[1]: title slide

Computing Folklore Studies

Into the Matrix

[2]: Still from Three Days of the Condor

In the opening scenes of 1975's *Three Days of the Condor*, a woman wanders through a room full of machines making noise and blinking lights. As we follow her, she stops to inspect the process, and we glimpse a red line of light sweep down the pages of an open book before a series of mechanical arms turn the page. The analyst shifts her attention from the book to the output of a telex machine which is, we are led to believe, simultaneously scanning and translating the text. As the film unfolds, we learn that the machines and humans work together to seek out patterns and plots that might reveal hidden operations by various clandestine organizations. (For those who don't know the film, the plot twist in the film is that it is the CIA itself whose plot its own unit has uncovered.)

What would it mean for a machine to read? Could it, for example, read a bunch of folklore scholarship and discern patterns and plots? How would it do that? Would it be as simple as an algorithm?

[3]: Levi-Strauss' formula

Trends in Folklore Studies

What if I told you that I had?

My colleague Jonathan Goodwin and I "fed" a computer program 125 years of folklore scholarship and asked it to make sense of it.

And the results looked like this:

[4]: 50 Topics

What you are seeing are fifty keyword clusters drawn from seven thousand articles from, principally, the *Journal of American Folklore* but also *Western Folklore* and the *Journal of Folklore Research* as determined by an algorithm which I will describe a little later. Before getting into any of the technical details, however, I thought it would be more interesting to explore with you some of the patterns that turned up and what plots they revealed.

We can discuss why fifty keyword clusters, and not twenty-five nor one hundred, in a moment, but for now, let's attend to the trend lines that accompany the keyword clusters and see if anything interesting has turned up. What the trend lines represent is the relative frequency that a keyword cluster occurred within a given year, or, in the case of the illustration in front of, over a five year period. (We have also smoothed the trend lines to make the graphs easier to read.)

As we surveyed the fifty historical graphs of one hundred years of folklore studies, Goodwin and I wondered what would we see. All of us here are familiar with general patterns that reveal a topic that has declined in interest or a topic that has gained in interest, and I am equally confident that all of us here would anticipate some constants. But which topics decline and which gain? Are there topics that gain and then decline, revealing a peak, or topics that begin with a great deal of interest, decline and then pick up again to reveal a dip. And when do these peaks and dips occur and what are their possible explanations? And are there any topics whose historical trends are not so confinable to such general patterns?

[5]: 7 topics among 50

Let me begin by isolating for you the topics which were the test subject of our investigation: those topics associated with performance studies. Goodwin and I had wondered not only if folklore studies' "turn towards performance" would show up, but when and how.

From a strictly descriptive point of view, we originally labeled this set of topics as "last quarter rise." We did this with all the graphs, trying to sort our results as best we could without knowing anything more than the trend lines themselves. (This seems a bizarre way of working in the humanities, where we lead with our subject matter expertise, but bear with me for a moment longer.) The other topics clustered into five broad trends, which we came to describe, sticking with our simple, objective schema, as follows:

[6]: Early Peak Topics among the 50

There were four topics whose trend lines revealed an early interest that *peaked early* in the history of American folklore studies and then declined as quickly as it arose. (It should be noted that at this point, all of our data is coming only from the *Journal of American Folklore*, so anything happening outside the journal is not reflected in any of these graphs.)

Curious to know more about this trend, and the others that follow, we pulled up a list of the articles most associated with this particular keyword cluster and discovered that these topics included treatments of 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.

[7]: First Half Activity Topics among the 50

The next trend, in terms of an historical timeline, were those that displayed a great deal of activity in the first half of the century. (Again, it should be noted that neither *Western Folklore* nor the *Journal of Folklore Research* are being published during this period, this trend, and so this trend like the one preceding it is a product of JAF's own history.) Even the briefest glances through the contents of these 7 topics reveals that these topics represent the great collection projects of the Journal as sketched out by William Wells Newell in its opening pages and as later developed by various members of the Society.

[8]: Mid-Century Peak Topics among the 50

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*. We labelled these two keyword clusters simply *mid-century peak*.

[9]: Outlier Topics among the 50

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.

[10]: 3 Outlier Topic Trends with Complete Keyword Cluster

A complete list of the word clusters associated with these topics 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 towards 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 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 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.

[11]: List of Dynamic Topics

These five trends — *early peak*, *first half activity*, *mid-century peak*, *last quarter activity*, and *outlier* — account for not quite half the keyword clusters identified by the topic modeling software we used. The remaining twenty-six topics were relatively constant in terms of the frequency of their occurrence across the historical record. (We have not, at this time, investigated these clusters very closely.) Our simple label for the 26 topics whose trend lines reflected a relatively stable presence throughout the history of folklore studies was *constant*.

A Closer Look at Performance Studies

[12]: 8 Performance Topics - Graphs

If we return to the intellectual history question with which we began — can we discern patterns and plots within a corpus of scholarly publications for a particular domain using nothing more than an algorithm? — we felt sure, having glimpsed the kinds of trends already described that topic modeling would highlight any increase in folkloristic discourse having to do with the the performative shift in the field. The five-year means of the *cultural performance discourse* topic's occurrence from 1888-2012 does not disappoint: the rise in the 1970s of those keywords frequency — in relationship to each other it should be noted — could not be more clear:

[13]: Topic 26 "cultural performance discourse"

It's clearly the case that in terms of terminology, the keywords associated with performance studies literally appear out of nowhere and their usage increases at first slowly and then quite quickly. (If there is an argument to be made for the application of paradigm shifts in the humanities, then surely it can be made using this graph, and the data lying behind it, alone.)

But there are other keyword clusters with similar rises in occurrence in the last quarter century or so:

[]: 8 Performance Topics - List

The first topic in the list above, topic 21 (does make world), confused us momentarily, but then, following our process of 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 to folklorists would be psychological approaches to folklore. 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.

Topic 21 & 26 & 9 layered

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.

It should be noted that 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 evaluate them, if folklore follows the pattern of other disciplines. 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.

Towards a Better Map of Folklore Studies

Co-Author Map with Projections

Jorge Luis Borges infamously once chronicled an empire who sought such perfection in their 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' account, the denizens of the empire eventually toss the map aside, its tattered remnants becoming shelter for animals and beggars. The irony of Borges' fictional world is that, as Gregory Bateson notes, is that the "territory is *Ding an sich* and you can't do anything with it" (454). A map is a representation, a compression of a larger entity which we wish both to conceive, and perhaps visualize, as a thing in itself and also with we wish to interact, navigate. A map asks us to put our

bodies where are 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 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 moving toward a finer and finer determination as the algorithm iterates over the body of texts, it will by nature generate slightly different topics with each pass over the same material.

That noted, we feel confident, based on the outcomes described above, 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.

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 towards 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 folklore studies a system. As a small society, it will be but one system among many, but that limitation is also an opportunity. (We have, for example, access to abstracts for the past ten years of the annual meeting. Etc.)

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:

slide

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 which they coauthor and a network of texts as a function of the authors which they have in common. The same kind of bimodal network applies to citations, and similar kinds of projections are commonly created.

Tripartite Graph

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 Figure) 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 de-personalizing scholarship, topic modeling as a form of distant reading makes possible 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.