin mapping cable traffic month by month (Figure 5).[70] From there, one can imagine overlaying material capabilities with military potential and alliance structure, for instance, or comparing those capabilities with the amount of bureaucratic attention paid to the countries in question. Classic questions of international history will suddenly become, quite literally, visible.

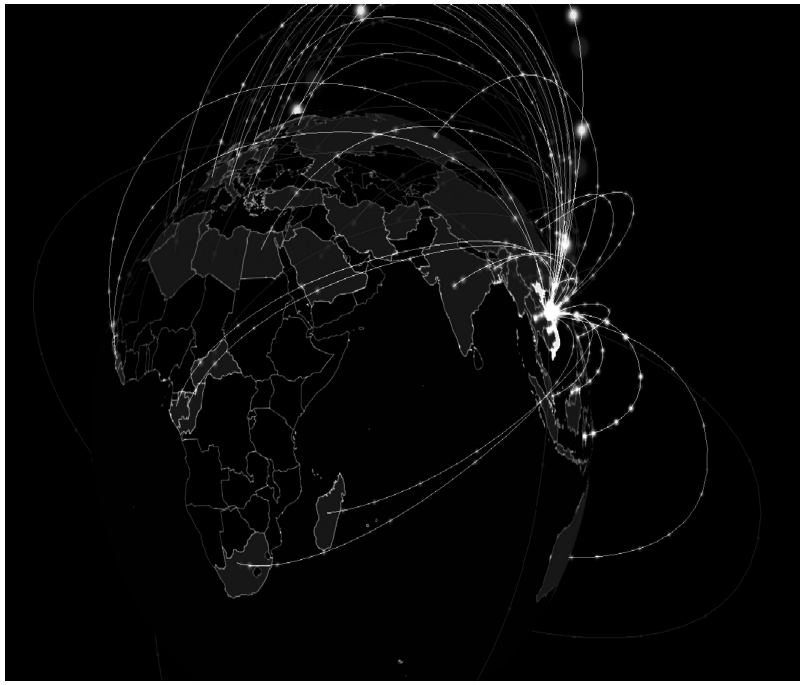

FIGURE 5. *Part of the Declassification Engine's "Sphere of Influence" project, this map shows the countries with which the American embassy in South Vietnam was communicating just before the fall of Saigon. Source: Cleaned AAD. We thank Kalev Leetaru and Dainis Kiusals for creating this map. For more, see* declassification-engine.org/index.py?section=sphere#.

CONCLUSION

Historians might understandably be nervous about the idea of having to learn an entirely new set of methods to explore contemporary diplomatic archives. In fact, at least some of the aforementioned approaches can and will produce tools that anyone will be able to use. But it could take some time. Meanwhile, there are tremendous opportunities for new discoveries to be made by historians who are willing to work collaboratively. To improve certain methods, such as network extraction and topic modeling, computer scientists need to consult with historians and others with deep knowledge of the documents. This research tends to be an iterative process, in which methods are continually refined to produce results that are both valid and significant. While computer scientists necessarily focus on making new discoveries in their own field – and are not usually eager merely to apply known techniques – many look to other disciplines to demonstrate that their work really does help us better to understand real-world phenomena. As computer scientist David Blei writes,

even if new statistical models "are meant to help interpret and understand texts," the models, and the texts they use, still need to be evaluated. "Using humanist texts to do humanist scholarship," he concludes, "is the job of a humanist."[71]

We as historians can do a better job once we realize, as even anti-cliometricians did four decades ago, that each time we write "most" or "likely" we are making quantitative or probabilistic judgments. At least some of those judgments could now be made more precisely, and in a way that can either be validated or disproven. Unlike a lot of the quantitative data used in social science research, ours will not come from coding by research assistants or self-reporting in surveys. We do not, in other words, have to join the international relations researchers who struggle to rate wars on a scale of one to five, or fall into the fallacy sociologists commit when they equate attitudes recorded in polls with actual behavior. We have the immense advantage of using primary sources, and can now use them in a whole new way. Even if we can never use the whole corpus, we have enough of it to mitigate the selection bias and out-of-sample issues that bedevil other disciplines. We need not become obsessed with running regressions or pursuing statistical significance as an end in itself. Instead we can combine computational methods with our traditional strength in closely reading our sources and attending to their context.

Whenever the next edition of *Explaining the History of American Foreign Relations* is published, much of this chapter will likely seem dated. That, at least, is our hope. Our expectation is that more and more digitized and born-digital documents will become available, more even than we envisage here, especially if historians begin to pool the scans and photographs we take in the course of our research into a virtual archive. We presume too that advances in data science will continue at an even more rapid pace, too rapid for us to imagine all the applications. Above all, we hope that historians of world politics take up these methods, determine which have practical utility, and help to develop new ones. Some will be found wanting, but we cannot be afraid to fail. If we do not at least try to come up with new means to cope with the infinite archive, we will not even realize what we have missed. If we start to work together and work across disciplines, we can begin an exciting new period of experimentation, and perhaps even lead the way in the reinvention of history as a data science.

NOTES

* As this chapter explains, the authors are part of a multidisciplinary team that has been exploring applications for text processing and machine learning in diplomatic history. Much of what we have learned has come from collaboration, especially with David Madigan, Daniel Krasner, Ian Langmore, Sasha Rush, Shawn Simpson, Thomas Nyberg, and Rex Douglass. Their individual contributions are also noted in the discussion of specific research

methodologies. Note that the graphs were produced with historical data but are provided for illustrative purposes only. We would also like to acknowledge the generous support of the Brown Institute for Media Innovation and the John D. and Catherine T. MacArthur Foundation.

1. Michael J. Hogan and Thomas G. Paterson (eds.), *Explaining the History of American Foreign Relations* (New York, 1991).
2. William McAllister, "The Documentary Big Bang, the Digital Records Revolution, and the Future of the Historical Profession," *Passport* 41:2 (September 2010): 12–17.
3. "Interchange: The Promise of Digital History," *Journal of American History* 95 (2008): 455.
4. See, for instance, Public Interest Declassification Board, "Transforming the Security Classification System" (December 2012), archives.gov/declassification/pidb/recommendations/transforming-classification.html.
5. William O. Aydelotte, "Quantification in History," *American Historical Review* 72 (1966): 803–25.
6. Robert P. Swierenga, "Computers and American History: The Impact of the "New" Generation," *Journal of American History* 60 (1974): 1045–70.
7. Robert W. Fogel, "The New Economic History," *Economic History Review* 19 (1966): 642–56; Robert W. Fogel, *Railroads and American Economic Growth: Essays in Econometric History* (Baltimore, 1964); Avner Greif, "Cliometrics After 40 Years," *American Economic Review* 87 (1997): 400–3; Douglass C. North, "Cliometrics—40 Years Later," *American Economic Review* 87 (1997): 412–14.
8. François Furet, "Quantitative History," *Daedalus* 100 (1971): 151–67; Emmanuel Le Roy Ladurie, *The Territory of the Historian* (Chicago, 1979).
9. Swierenga, "Computers and American History"; Jerome M. Clubb and Howard Allen, "Computers and Historical Studies," *Journal of American History* 54 (1967): 599–607; William G. Thomas, III, "Computing and the Historical Imagination," in Susan Schreibman, Ray Siemens, and John Unsworth (eds.), *A Companion to Digital Humanities* (Oxford, 2004), digitalhumanities.org/companion/.
10. See several of the articles in Robert P. Swierenga (ed.), *Quantification in History: Theory and Research* (New York, 1970).
11. Carl Bridenbaugh, "The Great Mutation," *American Historical Review* 68 (1963): 326.
12. Arthur Schlesinger, Jr., "The Humanist Looks at Empirical Social Research," *American Sociological Review* 27 (1962): 770.
13. Robert William Fogel and Stanley L. Engerman, *Time on the Cross: The Economics of American Negro Slavery* (Boston, 1974).
14. Alfred H. Conrad and John R. Meyer, "The Economics of Slavery in the Ante Bellum South," *Journal of Political Economy* 66 (1958): 95–130.
15. Herbert G. Gutman, "The World Two Cliometricians Made: A Review-Essay of F + E = T/C," *Journal of Negro History* 60 (1975): 53–7; C. Vann Woodward, "The Jolly Institution," *New York Review of Books* 21 (May 2, 1974); Eric Foner, "Redefining the Past," *Labor History* 16 (1975): 127–38.
16. Robert William Fogel, "The Limits of Quantitative Methods in History," *American Historical Review* 80 (1975): 329–50; Charlotte Erickson, "Quantitative History," *American Historical Review* 80 (1975): 351–65.
17. Furet, "Quantitative History," 160.

18. Joyce Appleby, "The Power of History," *American Historical Review* 103 (1998): 5–6.
19. Matthew K. Gold, "Introduction: 'The Digital Humanities Moment'," in Matthew K. Gold (ed.), *Debates in the Digital Humanities* (Minneapolis, 2012), ix. See also Anne Burdick, Johanna Drucker, Peter Lunenfeld, Todd Presner, Jeffrey Schnapp, *Digital_Humanities* (Cambridge, MA, 2012), and the collection edited by one of history's few representatives in the digital humanities, Daniel J. Cohen: Daniel J. Cohen and Tom Scheinfeldt (eds.), *Hacking the Academy: New Approaches to Scholarship and Teaching from Digital Humanities* (Ann Arbor, 2013).
20. Roy Rosenzweig, "The Road to Xanadu: Public and Private Pathways on the History Web," *Journal of American History* 88 (2001): 548–579. Rosenzweig was the key figure in earlier moves towards digital history. See Roy Rosenzweig, "'So, What's Next for Clio?' CD-ROM and Historians," *Journal of American History* 81 (1995): 1621–40; Roy Rosenzweig and Michael O'Malley, "Brave New World or Blind Alley? American History on the World Wide Web," *Journal of American History* 84 (1997): 132–55; Roy Rosenzweig, "Scarcity or Abundance? Preserving the Past in a Digital Era," *American Historical Review* 108 (2003): 735–62; Roy Rosenzweig, "Can History Be Open Source? Wikipedia and the Future of the Past," *Journal of American History* 93 (2006): 117–46; Daniel J. Cohen and Roy Rosenzweig, *Digital History: A Guide to Gathering, Preserving, and Presenting the Past on the Web* (Philadelphia, 2006); Roy Rosenzweig, *Clio Wired: The Future of the Past in the Digital Age* (New York, 2010).
21. Orville Vernon Burton, "American Digital History," *Social Science Computer Review* 23 (2005): 206–20.
22. "Interchange: The Promise of Digital History," 488, 442–51, 461, 464.
23. Cameron Blevins, "The Perpetual Sunrise of Methodology," January 5, 2015, cameronblevins.org/posts/perpetual-sunrise-methodology/.
24. Franco Moretti, *Graphs, Maps, Trees: Abstract Models for Literary History* (London, 2007); Franco Moretti, *Distant Reading* (London, 2013); Matthew L. Jockers, *Macroanalysis: Digital Methods and Literary History* (Urbana, 2013); Stephen Ramsay, *Reading Machines: Toward An Algorithmic Criticism* (Urbana, 2011). Moretti and Jockers work with the Stanford Literary Lab, the publications of which can be found at litlab.stanford.edu/pamphlets. For the backlash, see Stephen Marche, "Literature is Not Data: Against Digital Humanities," *Los Angeles Review of Books*, October 28, 2012, lareviewofbooks.org/essay/literature-is-not-data-against-digital-humanities; Scott Selisker and Holger S. Syme, "In Defense of Data: Responses to Stephen Marche's 'Literature is Not Data'," *Los Angeles Review of Books*, November 5, 2012, lareviewofbooks.org/essay/in-defense-of-data-responses-to-stephen-marches-literature-is-not-data.
25. See the "Humanities 2.0" series in the *New York Times*, especially Patricia Cohen, "Digital Keys for Unlocking the Humanities' Riches," *New York Times*, November 16, 2010, nytimes.com/2010/11/17/arts/17digital.html, and Patricia Cohen, "Digital Maps Are Giving Scholars the Historical Lay of the Land," *New York Times*, July 26, 2011, nytimes.com/2011/07/27/arts/geographic-information-systems-help-scholars-see-history.html.
26. Google Books, books.google.com; "Google Books History," google.com/googlebooks/about/history.html; HathiTrust Digital Library, hathitrust.org.

27. "Public Report of the Advisory Committee," 1961, John F. Kennedy Presidential Library, Arthur M. Schlesinger, Jr., Personal Papers, Box WH-12.
28. Gerald K. Haines and J. Samuel Walker, "Some Sources and Problems for Diplomatic Historians in the Next Two Decades," in Gerald K. Haines and J. Samuel Walker (eds.), *American Foreign Relations: A Historiographical Review* (Westport, 1981), 335–6.
29. McAllister, "The Documentary Big Bang," 16.
30. Mapping the Republic of Letters, "Visualizing Benjamin Franklin's Correspondence Network," republicofletters.stanford.edu/casestudies/franklin.html.
31. *Foreign Relations of the United States*, State Department Office of the Historian, history.state.gov/historicaldocuments; *Foreign Relations of the United States*, University of Wisconsin Digital Collections, uwdc.library.wisc.edu/collections/FRUS.
32. See, for instance, Central Intelligence Agency, "Freedom of Information Act Reading Room," foia.cia.gov.
33. Department of State, "Freedom of Information Act," foia.state.gov.
34. National Archives and Records Administration, "The Remote Archives Capture Program (RAC)," archives.gov/presidential-libraries/declassification/rac.html.
35. Rosenzweig, "The Road to Xanadu," esp. 564–77.
36. Digital National Security Archive, nsarchive.gwu.edu/publications/dnsa.html.
37. ProQuest History Vault, hv.conquestsystems.com/historyvault/.
38. Declassified Documents Reference System, galenet.galegroup.com/servlet/DDRS.
39. National Archives and Records Administration, Record Group 59, Central Foreign Policy Files, 1973–8, Access to Archival Databases (AAD), aad.archives.gov/aad/series-list.jsp?cat=WR43.
40. Information Security Oversight Office (ISOO), "Annual Report to the President" (2012), archives.gov/isoo/reports/2012-annual-report.pdf: 26.
41. David Langbart, William Fischer, and Lisa Roberson, "Appraisal of records covered by N1-59-07-3-P," June 4, 2007.
42. Calculations based on a reader working an eight-hour day with no days off.
43. This distinction between "close" and "distant" reading is borrowed from Franco Moretti. See Franco Moretti, "Conjectures on World Literature," *New Left Review* 1 (January/February 2000): 54–68; Moretti, *Distant Reading*. Matthew Jockers prefers to think in terms of "macro-" and "microanalysis," arguing that algorithms are not actually *reading* documents (a human activity if ever there was one), but analyzing them. See Jockers, *Macroanalysis*, 22–31.
44. See, for instance, "DocumentCloud," documentcloud.org/home; "The Overview Project," overview.ap.org.
45. Joshua Sternfeld, "Archival Theory and Digital Historiography: Selection, Search, and Metadata as Archival Processes for Assessing Historical Contextualization," *The American Archivist* 74 (2011): 544–75; Alexandra Chassanoff, "Historians and the Use of Primary Sources in the Digital Age," *The American Archivist* 76 (2013): 458–80.
46. Jessica Leber, "The Immortal Life of the Enron E-mails," *MIT Technology Review*, July 2, 2013, technologyreview.com/news/515801/the-immortal-life-of-the-enron-e-mails/.

47. See, for instance, James Belich, *Replenishing the Earth: The Settler Revolution and the Rise of the Anglo-World, 1783–1939* (New York, 2009), 151–2; Samuel Moyn, *The Last Utopia: Human Rights in History* (Cambridge, MA, 2010), 231.
48. "Google Ngram Viewer," books.google.com/ngrams and books.google.com/ngrams/info; Jean-Baptiste Michel, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K. Gray, The Google Books Team, Joseph P. Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, Steven Pinker, Martin A. Nowak, and Erez Lieberman Aiden, "Quantitative Analysis of Culture Using Millions of Digitized Books," *Science* 331 (January 14, 2011): 176–82, sciencemag.org/content/331/6014/176.full.html.
49. Some scholars have deployed Ngram graphs to back up their claims without appearing to be aware of this issue, an early warning of what awaits if historians do not take more of an interest in how these kinds of tools are developed. See, for example, Daniel Lord Smail and Andrew Shryock, "History and the 'Pre'," *American Historical Review* 118 (2013): 710–12.
50. For recent efforts, see Cameron Blevins, "Space, Nation, and the Triumph of Region: A View of the World from Houston," *Journal of American History* 101 (2014): 122–47, and Cameron Blevins, "Mining and Mapping the Production of Space," Spatial History Project, Stanford University, web.stanford.edu/group/spatialhistory/cgi-bin/site/pub.php?id=93; Franco Moretti and Dominique Pestre, "Bankspeak: The Language of World Bank Reports," *New Left Review* 92 (March-April 2015): 75–99.
51. See, for instance, the periodization debate present in Moyn, *The Last Utopia*; Sarah B. Snyder, *Human Rights Activism and the End of the Cold War: A Transnational History of the Helsinki Network* (New York, 2011); Elizabeth Borgwardt, *A New Deal for the World: America's Vision for Human Rights* (Cambridge, MA, 2005).
52. William H. Sewell, Jr., *Logics of History: Social Theory and Social Transformation* (Chicago, 2005), 9.
53. For the statistical basis of "burstiness," see Jon Kleinberg, "Bursty and Hierarchical Structure in Streams," *Data Mining and Knowledge Discovery* 7 (2003): 373–97.
54. For summaries of topic modeling, especially Latent Dirichlet Allocation (LDA), see Jockers, *Macroanalysis*, 118–53; David M. Blei, "Topic Modeling and Digital Humanities"; David M. Blei, "Probabilistic Topic Models," *Communications of the ACM* 55 (2012): 77–84. For the original statistical explanations, see David M. Blei, Andrew Y. Ng, Michael I. Jordan, "Latent Dirichlet Allocation," *Journal of Machine Learning Research* 3 (2003): 993–1022; Thomas L. Griffiths and Mark Steyvers, "Finding scientific topics," *Proceedings of the National Academy of Science* 101 (April 6, 2004): 5228–35. For an early historical use of topic modeling, see David J. Newman and Sharon Block, "Probabilistic Topic Decomposition of an Eighteenth-Century American Newspaper," *Journal of the American Society for Information Science and Technology* 57 (2006): 753–67.
55. Jonathan Chang, Jordan Boyd-Graber, Sean Gerrish, Chong Wang, David M. Blei, "Reading Tea Leaves: How Humans Interpret Topic Models," *Neural Information Processing Systems* (2009), cs.princeton.edu/~blei/papers/ChangBoyd-GraberWangGerrishBlei2009a.pdf. For more on the strings of words that topic modeling generates, and the coherence of the strings, see Jey Han Lau, Karl Grieser, David Newman, and Timothy Baldwin, "Automatic Labelling of Topic Models," *Proceedings of the 49th Annual Meeting of the Association for*

Computational Linguistics (June 19–24, 2011): 1536–45; Jey Han Lau, David Newman, Sarvnaz Karimi, and Timothy Baldwin, "Best Topic Word Selection for Topic Labelling," *Proceedings of the 23rd International Conference on Computational Linguistics Posters* (August 23–27, 2010): 605–13; David Newman, Jey Han Lau, Karl Grieser, and Timothy Baldwin, "Automatic Evaluation of Topic Coherence," *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the ACL* (2010): 100–8.

56. David M. Blei and John D. Lafferty, "Dynamic Topic Models," *Proceedings of the 23rd International Conference on Machine Learning* (2006): 113–20; Jonathan Chang and David M. Blei, "Relational Topic Models for Document Networks," *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics* (2009): 81–8; Jordan Boyd-Graber and David M. Blei, "Syntactic Topic Models," *Neural Information Processing Systems* (2008), cs.princeton.edu/~blei/papers/Boyd-GraberBlei2009.pdf; Thomas L. Griffiths, Mark Steyvers, David M. Blei, and Joshua B. Tenenbaum, "Integrating Topics and Syntax," *Neural Information Processing Systems* 17 (2005), psiexp.ss.uci.edu/research/papers/composite.pdf; Gabriel Doyle and Charles Elkan, "Accounting for Burstiness in Topic Models," *Proceedings of the 26th International Conference on Machine Learning* (2009), cseweb.ucsd.edu/~elkan/TopicBurstiness.pdf; Sean M. Gerrish and David M. Blei, "A Language-based Approach to Measuring Scholarly Impact," *Proceedings of the 26th International Conference on Machine Learning* (2010), cs.princeton.edu/~blei/papers/GerrishBlei2010.pdf.
57. We are greatly indebted to Ian Langmore for this idea, and for the modeling behind this technique.
58. Paris to SecState, "Middle East Situation: Soviet Embassy Approach," October 7, 1973, AAD, 1973PARIS26220.
59. Moscow to State, "SecState's Meetings with Brandt and Scheel," March 25, 1974, AAD, document number 1974MOSCOW04262; Geneva to State and Bonn, "Soviet Leadership Matters Bonn for Secretary's Party," February 15, 1975, AAD, 1975GENEVA01022.
60. Frederick Mosteller and David L. Wallace, "Inference in an Authorship Problem," *Journal of American Statistical Association* 58 (1963): 275–309.
61. For summaries of the various methods from literary scholars, see Jack Grieve, "Quantitative Authorship: And Evaluation of Techniques," *Literary and Linguistic Computing* 22 (2007): 251–70; Matthew L. Jockers and Daniela M. Witten, "A Comparative Study of Machine Learning Methods for Authorship Attribution," *Literary and Linguistic Computing* 25 (2010): 215–23; For technical summaries, see Patrick Juola, "Authorship Attribution," *Foundations and Trends in Information Retrieval* 1 (2006): 233–334; Moshe Koppel, Jonathan Schler, and Shlomo Argamon, "Computational Methods in Authorship Attribution," *Journal of the American Society for Information Science and Technology* 60 (2009): 9–26.
62. For a good summary, see David Easley and Jon Kleinberg, *Networks, Crowds, and Markets: Reasoning about a Highly Connected World* (New York, 2010), cs.cornell.edu/home/kleinber/networks-book/.
63. Franco Moretti, "Network Theory, Plot Analysis," *New Left Review* 68 (March-April 2011), newleftreview.org/II/68/franco-moretti-network-theory-plot-analysis.

64. Franco Moretti, "'Operationalizing': Or, the Function of Measurement in Literary Theory," *New Left Review* 84 (November–December 2013), newleftreview.org/II/84/franco-moretti-operationalizing.
65. Ian N. Gregory and Paul S. Ell, *Historical GIS: Technologies, Methodologies and Scholarship* (New York, 2007).
66. Anne Kelly Knowles (ed.), *Placing History: How Maps, Spatial Data, and GIS Are Changing Historical Scholarship* (Redlands, 2008); David J. Bodenhamer, John Corrigan, and Trevor M. Harris (eds.), *Spatial Humanities: GIS and the Future of Humanities Scholarship* (Bloomington, 2010).
67. Geoff Cunfer, *On the Great Plains: Agriculture and Environment* (College Station, 2005); "Digital Atlas of Roman and Medieval Civilizations," darmc.harvard.edu/icb/icb.do?keyword=k40248&pageid=icb.page188865; "A Vision of Britain through Time," visionofbritain.org.uk; William G. Thomas, III, and Edward L. Ayers, "An Overview: The Differences Slavery Made: A Close Analysis of Two American Communities," *American Historical Review* 108 (2003): 1299–1307; "The Differences Slavery Made," valley.lib.virginia.edu.
68. Richard White, "What is Spatial History?" (February 1, 2010), stanford.edu/group/spatialhistory/cgi-bin/site/pub.php?id=29. For one published example, see Caroline Winterer, "Where is America in the Republic of Letters," *Modern Intellectual History* 9 (2012): 597–623; stanford.edu/group/spatialhistory/cgi-bin/site/projects.php.
69. Henri Lefebvre, *The Production of Space* (Chicago, 1991).
70. Declassification Engine, "The Sphere of Influence," declassification-engine.org/index.py?section=sphere#. We thank Kalev Leetaru and Dainis Kiusals for putting together this project.
71. David M. Blei, "Topic Modeling and Digital Humanities," *Journal of Digital Humanities* 2 (Winter 2012), journalofdigitalhumanities.org/2-1/topic-modeling-and-digital-humanities-by-david-m-blei/.