Notes on Practice

John Laudun and Jonathan Goodwin

# Computing Folklore Studies: Mapping over a Century of Scholarly Production through Topics

*Folklorists, like most practitioners in a field, understand the history of their discipline through a combination of their own reading and the consensus inherited from their graduate training and professional interactions. Disciplinary history, an effectively oral form of communication, codifies quickly. Highly contingent and random processes become widely understood as historically inevitable. In this preliminary report on a larger project examining the application of computational methodologies in the service of intellectual history, we explore the use of topic modeling as a way to understand the ebb and flow of topics and paradigms within a domain. Using JSTOR's Data for Research application programming interface to access the contents of 6,778 articles from three folklore studies journals* (Journal of American Folklore, Western Folklore, Journal of Folklore Research), *we used one form of topic modeling, the Latent Dirichlet Allocation, to delineate 50 distinct topics drawn from 125 years of research publication. Of particular interest here was the legendary "turn toward performance" in our field.*

**Keywords**
AFS ENTHNOGRAPHIC THESAURUS: Intellectual history, computer analysis, folklore and folklife, journals (periodicals)

WHEN THE *CENTENNIAL INDEX* of this journal arrived in mailboxes toward the end of 1988, it was the culmination of an enormous undertaking. Its five hundred pages represented a considerable amount of work performed by a team that numbered, according to the editor's account, almost twenty. With support from two foundations, the journal's then home university, and additional funding for printing from the Society, the index was a remarkable accomplishment, but most of all, the journal's

JOHN LAUDUN is Associate Professor of English at the University of Louisiana
JONATHAN GOODWIN is Assistant Professor at the University of Louisiana at Lafayette

*Journal of American Folklore* 126(502):455–475

editor at the time stressed in his introduction, it was the product of a collaboration between folklorists and computer programmers, requiring as it did a fairly robust database infrastructure.[1]

The final files for the index weighed in at what was then an astonishing fifteen megabytes, and, as Bruce Jackson admitted, the project required that the team "confiscate an IBM 3081 mainframe" (Jackson 1988:1). All the data entry and careful collation of the coding done by the different indexers resulted in a codex that begins with a "serial listing of the 9,655 articles, reviews, notes, obituaries, and announcements published in [the journal's] first century" (Jackson and Taft 1988:5) and concludes with three indices to those listings keyed first by author, then by subject, and finally by title. Held in the hand, a folklorist felt as though she held not only the key to unlocking the knowledge of the field but also could, simply by leafing through the pages, have a real sense of where the field had been and how it had gotten there. That is, it was both a map and a history.

For a time, the index was the ready reference, the gateway to finding the various forms of folklore scholarship that had passed through the pages of the journal, but the index's demise is anticipated in its last few pages: as the last number of volume 101, the number of the journal containing the index also carries the table of contents for the entire volume. At the very bottom of the first page of the table is an announcement with a straightforward title: "JAF Goes On-Line." That simple announcement was the beginning of the transformation of publishing by various information technologies made possible by advances in computing, brought about by the dynamic relationship enjoyed between hardware and software developers and users. Twenty-five years later, in the present moment of this essay, most of us are familiar with searching the journal through JSTOR and probably spend more time with its graphical user interface than we do with the pages of the *Centennial Index*. The same computational power that made it possible to build that codex was already well on its way to displacing the *Index* as a tool for searching. At the same time, we recognize that despite its fluidity, or perhaps because of it, JSTOR's interface in no way supports the kind of browsing that the *Index*'s codex makes easy. And yet the fixed nature of the index means it is limited by the terms to which the editors subscribed at the time it was printed. Those terms, as well as the *Index* itself in some fashion, are as much historical artifacts as the articles that they index.

Let us stay with the codex for a moment, imagining ourselves making our way down the aisle of the library, seeking out a particular volume of the journal. We drag our fingertips across the hills and valleys of the spines of the bound volumes until we find the one we want, pluck it from the shelf, and flip to the page we seek. We admit that this image might be a bit archaic for some. We are, many of us, all too used to being able to quickly narrow our view of the field, of any field really, through typing a few words, terms, or names into a box on a computer screen, and so we are rarely struck by the sheer volume of material that has accrued over the past 120 years. While we have become used to the precision with which we can pull materials from our discipline's historical record, we tend not to think of the record itself in aggregate. Here, standing in the library, we can again be amazed.

While we step forward to pull a book from the stacks, we must step back to see the larger whole of the many volumes of the journal. It is to the larger whole that

we seek to attend in our current work. If we step in to read a particular article, what happens when we step back? Is there a way to read from this distance? There is a growing sense in the humanities that we can, and should, do both, read closely and read from a distance. As the literary scholar Franco Moretti notes, "[d]istant reading . . . *is a condition of knowledge*: it allows you to focus on units that are much smaller or much larger than the text: devices, themes, tropes—or genres and systems" (Moretti 2000:57). Such a way of thinking is, of course, very familiar to folklorists, whose previous indexical efforts in the twentieth century produced an elaborate typological system in the many tale-type and motif indices.

In some ways, all we are suggesting here is that we turn the same kind of typological curiosity that produced the great indices of the previous era toward a reflexive moment, in an effort to make objective our own intellectual history. We would like to suggest that folklorists, like most practitioners in a field, understand the history of their discipline through a combination of their own reading and the consensus inherited from their graduate training and professional interactions. Disciplinary history, an effectively oral form of communication, codifies quickly. Highly contingent and random processes become widely understood as historically inevitable.

The first objective of our search through 125 years of folklore scholarship is simply to map the historical shifts in topics. We have taken as our object only articles, believing them best positioned to reveal the kind of ideational topography we imagine emerging from this work, and we have made use of three journals in order to make the best use of the kinds of computational methods we detail below: the *Journal of American Folklore* (JAF), the *Journal of Folklore Research* (JFR) (including its prior incarnation as the *Journal of the Folklore Institute*), and *Western Folklore* (WF) (including the *California Folklore Quarterly*).

While we discuss this in more detail in the next part of the essay, we would like to note upfront that the total number of articles being examined here, while quite large from a humanistic perspective at 6,778, the total word count of these research articles is within the range of the size of corpora used in many computer science topic-modeling analyses. Such a number guarantees that you need to find a way to seek out patterns in order to visualize trends perhaps not immediately discernible were one simply to start with the first page of the first issue of a journal and read into the present, simultaneously reading the other journals as they arose, supposing that such a thing were even possible, let alone desirable.

We recognize that we are not the first to count, or compute, in the service of providing folklore studies with a more robust intellectual history. Both Jill Terry Rudy and Simon Bronner have been instrumental in making possible the kind of exploration we embark upon here. Interestingly, from our point of view, both in their recent work have been interested in the *turn toward performance*. Though their reasons for considering the moment and their methods for consideration differ, both Rudy and Bronner deploy some form of computational analysis that itself draws upon, in turn, data taken from a larger data store: in the case of Rudy, it was citational information from the Social Studies Citation Index and the Arts and Humanities Citation Index, and in the case of Bronner, it was JSTOR. Rudy's study, published a decade ago in the pages of this journal, was a product of her fascination with *Verbal Art as Performance*

and how "the text in its production and reception both informs and diverges from the history of folklore studies" (Rudy 2002:6). Her work, like ours, was interested in the various ways personal and ideational networks intersect and shape each other.[2] Bronner's study sought to establish the turn toward performance as a moment in which folklore studies fully embraced the post-structural turn toward interpretation, a turn away from an earlier emphasis on explanation. In order to substantiate his study beyond the many close readings of texts he provides and to emphasize that those texts are simply indices of a larger trend within the folklore studies corpus, he searched the JSTOR database for uses of the two words. He reported:

> It is possible to quantify my contention that a post-structural turn was associated with use of "interpretation" after earlier use of explanation. Using the JSTOR database, I found that between 1888 and 1946, references in folkloristic journals to interpretation appeared only 284 times, whereas in a comparable period from 1947 to 2005, interpretation produced 1,094 hits, a ratio of almost four to one. In fact, reference to explanation predominated over interpretation in the earlier period, with 388 references, whereas in the later period there were 321 more references to interpretation than explanation. (Bronner 2006:415)

Both Rudy and Bronner ask their audience to consider at least one part of folklore studies' historical record as a body of work with its own patterns and relations to be examined *eo ipso*, as an entity (a network, a collection, a body) in and of itself, unbound by our conventional tying of texts to particular authors and their histories.

Taking our cue from Rudy and Bronner, we found ourselves fascinated by the prospect of considering the corpus of American folklore studies as an object of study in and of itself. What would, we wondered, the *Centennial Index* look like now, with another twenty-five years and two journals added to it? And how would we read its expanded form? We took as our starting point that probably the most-thumbed part of any such collection of listings is the subject index. It is, after all, where most of us begin, and it is also the part of any index we really seek to advance—we are all, we realize, happy to see our names in print, but what drives us, for the most part, and how we get there is by engaging others in a dialogue over a particular topic.

For that reason, we decided to focus our efforts on topics, though perhaps the particular version of topics that we will describe next will be unfamiliar to many. The topics we will explore are not those predetermined by subject matter expertise but rather those drawn from the texts themselves, as the words of the texts instantiate and influence the topics. We use a form of analysis here known as *Latent Dirichlet Allocation* (LDA), which we think fits well with folklore studies within the framework of performance studies itself, in which reality is socially constructed, often through language. Psychologist Thomas Landauer described the approach to language of a closely related method—latent semantic analysis—as follows:

> The premise is quite straightforward. Much (but not all) of the meaning of discourse must depend on the meaning of the words of which it is composed. Much (but not all) of the meaning of words must depend on experience with words in discourse. LSA provides a model of how a great deal (but not all) of the relatively stable, shared,

> central meanings of words derives from experience with words in discourse, and LSA provides a model of how a great deal (but not all) of the meaning of discourse derives from the meaning of the words of which it is composed. (Landauer 1999:308)

What topic modeling makes possible is the "explor[ation] of documents based on the themes that run through them" (Blei 2012:77). Our interest here is to see how the themes/subjects/topics of folklore studies are embodied within the particular collection of texts available to us (and in the form they are available—more on this in a moment). If we start with nothing more than the texts themselves and then read, or feed, those texts with a suite of algorithms designed to trace connections using only the words of the texts themselves, what familiar patterns would we be able to see and name? Just as importantly, would the exploration turn up any under-recognized—even latent—dynamics or trends worth further consideration?

Because one of us is a folklorist and the other is a literary scholar, we want to emphasize the exploratory nature of this project. Both of us are invested in computational (or quantitative) methods. Ours is a partnership born of a mutual interest in exploring how information and tools interact. We believe that topic modeling offers us one way to understand the nature of one form of disciplinary discourse. It is not the only way to understand it, and as we will make clear later in this essay, we are interested in other means. And, just as clearly, the scholarly article is not the only form of discourse within any discipline. There are far more kinds of transactions, as recent attempts to open the journal and to augment the annual meetings have made clear, but we were concerned, in this, our current exploration, to focus on what was readily available and to use the tools with which we felt most comfortable and which we felt offered the most promise.

We want to emphasize up front that we think that quantitative approaches to disciplinary history such as those outlined here offer scholars new ways to understand and question received disciplinary narratives. In the case of folklore studies, topic modeling reveals a consistent pattern of performance-related terms becoming increasingly important during the 1970s, an observation that agrees with general disciplinary history. This agreement signals the potential of the method to reveal similar disciplinary changes that are not as well known or to reveal shifts that have not yet been acknowledged. While these techniques rely on complex Bayesian statistics, their often-suggestive results are best combined with qualitative assessment and theoretical context.

Scholars often challenge the conventional understandings of a disciplinary change—an easy enough task given how complex the primary documents are compared to the disciplinary narrative. If forceful enough, this intervention can itself lead to a new disciplinary shift. We propose that the performative turn in folklore studies fits these criteria. And we want to re-evaluate this disciplinary shift—with computers.

## *Topic Modeling*

As useful as distant reading is as a metaphor for our current enterprise, our analogical library patron backing up in order to take in all the spines of the journals involved is, in fact, confronting the rather established conundrum of the map-territory relation.

The problem, as Gregory Bateson noted, is really not about the literal truthfulness of a map, but its structure being sufficiently analogous for the purposes at hand. Perhaps the best example of this is the schematic transit map of the London Underground, a map that sacrifices geographical accuracy in order to make the relative positions of stations along the lines, what stations connect which lines and which stations are within what fare zones, as easily understood as possible. In other words, it emphasizes topological relations over geographical ones. The result, as those who have traveled in London know, is that sometimes stations are spaced farther apart than they actually are, and thus one spends more time in tunnels than one would have by simply walking above. The reverse can also occur: the regular spacing of stations along a route can lead unknowing users to believe that the distances above are shorter than they appear. Like the convex rearview mirrors in which similar words appear, compression always means a distortion of some kind: the goal must be a compression whose distortion makes ready sense to the user within the context of use.

It was very important to us that, in order to effect the kind of compression we sought and to do so in a manner consistent across all texts under consideration, we find a method of doing so that focused on the actual language used, that we would begin with the concrete and move toward abstraction. After all, our goal was to see if the established abstractions like *the turn toward performance*, would be discernible within the data itself. Machine-learning researchers and others in computer science have developed techniques to infer meaning from large numbers of texts. The various forms of latent semantic analysis, as these approaches are generally known, share three fundamental assumptions in the way they approach texts: that texts have a latent semantic structure (simply called "topics"), that topics can be inferred from the texts by attending to what words occur with other words, and that words are related to topics and topics are related to texts. What separates these forms of analysis is their mathematical frameworks.

To keep this methodological background as brief as possible, we will move quickly to our topic modeling method of choice, one that has proved fairly popular in the humanities and thus has achieved an interesting kind of prominence—and, thus, a curiosity on our part to try it out on a fairly familiar corpus, so we could see for ourselves its potential utility in other areas of inquiry. The technique we chose to use in our exploration of folklore's intellectual history in the form of some seven thousand articles published in three journals over 120 years is known as the Latent Dirichlet Allocation. "Latent" refers, as we noted above, to the selection process of the algorithm; "Dirichlet" is the name of a German mathematician (from a Belgian family), whose research in numeric analysis led to important discoveries in probability theory.[3] David Blei, Andrew Y. Ng, and Michael I. Jordan first described the algorithm in 2002. Their paper "Latent Dirichlet Allocation" was published the following year (Blei, Ng, and Jordan 2003). LDA-related algorithms have been the subject of much research since. Blei has written a useful introduction to the process that explains in a nonspecialist manner the underlying statistical method. In it Blei defines LDA as "a statistical model of document collections that tries to capture this intuition. It is most easily described by its generative process, the imaginary random process by which the model assumes the documents arose" (Blei 2011:2–3). Topics in this model are

grouped distributions of words of the entire collection. Initially this process is randomized with the prior assumption of a Dirichlet distribution. The algorithm then progressively refines the topic distribution as it scans each document.

It is important to note that LDA is context-insensitive. It uses what is known as a "bag-of-words" approach. That is, only the frequency of words in each document is measured. Commonly occurring words and words with questionable semantic relevance to the study at hand are best filtered before running the algorithm. This process often requires several iterations to identify artifacts caused by Optical Character Recognition (OCR) errors and uninformative words not on the original stoplist. JSTOR's Data for Research makes word-frequency lists of articles in its database available. Implementations of LDA translate individual documents into matrices of word-frequency counts before processing them. Since the tool we used for many of these topic-models, MALLET, expects to convert documents into matrices itself, it is necessary to convert the word-frequency tables into an unordered document. This process does not affect how the models operate, however.

What is more important than the question of whether or not the algorithm imitates human comprehension is if the statistical technique approximates the topics that a human reader would identify in the documents. For instance, we would also like to note that the regular version of LDA makes no attempt at tracking changes in the terms that make up topics over time. Blei and J. D. Lafferty developed a variant algorithm known as "dynamic topic modeling" that seeks to account for changes in topics over time (Blei and Lafferty 2006). Changes in the distribution of topics over a chronological corpus can be tracked with the regular LDA algorithm, however. Corpora ranging from *Wikipedia* to *The New York Times* have been analyzed successfully with LDA topic models. Blei's group has experimented extensively with JSTOR's digitized corpus of the journal *Science*, showing changes in topics over time with the dynamic topic model and Web-based browser of topics in *The American Political Science Review*.[4]

As noted above, JSTOR's Data for Research interface does not, for obvious reasons, return complete texts of articles. Instead, each text is rendered as a series of two-place entries, listing each word used in the text and the number of times it was used. Since MALLET was built to receive whole texts and not word-frequency pairs, such as "the, 345," we had to find some way to make texts out of two columns of data.[5] We used some simple code that duplicated a word the number of times listed in the second column, generating a new file that replicated the word-frequency of the original, as MALLET expects. The next step was to model the data using MALLET. The software allows the user to specify certain variables, such as the number of topics that will be generated. This parameter was the most important for our purposes, as LDA cannot offer any assistance in choosing a reflective range of topics. Too few, and distinct topics will be merged. Too many, and coherent topics will split. So choosing the correct number is a heuristic process. We settled on fifty as a number that illustrated the range of discourse in our corpus effectively.

In order to test our hypothesis that LDA topic modeling could reveal the underlying discourse structure of folklore's performative shift, we downloaded a corpus from JSTOR's Data for Research consisting of every research article published in

the *Journal of American Folklore*, *Western Folklore* and its predecessor, *California Folklore Quarterly*, and the *Journal of the Folklore Institute* (plus its successor, the *Journal of Folklore Research*). Early experiments revealed that the British journal *Folklore* contained too many results that were oriented toward British disciplinary practice, which was not as strongly affected by the shift to performance studies in the United States. We should also note that there are other important folklore journals that we have not currently been able to access in digital form, and so our results are naturally skewed by what was available at the time of this work. The corpora that we have, however, is large enough, we believe, to be representative of the disciplinary shift we seek to model.

MALLET outputs a list of terms classed together as topics, along with their relative proportion in the corpus as a whole. The topic most directly correlating to our interest in performance looks like this in the output:

> performance social verbal narrative text audience context language speech structure discourse ritual performances analysis genre event play meaning events.

You can see that "performance" and "performances" are both included as separate words. The corpus can be lemmatized, or reduced to its dictionary root, with existing software tools,[6] and running the model on a lemmatized corpus produced a similar topic:

> cultural performance discourse practice text study form social work context culture genre power identity perspective body issue experience space.

It's important to note that *performance* shows up in other topics as well. It is present, for example, both in what is clearly the festival topic:

> dance festival mardi dancer gras play day dancing celebration costume year carnival group performance custom procession event christmas local.

As well as in a topic so closely allied that we wondered if future tweakings of our parameters for forming topics would not resolve the topics into one:

> social group community performance member event individual ritual behavior people role context relationship society traditional structure function audience time.

We will have more to say about the relationships between these topics as part of a larger set of topics that make up a performance cluster in the next part of this essay.

It is important to note at this point in our exploration of topic modeling as a useful tool for mapping the intellectual history of a domain that the word groups above are straight from the output that the software gives us and that MALLET lists these words in order: the left-most words are those most strongly associated with the topic that the entire word group represents. LDA was designed to supplement document retrieval technologies. If too much weight is placed on fitting an interpretation to the

headwords of a given topic, the tendency for humans to find patterns in what may be a random or unstructured connection can lead to overgeneralizations.[7]

Being aware of the limitations of LDA allows us to use it for exploration of folklore's disciplinary corpus without using its model as stand-alone evidence for any claims about disciplinary trends. The clustering of topics and documents produced by a model with no knowledge of folklore or language beyond a basic sense of how categorized variables tend to be represented in large numbers allows us to test our *existing* notions of how the field changes over time. The performative turn in folklore was part of a larger trend in the humanities that moved away from classification and description toward increasingly sophisticated interpretations. And we use a machine-learning classifying tool to examine the ways folklore has understood its interpretive turn. Topic-modeling is far from an objective representation, but its results are free of the self-reflexive regress that often characterizes disciplinary history when it becomes suspicious of explanatory narratives.

We use topic-modeling, in other words, as a form of generative search of an incomplete-but-representative selection of folklore's research output. Traditional search could show us the relative frequency of keywords; this generative search maps the discourse structures that link keywords to each other. Combining this map with our existing knowledge of the field allows us to test existing hypotheses about how it has changed.

## *Computing Folklore Studies*

As we surveyed fifty historical graphs of over one hundred years of folklore studies, we wondered what could be gleaned. We anticipated that we would see a few general patterns: topics that had declined in interest and topics that had gained in interest, as well as some constants. We were also keen to see if there were any particular peaks or dips or other forms of dynamism among the histories, and we were not disappointed. The five-year means of the fifty topics, perhaps the easiest to see in this regard, returned some interesting results. Figure 1 shows the five-year means of all fifty topics in the lemmatized corpus. While we predicted that topic modeling would reveal an increase at the time corresponding to the performative shift in folkloristics, the five-year means of the *cultural performance discourse* topic's occurrence from 1888–2012 in figure 2 shows a clear rise in the 1970s.

If we turn to the actual texts involved and list the twenty articles that MALLET has ranked as having the highest percentage of the topic, we see the results in table 1. To some degree, at this point in our exploration, we have at least confirmed that the topic model aligns with the discipline's own sense of the topic's center: the articles above are largely written after the turn toward performance. In addition, the listing suggests that metacritical considerations have the highest concentration of terms that co-occur, and we see that here with very esoteric considerations of dimensions of performance theory and its application in particular arenas ranking highly. Four texts are, in fact, from the same issue of *Western Folklore*.[8]

What this quantification of folklore studies also makes possible is the ability to explore topics that trend similarly to performance. For the sake of simplicity and

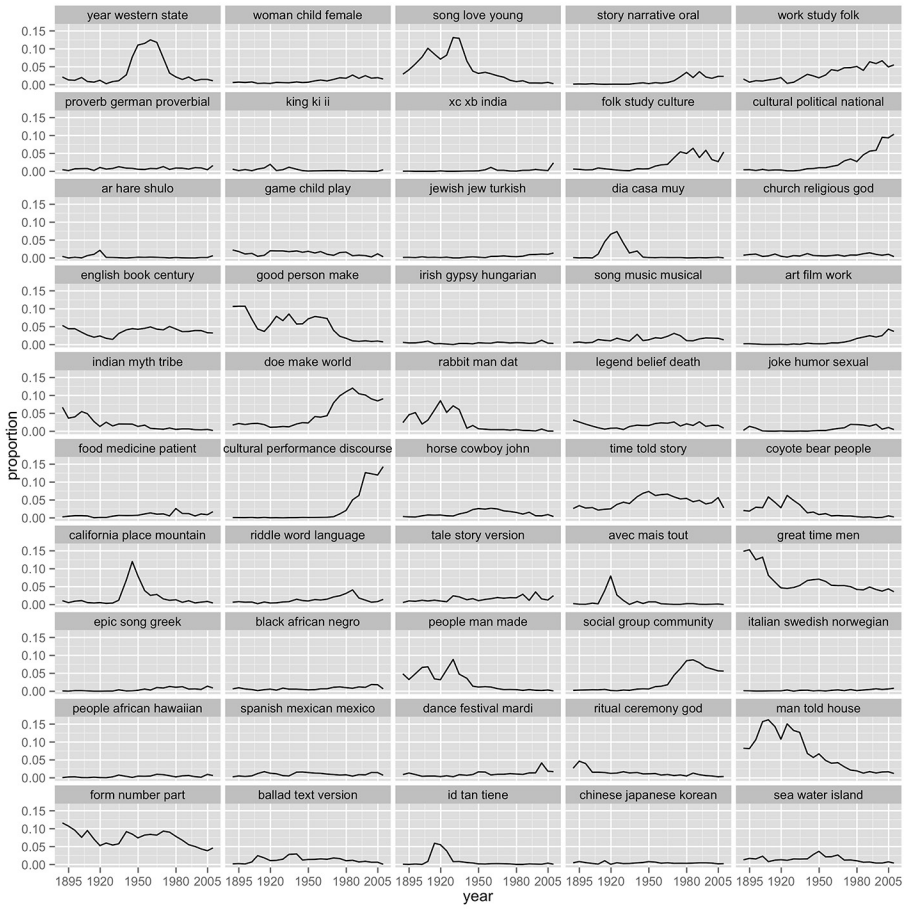


Figure 1: All Topics, 1888–2012.

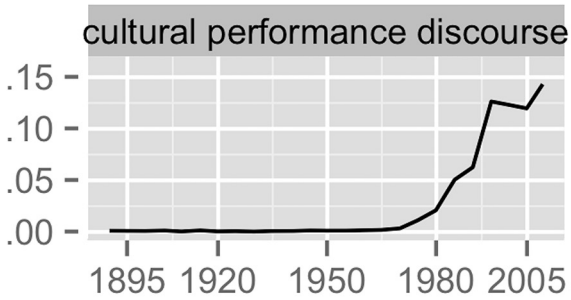


Figure 2: The *cultural performance discourse* topic, circa 1888–2012.

Table 1. The top twenty articles for *cultural performance discourse* listed in order of how prominently the topic figures in the article

| Author | Article | Journal | Year |
|---|---|---|---|
| Deborah A. Kapchan | "Hybridization and the Marketplace: Emerging Paradigms in Folkloristics" | WF | 1993 |
| Amy Shuman and Charles L. Briggs | "Introduction" | WF | 1993 |
| Katharine Young | "Whose Body? An Introduction to Bodylore" | JAF | 1994 |
| Richard Bauman | "Disciplinarity Reflexivity and Power in Verbal Art as Performance: A Response" | JAF | 2002 |
| Olivia Cadaval | "'Show Trial' or 'Truth and Reconciliation'? A Response" | JFR | 2000 |
| Deborah A. Kapchan and Pauline Turner Strong | "Theorizing the Hybrid" | JAF | 1999 |
| Margaret Mills | "Feminist Theory and the Study of Folklore: A Twenty-Year Trajectory toward Theory" | WF | 1993 |
| Charles L. Briggs | "Disciplining Folkloristics" | JFR | 2008 |
| Stephen Olbrys Gencarella | "Constituting Folklore: A Case for Critical Folklore Studies" | JAF | 2009 |
| Donald L. Brenneis | "Some Contributions of Folklore to Social Theory: Aesthetics and Politics in a Translocal World" | WF | 1993 |
| Deborah A. Kapchan | "Performance" | JAF | 1995 |
| Gerald E. Warshaver | "On Postmodern Folklore" | WF | 1991 |
| Charles L. Briggs | "Metadiscursive Practices and Scholarly Authority in Folkloristics" | JAF | 1993 |
| Charles L. Briggs | "Rethinking the Public: Folklorists and the Contestation of Public Cultures" | JFR | 1999 |
| Richard Bauman | "The Philology of the Vernacular" | JFR | 2008 |
| Jay Mechling | "On Sharing Folklore and American Identity in a Multicultural Society" | WF | 1993 |
| Harris M. Berger Giovanna P. Del Negro | "Bauman's Verbal Art and the Social Organization of Attention: The Role of Reflexivity in the Aesthetics of Performance" | JAF | 2002 |
| Amy Shuman | "Dismantling Local Culture" | WF | 1993 |
| Susan Ritchie | "Ventriloquist Folklore: Who Speaks for Representation?" | WF | 1993 |
| Paul W. Hanson | "Reconceiving the Shape of Culture: Folklore and Public Culture" | WF | 1993 |

because we wanted to keep our observations as objective as possible, we discerned eight topics as having "last quarter activity." (We will discuss the other trends in a moment.) Table 2 gives the complete list and orders the topics by the number of documents in which each topic is dominant.

The topic that we have labeled *performance* is third on the list and is flanked by topics in which the discipline's practitioners, *folklorists*, emerge as significant as well as, in the case of topic 4 (*work study folk*), the means by which disciplinary discourse proceeds: for example, *publication* and *program*. The first topic in the list above, topic 21 (*does make world*), confused us momentarily, but then, following our process of

Table 2. The eight topics exhibiting the same dynamics as *cultural performance discourse*

| Topic ID | Documents | Topic |
|---|---|---|
| 21 | 257 | does make world point human fact nature sense kind problem view work term life interpretation meaning thing question find |
| 04 | 252 | work study folk field student research folklorists collection material state society year published interest program publication folklorist article project |
| 26 | 169 | cultural performance discourse practice text study form social work context culture genre power identity perspective body issue experience space |
| 38 | 151 | social group community performance member event individual ritual behavior people role context relationship society traditional structure function audience time |
| 09 | 150 | cultural political national culture people social public history local ethnic tradition community state group folk united world life historical |
| 08 | 141 | folk study culture tradition theory cultural folklorists concept term research historical social traditional science approach belief method society analysis |
| 31 | 98 | riddle word language text speech line meaning form term english verbal answer formula linguistic phrase structure riddling oral speaker |
| 03 | 54 | story narrative oral experience narrator personal event text told life traditional audience telling tale storytelling storyteller tradition character listener |

beginning at a distance and then zooming in to understand the nature of a phenomenon, we checked the texts associated with it and recognized that the abstraction more familiar with readers would be *psychological approaches to folklore*. A short list of some of the authors that appear regularly in this list should be illustrative of our conclusion: David Hufford, Alan Dundes, Elliott Oring, Bruce Jackson, Roger Abrahams, Robert Georges, and Wolfgang Mieder.

Drawn into the graphs like this, we felt obliged to explore the peak and decline of interest in psychological approaches, only to realize that our initial topic of interest *performance* revealed a similar, if not decline, then at least plateau. Scanning all eight graphs we had coded for *last quarter activity*, we saw similar declines and plateaus in all but one topic, topic 9 (*cultural political national*). On the one hand, this collection of trends suggests that perhaps the psychological and the performative either found some resolution in the political or that focus on the political simply displaced previous concerns. Andrew Goldstone and Ted Underwood have described some of the interpretive dangers and opportunities here:

> As you change the number of topics (and other parameters), models provide different pictures of the same underlying collection. But this doesn't mean that topic modeling is an indeterminate process, unreliable as evidence. All of those pictures will be valid. They are taken (so to speak) at different distances, and with different levels of granularity. But they're all pictures of the same evidence and are by definition compatible. Different models may support different interpretations of the evidence, but not interpretations that absolutely conflict. (Goldstone and Underwood 2012)

The other topics clustered into five broad trends, which we came to describe, sticking with our simple, objective schema, as first half activity, early peak, middle peak, constant, and dynamic. The list of topics that populate each trend and the number of texts in each topic are in table 3.

Table 3. The rest of the topics, listed by the numerical identifier assigned to them by MALLET, the number of texts in which they are the dominant topic, the first three keywords in the topic, and the assigned topography

| Topic ID | Text Count | Keywords | Topography |
|---|---|---|---|
| 44 | 387 | man told house | first half activity |
| 34 | 352 | great time men | first half activity |
| 02 | 277 | song love young | first half activity |
| 20 | 127 | Indian myth tribe | first half activity |
| 22 | 123 | rabbit man dat | first half activity |
| 37 | 105 | people man made | first half activity |
| 29 | 50 | coyote bear people | first half activity |
| 47 | 49 | id tan tiene | early peak |
| 33 | 36 | avec mais tout | early peak |
| 10 | 28 | ar hare shulo | early peak |
| 13 | 28 | dia casa muy | early peak |
| 00 | 637 | year western state | mid-century peak |
| 30 | 280 | california place mountain | mid-century peak |
| 16 | 425 | good person make | dynamic |
| 28 | 317 | time told story | dynamic |
| 45 | 299 | form number part | dynamic |
| 15 | 209 | english book century | constant |
| 11 | 123 | game child play | constant |
| 18 | 120 | song music musical | constant |
| 27 | 111 | horse cowboy john | constant |
| 46 | 99 | ballad text version | constant |
| 32 | 95 | tale story version | constant |
| 49 | 95 | sea water island | constant |
| 42 | 94 | dance festival mardi | constant |
| 41 | 93 | spanish mexican mexico | constant |
| 43 | 84 | ritual ceremony god | constant |
| 19 | 75 | art film work | constant |
| 23 | 75 | legend belief death | constant |
| 35 | 75 | epic song greek | constant |
| 25 | 68 | food medicine patient | constant |
| 24 | 62 | joke humor secual | constant |
| 05 | 57 | proverb german proverbial | constant |
| 01 | 51 | woman child female | constant |
| 40 | 49 | people african hawaiian | constant |
| 48 | 49 | chinese japanese korean | constant |
| 12 | 46 | jewish jew turkish | constant |
| 14 | 46 | church religious god | constant |
| 36 | 42 | black african nefro | constant |
| 07 | 39 | xc xb India | constant |
| 17 | 26 | irish gypsy hungarian | constant |
| 39 | 23 | italian swedish norwegian | constant |
| 06 | 17 | king ki ii | constant |

As table 3 shows, there are seven topics with a great deal of activity in the first half of the period under examination. Because neither *Western Folklore* nor the *Journal of Folklore Research* are being published during this period, this trend and the one that follows are a product of JAF's own history. Even the briefest glances through the contents of these topics reveals that these topics represent the great collection projects of JAF as sketched out by William Wells Newell in its opening pages and as later developed by various members of the Society. We have, in the order listed above, folktales from around

the world (44), narratives and rituals found among natives of the Americas (34, 20, 22, 37, 29), American folk ballads (02), and African American folklore (22).

The next four trends are represented by far fewer topics. Having an early interest that typically declines as quickly as it arose, what we described above as an *early peak* are topics that treat Latin American folklore, including New Mexican traditions (47, 13); Francophone folklore, largely Canadian with some later work in Louisiana (33); and an interesting admixture of tale collections and considerations that span the old and new world (10). Interestingly, the peaks here are contemporaneous, with the rise and fall coming during the interwar years of the twentieth century, circa 1917 to 1942.

Peaking just after these topics, and in the middle of the period being mapped here, are two topics, *year western state* (00) and *california place mountain* (30) that can at least be partially understood as artifacts of the emergence of a major new journal in the field, *Western Folklore*.

Finally, of the topics showing significant dynamism during the study period, there are three whose behavior could not be readily captured in a short phrase. They are *good person make*, the lead topic in 425 articles of our corpus; *time told story* in 317 articles, and *form number part* in 299. A complete list of the word clusters associated with these topics (table 4) does nothing to reveal what they are. Just the opposite, they looked frighteningly similar to our eyes. But a look at the texts associated with each topic reveals that *good person make* addresses folk belief; that *time told story* is an admixture of jokes, legends, tall tales, and occupational folklore from contemporaneous historical settings; and that *form number part* encompasses collections of regional folklore, including place names, considerations of diffusion, some examinations of material culture forms, and treatments of myth.

The question remained, however, if there was any way to explain the particularly odd nature of their dynamic. In the case of folk belief, it seems to have reached its zenith in the early decades of the Society's history before succumbing to a precipitous fall during the First World War. Interest in folk belief seems to have remained quite low during the 1920s, picking up a bit in the thirties and falling again during the early years of the Second World War. Interest surged again in the fifties and sixties, with folk belief's popularity seeming to dim in the face of the turn toward performance.

The nexus of subjects that seems best encompassed by the label of *regional folklore* follows a somewhat similar pattern to *folk belief* above. Interest in the topic seems to

Table 4. Historical Trends in Outlier Topics, 1888–2012

| | |
|---|---|
| good person make water bad put luck day child sign house tree head cut cure white plant man eye |  |
| time told story back year people good thing heard day home man family friend make put big night house |  |
| form number part found time common case made present fact similar type general group material area point find study |  |

be greatest at the beginning of the Society's history. Regional interests seem to have declined in the twenties and remained low during the thirties, only emerging as a significant interest after the Second World War. In a trend that invites further exploration, we noted that interest in regional folklore remained strong and the topic does not participate in the same kind of decline as other topics as performance studies emerges. In fact, it increases slightly in proportion, peaking near the early nineties, a time during which, as we have observed above, performance studies became increasingly more reflexive in nature.

The dance festival topic's decline is not as steep and comes a bit later. Since it is not as strongly correlated with the discourse of performance, this result is what we expected. Articles representative of the "cultural performance" include "Hybridization and the Marketplace: Emerging Paradigms in Folkloristics" and "Whose Body? An Introduction to Bodylore" (Kapchan 1993; Young 1994). Those that are strongly represented by "social group community" include "Time Consciousness and Growing Up in Bamana Folk Drama" and "Tending Bar at Brown's: Occupational Role as Artistic Performance" (Brink 1982; Bell 1976).

Two topics that are Spanish language artifacts immediately reveal themselves with sharp spikes in the years of their publication. There have not been enough materials published in other languages for their topics to register independent of the inherent clustering of the foreign language terms in a mostly English corpus. The two most representative topics in the unlemmatized and lemmatized corpora, respectively, are "great men made time long day small end left large young good head order place held set hands make" and "form number part found time common case made present fact similar type general group material area point find study." The lemmatized corpus here is more specific, with more words indicative of the process of folkloric research. Words such as "made" and "end" are arguably worthy of being included in a stop-word list. Such inclusions are a matter of the art of progressive refinement. Another factor that can significantly affect the results is the minimum amount of times a word can occur in the corpus before it is considered in the topic model's probability matrix.

One intriguing detail from the looser non-lemmatized model was this topic: "folk study folklorists work studies research field society history published culture students folklorist material scholars materials collection literature dorson." That a topic so closely associated with folkloric practice would include the name of one of its most prominent practitioners is not surprising, but it is generally rare for even the most discussed scholar's names to appear in the topics created from journal articles. The only other author surname of a folklore scholar appearing in these topics was Dundes, unsurprisingly enough.[9] The lemmatized corpus does not contain any scholar's name.

Another artifact of the topic modeling process of some interest to readers of this journal is that publication places cluster together along with items relating to news accounts and other sources: "year western state news york angeles san california city item newspaper april reported day march daily letter july war." The journal *Western Folklore* contributes its name to this topic; we eliminated "american" and "folklore" from the corpus to avoid this kind of disproportionate representation. The word "western," however, was potentially meaningful enough in other contexts to warrant inclusion.

The degree to which a given article correlates with a topic representative of the disciplinary shift does not indicate how important or influential that article was in its formation. In fact, articles that signal an upcoming methodological shift will likely contain *fewer* keyword clusters later associated with it. It is not our goal in this paper to identify the sources of this shift, as existing scholarship has documented them well and will continue to re-evaluate them, if folklore follows the pattern of other disciplines.[10] We are concerned, rather, with finding signs of an impending shift and seeing if computational methods assign the same chronology to these ruptures as do disciplinary narratives.

## *Toward a Better Map of Folklore Studies*

Jorge Luis Borges infamously once chronicled an empire who sought such perfection in its representations of things that the "Cartographers Guilds struck a Map of the Empire whose size was that of the Empire, and which coincided point for point with it" (Borges 1972:131). In Borges's account, the denizens of the empire eventually toss the map aside, its tattered remnants becoming shelter for animals and beggars. The irony of Borges's fictional world is that, as Gregory Bateson notes, the "territory is *Ding an sich* and you can't do anything with it" (1972:454). A map is a representation, a compression of a larger entity that we wish both to conceive, and perhaps visualize, as a thing in itself and also with which we wish to interact, navigate. A map asks us to put our bodies where our ideas are, to commit to a representation and to trust that its paths will lead us where we wish to go.

The metaphor of the map is a useful one, if only because, in this instance, the mathematics lying behind the cartography is fairly complex. Here, the map is a topic model, and the cartography behind it is the implementation of the Latent Dirichlet Allocation found in MALLET. This particular technology proceeds through a probabilistic process that eventually determines a coherent cluster of words that regularly occur with each other. These clusters are our topics. Because the process is, purposefully, somewhat random in its beginning and moves toward a finer and finer determination as the algorithm iterates over a body of texts, it will by nature generate slightly different topics with each pass over the same material. Were this a more scientific exploration, it is quite possible that we would have run through the collection dozens, perhaps hundreds, of times in hopes of achieving a kind of statistical regression to a stable composition of the topics.

That noted, we feel confident, based on the outcomes described in the previous section, that we arrived at a reasonable sketch of folklore studies in its first century or so—a sketch with great potential to become a more robust map. Future explorations of the collection may very well draw us to modify the number of topics in an attempt to determine if there is an ideal number for this particular corpus, but any such ideal would only be temporary against an ever-expanding and ever-changing body of work. What is important, we believe, is that the current model captures the general topography of the discipline as gleaned through these particular texts. Another model might very well switch a word or two from one topic to another, but this amounts to worrying, in our minds, about whether or not a map has described a particular topographical feature as

a foothill or a mountain. As Bateson notes later in the chapter "Form, Substance, and Difference" in *Steps to an Ecology of Mind*: "Every effective difference denotes a demarcation, a line of classification, and all classification is hierarchical. In other words, differences are themselves to be differentiated and classified" (Bateson 1972:457). What is important, from the traveler's point of view, is the landscape has changed. Clean boundaries are the stuff of political necessity, not reality itself. We are confident that folklorists, of all disciplinary practitioners, will be comfortable with the bit of potential slippage that this particular methodology introduces.

This initial work is limited to an examination of historical trends. There is more work to be done. Having isolated not one but eight potential topics intertwined within the *turn toward performance*, we now face the challenge of trying to understand how these ideas are related. With these topics in hand, we can also chart the network of citations that are, perhaps, to be found within and across these texts. Such an approach would allow readers to see not only the network of ideas, the ideology or ideologies of folklore studies, but also the networks of authors. The whole, the culture and society that lie behind the discipline and the field, may very well give us the opportunity to understand ourselves and our work as a system. As a small society, ours will be but one system among many, but that limitation is also our opportunity.

Of the work to be done, we have only just begun to sort out, as Goldstone and Underwood observe, which of these topics are more like subject categories and which are simply clusters of words associated with particular theoretical approaches. We have used *clusters* in this essay in the intellectual history of folklores studies repeatedly, as a way to think about the branched nature of the topography we are trying to trace, as well as a way to begin to think about any potential taxonomies that might be present but not yet fully articulated. Topics are themselves simply clusters of words that co-occur in texts. Mapped out as either historical graphs or network visualizations, we can see that topics also build into clusters—the network studies term for this is *components*—which help us to understand the subdomains within a larger domain. Such an approach, we believe, not only offers a more subtle view of folklore studies but might allow us to understand how subdomains within adjacent domains actually approach or overlap each other, perhaps creating a bridge across which other ideas might be communicated.[11]

In an expanded version of this work, we anticipate not only a more detailed examination of the historical trends but also a better assessment of the way authors, ideas, texts, and domains of inquiry intersect and interact. A very common way to examine the relationships between texts, especially in those fields where co-authorship is a common practice, is to begin with a bimodal graph of authors and texts. (See fig. 3.) Using such a graph as a starting point, it is possible to project two networks: the network of authors as a function of the texts that they co-author, and a network of texts as a function of the authors that they have in common. The same kind of bimodal network applies to citations, and similar kinds of projections are commonly created.[12] Goldstone and Underwood, among others, pursue a very similar scheme in their work with PMLA materials, creating a bimodal graph in which "topics are linked to each other if they tend to appear in the same articles" (Goldstone and Underwood 2012).

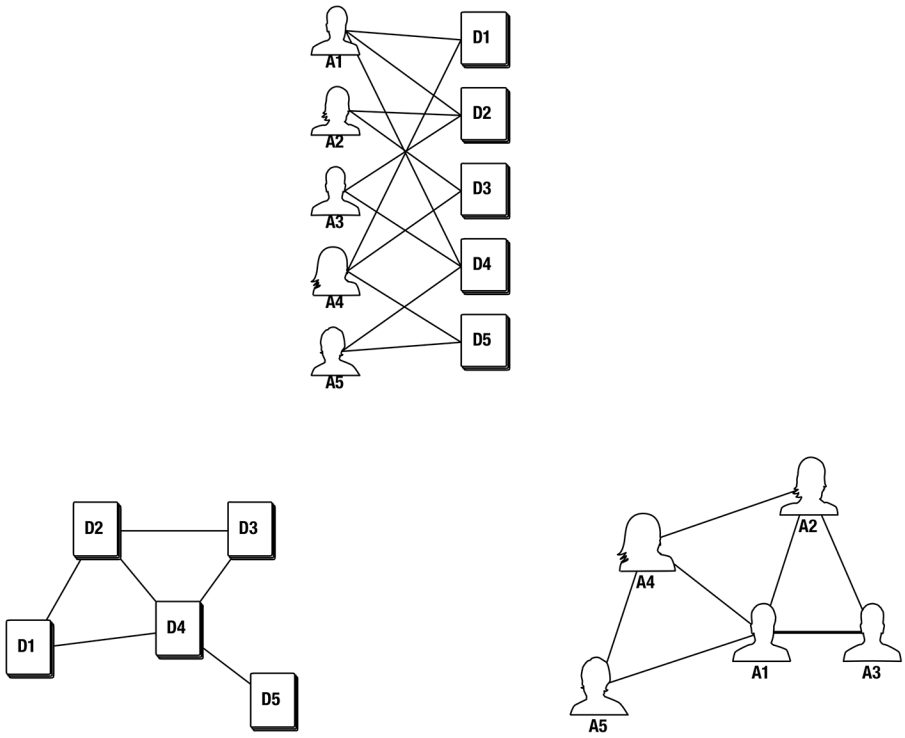


Figure 3: An example bipartite graph of authors and texts with projections of author and text networks with edges created and weighted by the number of collaborations.

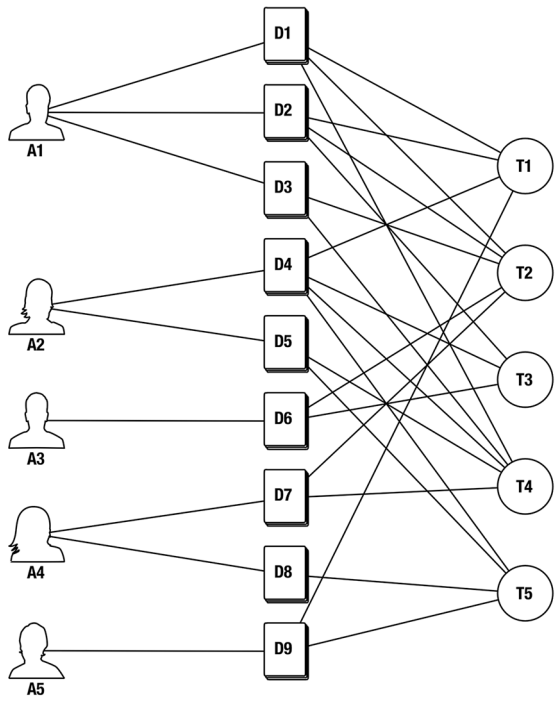


Figure 4: Authors, Texts, Topics.

Even from these few examples here, it is easy to see that there are a number of possibilities for spanning from one bipartite graph to another such that one could explore the relationship between authors and particular topics. (See fig. 4.) One could just as easily imagine doing something similar for the respective journals involved, or for comparing citation networks over and against topical networks. What all of these relationships, and their projections into networks, makes possible is an examination of the possible convergences and divergences, or slippage, between one set of relationships and another. Our belief is that the convergences are likely to confirm conventional understandings of folklore studies as a domain and as a field, and that the slippages will perhaps raise interesting questions that will require a closer look at the intellectual history of the discipline and its societies. Do, for example, slippages lead simply to the leaking of disciplinary focus or intensity or do they lead to innovation at the fringes of the domain that will make their way back to the center, and how quickly does that movement, as a move from the center to the margin or from the margin to the center, occur?

Obviously, this new way of mapping the intellectual topographies and histories of a domain raises a great number of possibilities for correlation. Which of those correlations will reveal new ways of understanding the work we do as individual practitioners or in groups is yet to be revealed. As we noted above, following Goldstone and Underwood, topic modeling is just as good at revealing what is being written about as it is at revealing how something was written, which might open an avenue for those interested in stylistics to examine modes of scholarly discourse. While an approach like topic modeling begins by working on a very large scale, it enables new ways to regard things up close. Far from depersonalizing scholarship, topic modeling as a form of distant reading enables the kind of close reading that makes it possible for us to see texts and their authors much more personally. The role of the individual in tradition has always been a great concern for folklore studies. At long last there seems to be methods of intellectual historiography that make it possible to see ourselves in the mirror the same way we see the world through the windows of our practice.

## *Acknowledgments*

John Laudun wishes to thank Tim Tangherlini for organizing the Seminar on Networks and Networking in the Humanities in August 2010, which was sponsored by the National Endowment for the Humanities and the Institute for Pure and Applied Mathematics at UCLA, where this project was first conceived. He also wishes to acknowledge the important contribution of Kyle Felker in an earlier phase of the project (Laudun and Felker 2011). Jonathan Goodwin wishes to thank Jennifer Guiliano, Travis Brown, and the staff at the Maryland Institute for Technology in the Humanities for organizing the NEH-supported Topic Modeling in the Humanities workshop in November 2012, where he learned much about the subject from the presenters and other attendees.

## *Notes*

1. The editor, Bruce Jackson, noted that the index was made possible by support from the L. J. and Mary C. Skaggs Foundation, the John W. and Clara C. Higgins Foundation, and the Faculty of Arts and Letters of the State University of New York at Buffalo, and even that "the AFS Executive Board increased the *Journal*'s ordinary printing and postage allocation for this issue" (Jackson 1988:1).

2. Indeed, there has been considerable exploration using the DIFI framework of *domain, individual, field* to establish a computational model of creative culture. Such efforts follow closely on the work of folklorists, who have long focused on the fact that, as Henry Glassie once observed, "[c]ulture is made up of ideas; society of people." For an early overview of work in this area, see Csikszentmihalyi, Feldman, and Gardner (1994).

3. Blei noted that the name of the previous version of the algorithm was the more prosaic "Model Three."

4. The *Science* browser is available here: http://topics.cs.princeton.edu/ Science/; the *American Political Science Review* here: http://topics.cs. princeton.edu/polisci-review/.

5. The most extensive tool for LDA is MALLET (http://mallet.cs.umass.edu/), though other implementations of LDA are available.

6. We used the WordNet interface in the Python Natural Language Toolkit library (NLTK) (http://nltk.org) to lemmatize the corpus. The NLTK is a robust collection of Python libraries for both basic and advanced linguistic computation of texts and corpora.

7. Ben Schmidt shows this is an interesting way with data from whaling voyages, though, as Schmidt notes, LDA is not a text-restricted model (Schmidt 2012).

8. As a matter of note, Jill Terry Rudy's work cited earlier in this essay is only four spaces down from being included in this list of top documents for this topic, confirming that metadiscursive reflections contain the greater density of a topic's keywords.

9. "MacPherson" and "Yeats" appear in a topic devoted to Celtic folklore: "irish ireland gaelic celtic macpherson ki ossian scottish tut john sean fi ba welsh ii mun yeats scotland fe," but we think these are not scholars but rather the subjects of the research in many cases. First names appear scattered in several topics, but "Richard" is too ambiguous to determine if it refers to a given scholar. Since surnames occur far more frequently in citations, it is unlikely that any first name refers to a particular scholar.

10. See, for instance, Barbara Kirshenblatt-Gimblett's "Folklore's Crisis" (1998).

11. Some of this work is not yet possible while a great deal of anthropological journal data lies behind firewalls and is not accessible in ways required by the approach we model here. It is certainly the case that we will be able to understand folklore studies as more of its journals become available through such interfaces as JSTOR's Data for Research program.

12. Earlier work in this project focused on creating a citation network using data accessed during the brief period that JSTOR made citation data available. Preliminary assessment of the material was that the data was extremely complex and required a great deal of structuring to be effective (Laudun and Felker 2011).

## *References Cited*

Bateson, Gregory. 1972. *Steps to an Ecology of Mind*. New York: Ballatine Books.

Bell, Michael J. 1976. Tending Bar at Brown's: Occupational Role as Artistic Performance. *Western Folklore* 35(2):93–107.

Blei, David M. 2011. Introduction to Probabilistic Topic Models. *Communications of the ACM*:1–16. http://www.cs.princeton.edu/~blei/papers/Blei2011.pdf.

———. 2012. Probabilistic Topic Models. *Communications of the ACM* 55(4):77–84.

Blei, David M., and John D. Lafferty. 2006. Dynamic Topic Models. In *Proceedings of the 23rd International Conference on Machine Learning* 138:113–20. Pittsburgh, PA, June 25–29.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3:993–1022.

Borges, Jorge Luis. 1972. *A Universal History of Infamy*, trans. Norman Thomas de Giovanni. New York: Penguin.

Brink, James T. 1982. Time Consciousness and Growing Up in Bamana Folk Drama. *Journal of American Folklore* 95(378):415–34.

Bronner, Simon J. 2006. Folk Logic: Interpretation and Explanation in Folkloristics. *Western Folklore* 65(4):401–34.

Csikszentmihalyi, Mihaly, David Henry Feldman, and Howard Gardner. 1994. *Changing the World: A Framework for the Study of Creativity*. Westwood, CT: Praeger.

Goldstone, Andrew, and Ted Underwood. 2012. What Can Topic Models of PMLA Teach Us about the History of Literary Scholarship? *Stone and the Shell*, December 14. http://tedunderwood.com/2012/12/14/what-can-topic-models-of-pmla-teach-us-about-the-history-of-literary-scholarship/.

Jackson, Bruce. 1988. Foreword. *Journal of American Folklore* 101(402):1–3.

Jackson, Bruce, and Michael Taft. 1988. A Century of Folklore. *Journal of American Folklore* 101(402):5–19.

Kapchan, Deborah A. 1993. Hybridization and the Marketplace: Emerging Paradigms in Folkloristics. *Western Folklore* 52(2–4):303–26.

Kirshenblatt-Gimblett, Barbara. 1998. Folklore's Crisis. *Journal of American Folklore* 111(441):281–327.

Landauer, Thomas. 1999. Latent Semantic Analysis. *Discourse Processes* 27(3):303–10.

Laudun, John, and Kyle Felker. 2011. Visualizing a Paradigm Shift through a Discipline's Citation Network. Paper presented at the annual meeting of the American Folklore Society, Bloomington, IN, October 12–15.

Moretti, Franco. 2000. Conjectures on World Literature. *New Left Review* 1:54–68.

Rudy, Jill Terry. 2002. Toward an Assessment of Verbal Art as Performance: A Cross-Disciplinary Citation Study with Rhetorical Analysis. *Journal of American Folklore* 115(455):5–27.

Schmidt, Ben. 2012. When You Have a MALLET, Everything Looks Like a Nail. *Sapping Attention*, November 2. http://sappingattention.blogspot.com/2012/11/when-you-have-mallet-everything-looks.html.

Young, Katharine. 1994. Whose Body? An Introduction to Bodylore. *Journal of American Folklore* 107(423):3–8.