

Infectious Texts: Modeling Text Reuse in Nineteenth-Century Newspapers

David A. Smith*, Ryan Cordell†, and Elizabeth Maddock Dillon†

**College of Computer and Information Science and* †*Department of English*
Northeastern University
Boston, MA, U.S.A.
dasmith@ccs.neu.edu, {r.cordell,e.dillon}@neu.edu

Abstract

Texts propagate through many social networks and provide evidence for their structure. We present efficient algorithms for detecting clusters of reused passages embedded within longer documents in large collections. We apply these techniques to analyzing the culture of reprinting in the United States before the Civil War. Without substantial copyright enforcement, stories, poems, news, and anecdotes circulated freely among newspapers, magazines, and books. From a collection of OCR'd newspapers, we extract a new corpus of reprinted texts, explore the geographic spread and network connections of different publications, and analyze the time dynamics of different genres.

Index Terms

clustering algorithms; nearest neighbor searches

1. Introduction

Many studies of social networks and interactions use surveys and observational data on nodes and links. This metadata is often augmented by the language the members of a social network use to interact with each other—e.g., the text of email messages within a company or of posts and tweets exchanged among online social network participants. In some cases, however, we cannot directly observe network links, or even get a census of network nodes, and yet we can still observe text that provides evidence for social interactions.

A particularly useful form of evidence comes from copying and quotation of text by different actors in a social network. The social ties indicated by text reuse may be overt, as with acknowledged quotations of earlier scholarly work, or unacknowledged, as with legislators' adoption, when drafting bills, of language from interest groups' policy documents. In aggregate,

text reuse patterns can provide evidence for direct or indirect contact between different social actors. Examination of the reused passages can provide evidence for the shared interests of groups. At a finer level of analysis, the temporal and spatial dynamics of text reuse can provide evidence for individual links and propagation patterns in networks [1, 2].

Our approach to local text reuse detection and clustering is motivated by analysis of a corpus of nineteenth-century U.S. newspapers to investigate the “culture of reprinting” that existed before effective copyright enforcement and wire services. While this paper focuses on the textual evidence alone, a more general setup could employ both text and information about some *related* network. Namata et al. [3], for instance, considered the task of inferring corporate hierarchies from the network of email communications. In the case study of 19c newspapers, we might use the railroad and telegraph networks as additional evidence for how news and culture propagated.

Researchers in NLP and information retrieval have often employed text reuse detection to remove duplicates from web crawls and search results and to detect plagiarism in documents and source code. These methods can achieve quite high precision when the majority of documents' contents are implicated in reuse. Performance is often tuned for ad-hoc retrieval, with one or a few query documents (e.g., student submissions) being run against a large corpus of potential matches. We are interested, however, in the overall patterns of text reuse and focus here on a different task: all-pairs near-neighbor search for overlapping passages. For a web crawler, lower recall simply leads to a larger index; in our case, low recall on the text reuse task can lead to inaccurate output.

At a high level, the method proposed in this paper (1) finds pairs of documents likely to contain substantial overlap, (2) performs an alignment to find passages of local similarity even if the rest of the documents are unrelated, and (3) uses that pairwise passage informa-

tion to infer larger clusters of text reuse. As we explain in more detail below, there are two primary sources of error in this process, in addition to the appearance of resemblance arising by chance that affects other work on sequence similarity from NLP to bioinformatics. First, since our collections contain multiple documents from the same source, precision can be hurt by often substantial amounts of intra-source text reuse, which are not as important when looking for evidence of interaction among sources. Second, since we are often interested in mapping networks in historical data, text that has been poorly transcribed by optical character recognition can depress recall.

Our approach, like many others, involves “shingling”—evaluating document pairs by the intersection of their n-gram features—, but these considerations require further refinements to optimize effectiveness and efficiency. In particular, we employ hashing for space-efficient indexing of repeated n-grams and the use of overlapping n-grams to prune dynamic programming for local alignment.

Before describing methods for detecting and clustering passages of local text reuse in §3, we provide background on reprinting in 19c newspapers in order to motivate some of our design choices (§2). In §4, we describe an experimental evaluation of the effectiveness and efficiency of various text reuse detection approaches. Finally, we investigate interesting properties of the most popular texts, explore the geographic spread and network connections of different texts and publications, and analyze some of the text reuse dynamics of different genres of texts (§5).

2. Tracking Viral Texts in 19c Newspapers

In *American Literature and the Culture of Reprinting*, McGill [4] argues that American literary culture in the nineteenth century was shaped by the widespread practice of reprinting stories and poems, usually without authorial permission or even knowledge, in newspapers, magazines, and books. Without substantial copyright enforcement, texts circulated promiscuously through the print market and were often revised by editors during the process. These “viral” texts—be they news stories, short fiction, or poetry—are much more than historical curiosities. The texts that editors chose to pass on are useful barometers of what was exciting or important to readers during the period, and thus offer significant insight into the priorities and concerns of the culture. Nineteenth-century U.S. newspapers were usually associated with a particular political party, religious denomination, or social cause (e.g., temperance or abolition). Mapping the specific locations and

venues in which varied texts circulated would therefore allow us to answer questions about how reprinting and the public sphere in general were affected by geography, communication and transportation networks, and social, political, and religious affinities.

To study the reprint culture of this period, we crawled the online newspaper archives of the Library of Congress’s *Chronicling America* project (chroniclingamerica.loc.gov). Since the *Chronicling America* project aggregates state-level digitization efforts, there are some significant gaps: e.g., there are no newspapers from Massachusetts, which played a not insubstantial role in the literary culture of the period. While we continue to collect data from other sources in order to improve our network analysis, the current dataset remains a useful, and open, testbed for text reuse detection and analysis of overall trends.

Another difficulty with this collection is that it consists of the OCR’d text of newspaper issues without any marking of article breaks, headlines, or other structure. The local alignment methods described in §3.2 are designed not only to mitigate this problem, but also to deal with partial reprinting. One newspaper issue, for instance, might reprint chapters 4 and 5 of a Thackeray novel while another issue prints only chapter 5.

Since our goal is to detect texts that spread from one venue to another, we are not interested in texts that were reprinted frequently in the same newspaper, or *series*, to use the cataloguing term. This includes material such as mastheads and manifestos and also the large number of advertisements that recur week after week in the same newspaper.

The methods we discuss in this paper could also be applied, of course, to other domains. Most closely related to previous work on plagiarism detection is the investigation in historical and literary studies of the sources that influenced a text, sometimes referred to as *source criticism*.

3. Text Reuse Detection

As noted above, we are interested in detecting passages of text reuse (individual articles) that comprise a very small fraction of the containing documents (newspaper issues). Using the terminology of biological sequence alignment, we are interested in *local alignments* between documents. Henzinger [5] provides a good overview of text reuse detection research and provides an empirical evaluation of the two primary methods—n-gram shingling and locality-sensitive hashing (LSH)—on near-duplicate webpage detection tasks. The need for local alignments makes LSH less practical without performing a large number

of sliding-window matches.¹ We therefore base our approach on n-gram shingling.

Several attributes distinguish the present approach from previous work:

- The boundaries of the reused passages are not known, in contrast to near-duplicate document detection and to work on “meme tracking” that takes text between quotation marks as the unit of reuse [1, 7].
- Also in contrast to this work on the contemporary news cycle and blogosphere, we are interested both in texts that are reprinted within a few days and after many years. We thus cannot exclude potentially matching documents for being far removed in time.
- We are looking for reuse of substantial amounts of text, on the order of 100 words or more, in contrast to work on detecting shorter quotations [8, 9, 10, 2].
- We wish to compute all nearest neighbor passages, instead of running a small number of query documents against a corpus.
- Text reuse that occurs only among documents from the same “source” should be excluded. Henzinger [5] notes, for instance, that many of the errors in near-duplicate webpage detection arose from false matches among documents from the same website that shared boilerplate navigational elements.
- We require methods that are robust to noisy OCR transcripts. While we could adopt a pipelined approach and perform OCR correction before text reuse detection, it is worth noting that a promising source of data for OCR correction are the clusters of similar passages found in other documents.

In common with work in duplicate detection, we start the search for reused passages by first finding likely document pairs. For each document pair returned by this first step, we then identify a passage embedded within each document that is a good local match for the other. We then greedily aggregate this pairwise data into large clusters of reused text passages. We now describe the details of each step in turn.

3.1. Search for Candidate Document Pairs

We begin with a scan for document pairs likely to contain reprinted passages. We use “shingling”, which represents documents by an unordered collection of its

1. While there is some work on embedding string edit distances for LSH, it applies to measures of *global* alignment such as Hamming distance and cyclic shifts [6].

n-grams, to provide features for document similarity computations. We balance recall and precision by adjusting the size of the shingles and the proportion of the shingles that are used to calculate document similarity. After determining which features will be used to represent documents, we build an index of the repeated n-grams and then extract candidate document pairs from this index. Table 1 shows the parameters used in our approach.

3.1.1. N-gram Document Representations. The choice of document representation affects the precision and recall, as well as the time and space complexity of a text reuse detection system. In our experiments, we adopt the popular choice of n-grams, contiguous subsequences of n words. Very similar passages of sufficient length will share many long subsequences, but long n-grams are less robust to textual variation and OCR errors. As we see in the experiments below (§4), values of n between 5 and 7 provide a reasonable tradeoff between accuracy and efficiency for the newspaper corpus.

In previous work on detecting short quotes or “memes”, short contiguous n-grams have been proven to be quite effective. Suen et al. [7], for instance, use 4-grams to identify similar quotes in newspaper articles and blogs. Many fixed phrases, however, will also be detected by that approach. For instance, two documents are not much more likely to be derived from the same source if they both include the 5-grams “it is no accident that” and “in the final analysis it”. As describe below, we can mitigate this problem by putting an upper bound on the document frequency of the n-grams we use.

In addition to standard contiguous n-grams, therefore, we also explored the use of *skip n-grams*, or non-contiguous subsequences of the document. Skip n-grams allow us to look at subsequences of words that fall in more widely separate parts of the document. We parameterized our skip n-grams by n , the number of terms included; g , the minimum gap between terms; and w , the maximum number of terms covered by the skip n-gram. A contiguous 5-gram would thus be $n=5$ $g=1$ $w=5$. In this paper, we confine ourselves to skip bigrams ($n=2$) with width up to 105.

3.1.2. Downsampling Document Features. Since duplicate texts will share many matching n-grams, many duplicate detection systems use only a small fraction of the n-grams to represent each document.

One of the most popular downsampling techniques is $0 \bmod p$ [11]. When indexing a document, this technique hashes each n-gram to an integer and then

determines whether the hash value is divisible by some integer p . Since instances of the same n-gram in all documents hash to the same value, the algorithm does not need to maintain a dictionary of the downsampled features. This also permits easy parallelization since the hash functions can be computed independently in each batch. Many duplicate detection systems for the web use values of p of 50 or above, which drastically reduces index sizes [5]. This downsampling technique is less effective for local as opposed to global text reuse [9] and can also hurt recall for noisily OCR'd documents, as we see below.

For bigrams, we also adopt filtering on stopwords. While indexes of longer n-grams are not dominated by stopword-only entries, pairs of stopwords can increase the size of a bigram index significantly. Due to the prevalence of OCR errors in the newspaper dataset, we treat all words of four or fewer characters as stopwords. Also on the stoplist are the terms with the highest document frequency that make up 2/3 of the tokens in the corpus. On the newspaper collection, leaving out the four-character-or-less terms, this amounted to 1440 stopwords. We ignore any bigram with at least one stopword. This filter proved to be too stringent with longer n-grams, though some variants might be successful.

3.1.3. Efficient N-gram Indexing. The next step is to build for each n-gram feature an inverted index of the documents where it appears. As in other duplicate detection and text reuse applications, we are only interested in the n-grams shared by two or more documents. The index, therefore, does not need to contain entries for the n-grams that occur only once. We use the two-pass space-efficient algorithm described in Huston et al. [12], which, empirically, is very efficient on large collections. In a first pass, n-grams are hashed into a fixed number of bins. On the second pass, n-grams that hash to bins with one occupant can be discarded; other postings are passed through. Due to hash collisions, there may still be a small number of singleton n-grams that reach this stage. These singletons are filtered out as the index is written.

In building an index of n-grams, an index of (n-1)-grams can also provide a useful filter. No 5-gram, for example, can occur twice unless its constituent 4-grams occur at least twice. We do not use this optimization in our experiments; in practice, n-gram indexing is less expensive than the later steps.

3.1.4. Extracting and Ranking Candidate Pairs.

Once we have an inverted index of the documents that contain each (skip) n-gram, we use it to generate and

Table 1: Parameters for text reuse detection

n	n-gram order
w	maximum width of skip n-grams
g	minimum gap of skip n-grams
p	modulus for downsampling n-grams using $0 \bmod p$ [11]
u	upper bound on distinct series in the posting list

rank document pairs that are candidates for containing reprinted texts. Each entry, or *posting list*, in the index may be viewed as a set of pairs (d_i, p_i) that record the document identifier and position in that document of that n-gram.

Once we have a posting list of documents containing each distinct n-gram, we output all pairs of documents in each list. We suppress repeated n-grams that appear in different issues of the same newspaper. These repetitions often occur in editorial boilerplate or advertisements, which, while interesting, are outside the scope of this project. We also suppress n-grams that generate more than $\binom{u}{2}$ pairs, where u is a parameter.² These frequent n-grams are likely to be common fixed phrases. Filtering terms with high document frequency has led to significant speed increases with small loss in accuracy in other document similarity work [13]. We then sort the list of repeated n-grams by document pair, which allows us to assign a score to each pair based on the number of overlapping n-grams and the distinctiveness of those n-grams. When downsampling n-grams with $0 \bmod p$, we use the entire ranked list; otherwise, we confine our evaluation to document pairs with 5 or more shared n-grams.

3.2. Local Document Alignment

The initial pass returns a large ranked list of candidate document pairs, but it ignores the order of the n-grams as they occur in each document. We therefore employ local alignment techniques to find compact passages with the highest probability of matching. The goal of this alignment is to increase the precision of the detected document pairs while maintaining high recall. Due to the high rate of OCR errors, many n-grams in matching articles will contain slight differences.

Unlike some partial duplicate detection techniques based on global alignment [14], we cannot expect all or even most of the articles in two newspaper issues, or the text in two books with a shared quotation, to align. Rather, as in some work on biological subsequence alignment [15], we are looking for regions of high overlap embedded within sequences that are otherwise

2. The filter is parameterized this way because it is applied after removing document pairs in the same series.

unrelated. We therefore employ the Smith-Waterman dynamic programming algorithm with an affine gap penalty. We use the simplest form of a weight matrix common in the bioinformatics community: matching characters have a weight of 2, mismatches have -1, opening a gap is -5, and continuing a gap is -0.5. This use of model-based alignment also distinguishes this approach for other work, for detecting shorter quotations, that greedily expands areas of n-gram overlap [8, 10]. We do, however, prune the dynamic programming search by forcing the alignment to go through position pairs that contain a matching n-gram from the previous step, as long as the two n-grams are unique in their respective texts.

Even the exact Smith-Waterman algorithm, however, is an approximation to the problem we aim to solve. If, for instance, two separate articles from one newspaper issue were reprinted in another newspaper issue in the opposite order—or separated by a long span of unrelated matter—the local alignment algorithm would simply output the better-aligned article pair and ignore the other. Anecdotally, we only observed this phenomenon once in the newspaper collection, where two different parodies of the same poem were reprinted in the same issue. In any case, our approach can straightforwardly align two passages in the same document to different passages in two other documents.

3.3. Passage Clustering

We now have a set of aligned passage pairs. We sort the passage pairs in descending order by length and perform greedy single-link clustering. If two passages in the same document overlap by 80%, they are taken to be the same passage for purposes of clustering; otherwise, they are treated as different passages. Given this placement of passages into equivalence classes, we can view this step as detecting connected components in the graph of aligned passages. The V equivalence classes form the vertices; the edges are the E pairwise connections between passages determined by the previous alignment step. When clustering in a single pass through the pairwise data, using Tarjan’s disjoint set algorithm required $O(V)$ space and $O(V+E\cdot\alpha(E, V))$ time, where $\alpha(m, n)$ is the inverse Ackermann function [16]. This function grows so slowly that the algorithm is effectively linear in the input size. Table 2 shows an example cluster.

4. Evaluating Reprint Detection

For the pre-Civil War period, our corpus contains 1.6 billion words from 41,829 issues of 132 newspapers.

The collection was created with OCR of middling accuracy, as can be seen in table 2.

To evaluate the precision and recall of text reuse detection, we create a pseudo-relevant set of document pairs by pooling the results of several runs with different parameter settings. For each document pair found in the union of these runs, we observe the length, in matching characters, of the longest local alignment. (Using matching character length allows us to abstract somewhat from the precise cost matrix.) We can then observe how many aligned passages each method retrieves that are at least 50,000 character matches in length, at least 20,000 character matches in length, and so on. The candidate pairs are sorted by the number of overlapping n-grams; we measure the pseudo-recall at several length cutoffs. For each position in a ranked list of document pairs, we then measure the precision: what proportion of documents retrieved are in fact 50k, 20k, etc., in length? Since we wish to rank documents by the length of the aligned passages they contain, this is a reasonable metric. As a summary of these precision values, we employ the *average precision*: the mean of the precision at every rank position that contains an actually relevant document pair. One of the few earlier evaluations of local text reuse, by Seo and Croft [9], compared fingerprinting methods to a trigram baseline. Since their corpus contained short individual news articles, the extent of the reused passages was evaluated qualitatively rather than by alignment.

Table 3 shows the average precision of different parameter settings on the newspaper collection, ranked by the number of pairs each returns. If the pairwise document step returns a large number of pairs, we will have to perform a large number of more costly Smith-Waterman alignments. On this collection, a good tradeoff between space and speed is achieved by skip bigram features. In the best case, we look at bigrams where there is a gap of at least 95, and not more than 105, words between the first and second terms ($n=2$ $u=100$ $w=105$ $g=95$).

5. Exploratory Data Analysis

For a reader of today, one of the most striking aspects of 19c U.S. newspapers is the mixture of what we would recognize as news items with other genres such as poetry, chapters of novels, jokes, and morally edifying anecdotes of vague origin. After looking in detail at some of the more popular reprinted texts (§5.1), we look in aggregate at the geographic and network connections among newspapers that can be inferred from the reprinting data (§5.2). Finally (§5.3), we explore how different genres may be distinguished

Table 2: Example cluster from newspaper corpus. Only the beginnings of the text of the story, by popular temperance writer T. S. Arthur, are shown. In addition to the many obvious OCR errors, there are editorial changes, as well, such as omitting “Mr.” in the third text or changing “backwards” to “backward” in the second two texts. (These changes were checked against the page images.) Such changes provide evidence for how texts propagate.

1841-04-17 Sunbury American	soft answer ny t s arthur ill give him law to his hearts content the scoundrel said mr singleton walking backwards and forwards
1847-02-05 Anti-slavery Bugle	soft aasffch bv t s arthuit 1 4 ill give him law to his hearts coii ent fhe scoundrel said mr single on walking backwards and forwards
1847-05-04 Somerset Herald	soft answer by t s arthur ill ffive him inw to his hearts content the scoundrel said singleton walking bick ward and forward
1847-07-22 Vermont Watchman	soft answer ey t s arthur ill give him law to his hearts content the scoundrep said singleton walking backward and forward

Table 3: Average precision on detecting reprinted passages of different minimum lengths, measured in the number of *matching* characters. See Table 1 for explanations of the parameters of different runs.

Method	Pairs	Average precision on passages at least					
		50k	20k	10k	5k	2k	1k
n=10 u=100	3,342,162	0.18	0.09	0.10	0.10	0.24	0.31
n=7 u=50	3,579,992	0.22	0.18	0.20	0.20	0.22	0.30
n=2 u=50 w=55 g=45	4,297,764	0.32	0.30	0.28	0.27	0.21	0.27
n=6 u=50	4,433,792	0.21	0.20	0.22	0.21	0.21	0.31
n=7 u=100	5,971,269	0.20	0.11	0.12	0.14	0.30	0.44
n=5 u=50	6,443,100	0.21	0.22	0.25	0.22	0.20	0.31
n=2 u=100 w=105 g=95	7,512,442	0.34	0.30	0.25	0.29	0.40	0.47
n=2 u=100 w=55 g=45	9,756,985	0.31	0.26	0.21	0.24	0.32	0.45
n=2 u=100 w=25 g=15	12,798,056	0.28	0.24	0.19	0.21	0.26	0.39
n=5 u=100	15,258,185	0.19	0.15	0.15	0.17	0.30	0.47
n=4 u=50	17,954,922	0.19	0.23	0.27	0.23	0.18	0.28
n=5 u=100 p=10	43,701,236	0.19	0.15	0.15	0.16	0.26	0.35
n=4 u=100	71,263,521	0.16	0.17	0.17	0.18	0.26	0.42
n=5 u=100 p=5	77,771,798	0.18	0.15	0.15	0.17	0.28	0.41

not only by their content but also by the timing of their spread from one publication to another.

5.1. Reprinting and Textual Repurposing

Our initial list of the most reprinted articles is generically quite diverse, indicating that reprinting was not limited to particular genres of writing but reflected the range of materials included in most 19c newspapers. Top reprints from the corpus we have analyzed (from the period of 1830–1860) include the following:

- the March 4, 1857, inaugural address of President James Buchanan (over 150 reprints);
- Washington’s Farewell Address (53 reprints);
- a recipe for making starch with gum Arabic, to refresh and renew lawns, linens, and muslins (46);
- “A Beautiful Reflection” by Edward Bulwer-Lytton on the certain immortality of life (43);
- message of Queen Victoria to Pres. Buchanan conveyed over the Atlantic telegraph, upon the completion of the Atlantic cable in 1858, (43);
- December 1860 speech by President Buchanan proposing a (subsequently failed) amendment to

the constitution guaranteeing the right to own slaves (36);

- the text of the 1848 Treaty of Guadeloupe-Hidalgo ending the Mexican War and granting the US ownership of Texas and California (35);
- the “Soldier’s Bill,” passed by Congress in 1850, guaranteeing pension rights to soldiers and their widows (34);
- report on the defeat of the Austrian army in the Battle of Montebello during the Second Italian War for Independence in 1859 (33);
- a story titled “Household Economy” proclaiming the educational and moral value to children of reading newspapers in the home (32).

As this list indicates, reprints provide a vivid and engaging picture of the quotidian concerns of the American newspaper public in the nineteenth century as well a sharp historical lens onto the period. But given the diversity of genres and interests included on this list of heavily reprinted articles, what might account for the virality of heavily recirculated texts? One hypothesis is the following: quite a number of texts that are frequently reprinted have resonances in

multiple contexts—that is, these texts have “legs” that allow them to migrate into new publication venues. The second most reprinted article on our list, Washington’s Farewell Address, provides an instance of this phenomenon. Unlike Buchanan’s inaugural address, which clearly had news value in 1857, Washington’s address—initially delivered in 1796—did not. But whereas Buchanan’s address (or excerpts from it) was often reprinted as a stand-alone article, Washington’s address (or more accurately, significant excerpts from it) typically appears embedded within an article on another topic. Further, the excerpts from Washington’s address are most often those concerning the need for national unity—a topic with significant political and rhetorical currency in the period preceding the Civil War. However, the articles in which Washington’s address appears are themselves diverse: for instance, portions of the address appear in a report on the Republican convention denouncing the slave power of the south in the *Vermont Phoenix* (August 25, 1855) in a context one might expect; but also in a less anticipated context—namely a piece on the swearing in of Freemasons to the third degree that appeared in the *Nashville Union and American* (May 29, 1855).

A second example shows evidence of “legs,” as well, but of a different sort. In the case of the message from Queen Victoria to President Buchanan transmitted over the newly completed Atlantic Telegraph, the reprinting of the Queen’s words involved an unfolding news story that hinged upon the error of a hapless telegraph operator. Initially, only half of the Queen’s greeting was transmitted because the telegraph operator snipped the message in half. Thus, following an initial round of stories including half of the Queen’s words, a second round of articles appeared in which the Queen’s full message was included. And finally, a third round of articles appeared that repeated the Queen’s message following a 60,000-person march in Chicago celebrating the completion of the transatlantic telegraph. The words of both Queen Victoria and George Washington, then, were put to multiple uses and went viral in part because they had meaning in the context of diverse narratives of news and politics.

5.2. Geographic and Network Analysis

Looking beyond the details of individual texts, we can use our large and growing index of “viral” texts to begin modeling the systems that underlay nineteenth-century print culture. If we take reprinted texts as strong indicators of connection and influence among newspapers and among their editors, we can use our database of recirculated texts to map those connections,

both geographically and as a network. Literary scholars and historians have in the past been limited in their analyses of print culture by the constraints of physical archives and human capacity. A lone scholar cannot read, much less make sense of, millions of newspaper pages. With the aid of computational linguistics tools and digitized corpora, however, we are working toward a large-scale, systemic understanding of how texts were valued and transmitted during this period.

Using geographic information system software, our uncovered print histories can be correlated with other spatial data, such as historical county boundaries, census reports, and transportation data, to construct individual and comparative models of geographic dispersal. Using only the broadest data point from the census—raw population—we can compare the potential audiences for different “viral” texts. For instance, John Greenleaf Whittier’s agricultural hymn, “The Huskers,” was reprinted 30 times within our data set, and 634,499 people lived within 5 miles of its sites of reprinting. By contrast, the Scottish poet Charles MacKay’s religious poem, alternately titled “The Inquiry” and “The Resting Place,” was reprinted 30 times with 5 miles of 609,709 people. There are a range of other interesting data points in the census reports that we plan to explore more fully as the project develops. In 1840, for instance, census takers were interested in the print industry, and they recorded the number of daily, weekly, and monthly newspapers in each county, as well as the number of magazines and other periodicals. These data will help us understand how the print culture reconstructed through recirculation data lines up (or does not line up) with 19c Americans’ understanding of print culture.

Perhaps most compelling in our early spatial analyses has been the close correlation between the growth of reprinting networks and the growth of transportation technologies, particularly the railroad. As the rail network grows, connecting ever-smaller pockets of the country to the major hubs in the east, so to do the reprinted stories spread further and further outside major cities. Indeed, traffic seems to have moved in both directions, though unevenly. While the majority of the “viral” texts in our data set do seem to begin in major cities, such as New York, and spread from there to the south and west, a more significant minority of texts than we expected move in the opposite direction, beginning in smaller country newspapers and spreading from there into more metropolitan papers. In future, we plan to develop more robust techniques to reconstruct stemmas, or phylogenetic trees, of textual transmission.

Finally, visualizing our data as a network exposes potential connections between newspapers beyond ge-

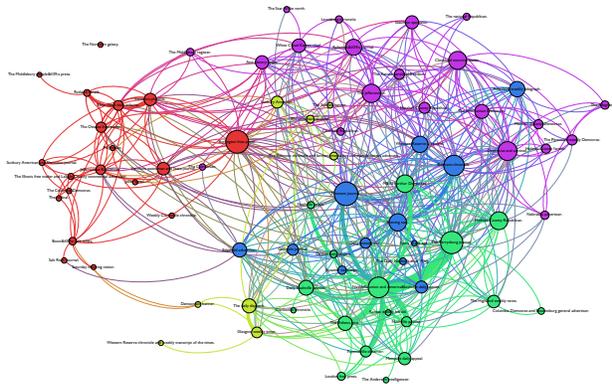


Figure 1: Visualization of the links among newspapers induced by shared reprinted texts. Colors indicate communities derived from graph modularity.

ography (figure 1). In these graphs, the nodes represent individual newspapers, while the edges are weighted based on how many texts two newspapers share. The communities of practice highlighted in these graphs often stretch across antebellum America. One community of reprinting partners, for instance, includes newspapers in Vermont, New York, Ohio, and Missouri, which was in the west of the United States during the period of our study. These connections suggest further research. When the network graphs suggested a close connection between the *Vermont Phoenix* (Brattleboro, Vermont) and the *Fremont Journal* (Fremont, Ohio), for instance, we investigated further, only to discover that the editors of these papers were brothers-in-law. These geographic and network modeling techniques, then, have pointed to new and relevant questions for humanities fields such as history and literary studies.

5.3. Reuse Dynamics: Reading Fast and Slow

The corpus of reprinted texts also reveals interesting temporal patterns. Some texts are reprinted frequently at first and then become less popular; others continue to attract interest from newspaper editors—and presumably readers—for years afterwards. Figure 2 shows the median lag time, in log number of days, between the first observed appearance and later reprints. The lag times are plotted by the year of a text’s first occurrence. (We omit plots for 1856–60 due to boundary effects.) Most obviously, there are simply more texts (and more newspapers) as time goes on. In many years, there are two peaks in the frequency of lag time, at around one week and over 3 years. In 1846–7, for example, there is an increase in short lag-time, “fast” texts, due in part to coverage of the Mexican-American War. In later years,

these fast texts become more frequent. We suspect that this is in part explained by the rise of wire services, though more research is needed.

Not surprisingly, different kinds of texts travel fast and slow. We divided reprinted texts into a fast set, with median lag times under one year, and a slow set, with median lag times over five years, which gave approximately equal sets. We trained a logistic regression model to predict whether a text would be fast or slow. Features were all words in a text’s first occurrence that (1) were five characters or more in length and (2) occurred five or more times in the corpus. The year of first occurrence was included as a categorical feature, to account for the different proportions in fast and slow texts over time. The training objective function was regularized with an L_1 (Lasso) penalty to achieve a sparse set of predictors [17].

Table 4 shows the top negative (“fast”) and positive (“slow”) coefficients. Highly correlated with fast textual propagation, for example, are terms related to the Mexican-American War, such as *Texas*, *Mexico*, and [Zachary] *Taylor*. Also interesting are *government* and *tariff*, *cases* and *corpse*. Slow texts focus on *love* and other affective terms, on *heaven* and interestingly *woman*. The year 1840 is also a useful predictor since fewer texts from that year were fast. With a random split between training and test, logistic regression achieved 73% accuracy on test. When earlier years (1840–1849) were used to predict a later one (1850), accuracy fell to 60%. Taking the cross product of year and word features only slightly diminished this overfitting effect, and more analysis of the rhetorical patterns of different genres should be helpful. Linear regression on the log median lag produced similar coefficients but fairly poor predictive accuracy.

6. Conclusions

We have described efficient algorithms that exploit shingling, space-efficient hashed indexing, and local alignment to detect clusters of passages reprinted within large collections of noisily OCR’d texts without a priori knowledge of which texts were repeated. Applied to a collection of 19c U.S. newspapers, these techniques uncover a corpus of popular but now often obscure texts of diverse genres, geographic distribution, and temporal dynamics. By analyzing patterns of reprinting, we hope to explore how ideas spread, which ideas spread, and which segments of the public shared ideas. These techniques thus stand to yield insight into the nature of the print public sphere in the nineteenth century at a level of specificity heretofore unavailable.

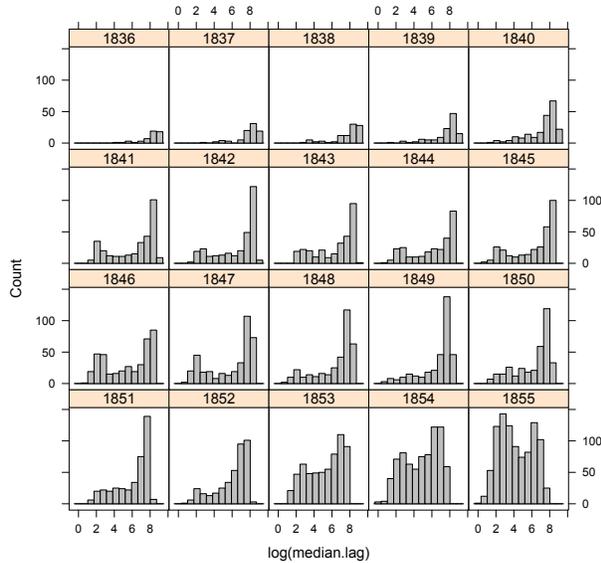


Figure 2: Median lag time, in log number of days, of reprint clusters from first appearance, by year of first appearance. Many years show peaks around 2 (\approx 1 week) and 7 (\approx 3 years).

Table 4: Top features from model predicting the lag in reprinting. “Fast” texts have a median reprinting time of a year or less; “slow” texts have a median reprinting time of at least five years.

“Fast” texts		“Slow” texts	
texas	-0.700	love	0.459
mexico	-0.692	young	0.154
pills	-0.672	earth	0.138
taylor	-0.649	1840	0.129
tariff	-0.534	awoke	0.121
government	-0.502	fine	0.098
board	-0.463	silence	0.097
effect	-0.428	benevolence	0.078
whig	-0.426	soul	0.057
mate	-0.418	sweet	0.048
prices	-0.418	grant	0.035
goods	-0.416	hang	0.033
corpse	-0.387	behold	0.026
cases	-0.383	bright	0.025
general	-0.370	woman	0.020
public	-0.368	things	0.020
cure	-0.367	heaven	0.019

Acknowledgment

This work was supported in part by NEH Digital Humanities Start-Up Grant HD-51728-13 and by Northeastern’s NULab for Texts, Maps, & Networks. Any views, findings, conclusions, or recommendations expressed do not necessarily reflect those of the NEH.

References

- [1] J. Leskovec, L. Backstrom, and J. Kleinberg, “Meme-tracking and the dynamics of the news cycle,” in *KDD*, 2009.
- [2] H. Ryu, M. Lease, and N. Woodward, “Finding and exploring memes in social media,” in *Hypertext*, 2012.
- [3] G. M. Namata, S. Kok, and L. Getoor, “Collective graph identification,” in *KDD*, 2011.
- [4] M. L. McGill, *American Literature and the Culture of Reprinting, 1834–1853*. U. Penn. Press, 2003.
- [5] M. Henzinger, “Finding near-duplicate web pages: A large-scale evaluation of algorithms,” in *SIGIR*, 2006.
- [6] A. Andoni, A. Goldberger, A. McGregor, and E. Porat, “Homomorphic fingerprints under misalignments: Sketching edit and shift distances,” in *STOC*, 2013.
- [7] C. Suen, S. Huang, C. Eksombatchai, R. Sosič, and J. Leskovec, “NIFTY: A system for large scale information flow tracking and clustering,” in *WWW*, 2013.
- [8] O. Kolak and B. N. Schilit, “Generating links by mining quotations,” in *Hypertext*, 2008.
- [9] J. Seo and W. B. Croft, “Local text reuse detection,” in *SIGIR*, 2008.
- [10] R. Horton, M. Olsen, and G. Roe, “Something borrowed: Sequence alignment and the identification of similar passages in large text collections,” *Digital Studies / Le champ numérique*, vol. 2, no. 1, 2010.
- [11] Y. Bernstein and J. Zobel, “A scalable system for identifying co-derivative documents,” in *SPIRE*, 2004.
- [12] S. Huston, A. Moffat, and W. B. Croft, “Efficient indexing of repeated n-grams,” in *WSDM*, 2011.
- [13] T. Elsayed, J. Lin, and D. W. Oard, “Pairwise document similarity in large collections with MapReduce,” in *ACL Short Papers*, 2008.
- [14] I. Z. Yalniz, E. F. Can, and R. Manmatha, “Partial duplicate detection for large book collections,” in *CIKM*, 2011.
- [15] D. Gusfield, *Algorithms on Strings, Trees, and Sequences*. Cambridge University Press, 1997.
- [16] R. E. Tarjan, “Efficiency of a good but not linear set union algorithm,” *J. ACM*, vol. 22, no. 2, pp. 215–225, 1975.
- [17] J. Friedman, T. Hastie, and R. Tibshirani, “Regularization paths for generalized linear models via coordinate descent,” *J. Statistical Software*, vol. 33, no. 1, 2008.