Mapping Research Topics
at Multiple Levels of Detail

Sara Lafia[*1], Werner Kuhn[1], and Kelly Caylor[2]

[1]*Department of Geography and Center for Spatial Studies, UC Santa Barbara*
[2]*Department of Geography and Earth Research Institute, UC Santa Barbara*

Abstract: Institutional reviews typically rely on scientometrics, like the h-index and impact factors of their participants, to assess research productivity. Productivity is not the only review criterion however, and scientometrics can be difficult to generate and compare in multidisciplinary settings. "Distant reading" methods from the Digital Humanities can complement the current quantitative evaluation paradigm; these methods support qualitative narratives, comprehension, and discovery of knowledge by arranging vast bodies of text into graphs, maps, and trees. To test this idea, we apply distant reading methods to a multidisciplinary body of research authored by 240 researchers from the Earth Research Institute (ERI) at UC Santa Barbara over the past decade. We model cross-disciplinary topics of research publications and projects emerging at multiple levels of detail. From these, we design maps that reveal the latent thematic structure of multidisciplinary research. ERI's researchers use and evaluate these maps of research topics in the context of an institutional review to "read" ERI's body of research at a distance, i.e. at multiple levels of detail. We find that our approach strengthens the institutional review process by exposing thematic expertise, relationships between researchers, topical distributions and clusters of work, and the evolution of these aspects over time.

1   Introduction and Motivation

Universities and funding agencies request that organized, multidisciplinary research units summarize and report on their research, collaboration, and growth as part of periodic institutional reviews. Such reviews conduct "meta-research" on grant proposals, publications, and other products to provide a high-level view of research impact and productivity (Ioannidis, 2018). Yet, it is unclear how best to summarize and present multidisciplinary bodies of research in ways that generate useful insights and support effective reviews.

Bibliometrics and scientometrics supporting the quantitative study of published documentation and academic disciplines (Bawden & Robinson, 2015) have

---

*Correspondence: slafia@ucsb.edu

become cornerstones of institutional research assessments. Research administrators and funding agencies often use metrics, like the Hirsch Index (h-index) and the journal impact factor (JIF), to assess the impact and performance of departments or individual researchers and to monitor their collaborators or competitors (Jappe, Pithan, & Heinze, 2018). Such metrics are trusted due in part to their perceived scientific legitimacy and because they offer indicators, which if appropriately selected and applied, can yield data to support performance monitoring and the selection of research priorities (Vinkler, 2010).

There are limitations to evaluative metrics, however. For one, metrics based on citation analysis are often not comparable across academic disciplines (Radicchi, Fortunato, & Castellano, 2008). Incentivized metrics like citation counts, impact factors, and amounts of funding have been found to be vulnerable to manipulation (Moher et al., 2018). Academics subject to performance evaluations have also criticized institutions for relying on such metrics at the expense of narratives (McGuire, 2019).

Given these concerns, there is a growing desire to develop principles for multidisciplinary research evaluation that more adequately capture impact and quality. One strategy has been to complement evaluative metrics with high-level characterizations and narratives (Polonioli, 2020). Another has been to develop knowledge domain visualizations, or "maps of science", that chart the structure of vast realms of scientific literature by showing the development of research areas, their interconnections, and evolution (Börner, Chen, & Boyack, 2003).

Mapping is indispensable in many monitoring and planning contexts; without maps of the physical territory, it would be challenging to plan and manage the development of cities, landscapes, and infrastructure. Cadastral maps, for example, document ownership and other rights to the land; they also inform and communicate numerous planning interventions including strategic land use decisions, economic investment, and mitigation measures (Andrienko et al., 2007).

In general, space and time are fundamental ordering relations for knowledge representation (Janowicz, 2010). The "spatial turn" observed in the social sciences and humanities has exploited the idea of spatial organization to facilitate cross-disciplinary exchange, allowing many lines of thought to converge (Arias & Warf, 2009). In cognitive science, it has been claimed that conceptual spaces in which nearby concepts are similar, underlie human thinking and learning (Gärdenfors, 2000). These powers of spatial representation underpin the idea of *spatialization*, which maps abstract domains to map spaces in which nearby elements are similar (Kuhn, 1996).

In this article, we examine the utility and benefits of spatialization for mapping the latent thematic structure of a multidisciplinary body of research. We further explore its potential to complement current quantitative approaches by aggregating research topics at multiple levels of thematic detail. To develop and test our ideas, we situate our research in the context of an organized research unit at UC Santa Barbara (UCSB): the Earth Research Institute (ERI)[1]. ERI's stated mission is to "support research and education in the sciences of

_____

[1] https://www.eri.ucsb.edu/

2

the solid, fluid, and living Earth". Core areas of research within the institute consist of Natural Hazards, Human Impacts, Earth System Science, and Earth Evolution. ERI supports faculty and researchers spanning 24 academic departments who are supported by 145 different funding agencies, which encompass the full breadth of Earth and Environmental Sciences. To date, ERI has taken an ad-hoc approach to characterizing their research based in part on anecdotal observations. For example, ERI expertise has recently broadened from earth science and crustal studies to include conservation and biodiversity topics.

To formally capture this kind of institutional knowledge about ERI's evolving research, we propose a systematic, reproducible, and data-driven approach for eliciting cross-cutting research topics. We analyze research documents (publications and funded projects) by 240 researchers affiliated with ERI between 2009 and 2019 to produce research maps at multiple levels of detail. We then evaluate whether they are useful to researchers within the institution.

In the remainder of this article, we first situate our work in relation to existing approaches for abstracting and mapping information. We then describe our method to produce maps of a body of research at selected levels of detail. Finally, we report how leading ERI researchers evaluate the potential for our maps to support an institutional review[2]. We find that our approach complements the review process by exposing and relating thematic expertise, highlighting relationships between projects or teams of authors, analyzing topical distributions and clusters of work, and tracking the evolution of all these aspects.

2   Background

Topic modeling and spatialization jointly support exploration and discovery at multiple levels of detail by performing dimensionality reduction. This is analogous to cartographic generalization for maps: computational and cognitive issues of complexity are addressed by deliberately reducing the level of detail in the representation (Skupin & Fabrikant, 2003). Topic modeling offers a way to identify research topics latent in articles and projects that are not bounded by traditional silos, like academic departments and their terminologies. We use topic modeling to choose levels of detail and identify major topics shared by research documents. Spatialization is a means of abstracting large amounts of data to give a visual impression of their underlying structure, in particular similarity between concepts. Spatialization allows us to visualize these topics in such a way that nearby topics are more similar than distant ones. These methods jointly and sequentially support our goal of summarizing large amounts of thematic data and connections taking place at multiple levels of detail.

---

[2]A slight delay in the actual institutional review (resulting entirely from COVID-19) precluded feedback from external reviewers in time for our research project.

2.1 Topic Modeling

Topic models are statistical machine learning techniques that can uncover hidden structures in collections of documents; topic models group documents in which similar terms co-occur, forming sets of semantically related documents (Boyd-Graber, Mimno, & Newman, 2014). Latent Dirichlet Allocation (LDA) is a ubiquitous type of topic model in which documents are described by multiple topics (Blei, Ng, & Jordan, 2003). Topic models have been applied to classify and summarize large collections of documents, as well as solve similarity judgement problems (Blei et al., 2003). Topics themselves can also be of interest; for example, the National Institutes of Health and National Science Foundation have developed topic-based search interfaces to explore trends in related research projects (Boyd-Graber et al., 2014).

2.2 Spatialization

Spatializations offer high-level views of contents through the familiar visual modality of maps (Wise et al., 1995). The first law of cognitive geography, or distance-similarity metaphor, references the first law of geography, which states that "everything is related to everything else, but nearby things are more related than distant things" (Tobler, 1970). The distance-similarity metaphor treats distance in abstract spaces as metaphorically equivalent to dissimilarity (Fabrikant, Ruocco, Middleton, Montello, & Jörgensen, 2002). Spatialization has been applied to organize large, thematically diverse, and multidimensional collections; for example, Figure 1 shows political survey data from Switzerland spatialized as a political landscape, resulting in a map of continuous ideological values with nearby cities sharing preferences (Hermann & Leuthold, 2003).

Figure 1: Spatialization of Swiss political ideology as a landscape (*Source*: Hermann and Leuthold, 2003; reprinted with author permission)

4

2.3 Organizing Spatializations by Core Concepts of Spatial Information

To further systematize spatializations, we apply the theory of core concepts of spatial information (Kuhn, 2012). The concepts provide a high-level vocabulary with which to ask and answer questions about phenomena in space and time (Table 1). They capture distinct ways of thinking about space and computing with spatial information; thus, they are applicable to geographic as well as other spaces. They provide a set of interchangeable lenses through which research data can be spatialized and viewed (Allen et al., 2016).

Concept	Definition	Intuition	Question	Examples
Location	a description of where things are by spatial relations	things are located relative to each other or in reference frames	where is it?	in the center of town; at a latitude and longitude
Field	an attribute with values everywhere in a region and at all times during a period	fields continuously map positions and times to attributes	what's the value at a given position and time?	today's air temperatures at 8am everywhere in the state
Object	an individual in space and time, with properties and relationships	objects have identity	what's the value of an object property at a given time?	a building with address and owner
Network	a set of objects with links between pairs	networks capture connectivity	what connects two objects at a given time?	a bus network in a city
Event	an occurrence at some time, involving participants (fields, objects, and/or networks)	events change participants	what's the value of an event property at a participant location?	a tornado, an epidemic, a house sale, a road closure

Table 1: Summary of core concepts of spatial information

We produce a *field* of continuous topic values from the texts of research documents with a topic value at each position. This can be thought of as a landscape or surface of topic values. Research documents conceptualized as *objects* can then be embedded in this topical landscape. *Networks* of research documents may be linked, for example by shared topics or authors. *Events* may be introduced to answer questions about what has happened, is happening, or may happen to researchers and departments, for example.

In our previous work, we found that mapping research documents as objects embedded in fields revealed their distribution and coverage, while linking them in a network revealed their topical connectivity and centrality in the collection, among other insights (Lafia, Last, & Kuhn, 2019). In this article, we concentrate on fields and objects, but further demonstrate the relevance of core concepts of spatial information by enabling the detection of events through a timeline.

2.4 Level of Detail

Previous studies have shown that levels of detail in spatialized displays, such as hierarchical regions, shape viewers' interpretation of the similarity of elements like news articles (Fabrikant et al., 2002). The problem of choosing adequate levels of detail for topic models and spatializations is unresolved.

The selection of an appropriate number of topics is a central challenge in topic modeling; selecting too few leads to overly broad topics while selecting too many leads to redundancy (Greene, O'Callaghan, & Cunningham, 2014). Best practices recommend a combination of human evaluation strategies and topic coherence measures (Boyd-Graber et al., 2014).

Spatialization relies on generalization methods for merging individual features into groups. This is analogous to cartographic generalization, which performs hierarchical clustering based on feature similarity and results in changing geometric representations and labels for the features at each level of detail (Skupin & Fabrikant, 2003). The strict membership requirements imposed by hierarchical clustering may not be the best way to represent relationships between levels of detail in abstract spaces, however. Alternative tree-like structures, like semi-lattices, could more adequately represent concept organization, for example as sets of partially overlapping information (Hirtle, 1995).

2.5 Distant Reading as an Abstraction Paradigm

"Distant reading" (Moretti, 2005) provides methods for deliberately abstracting and visualizing unstructured text; in order to analyze hundreds of novels for example, it is necessary to render fewer elements in order to offer a sharper sense of high-level themes and their interconnection (Underwood, 2016). Distant reading uses graphs, maps, and trees to spatially configure units, like genres and novels, and reveal latent structures in their source material (Moretti, 2005). These methods are part of a suite of "macroscopic research" devices, including knowledge domain visualizations (Börner et al., 2003) and spatializations, that enable the study of patterns at multiple levels of detail over time.

Graphs, maps, and trees offer a typology of spatial and temporal views that support the exploration of large bodies of literary texts. Graphs are quantitative diagrams that produce temporal signatures revealing patterns and cycles like the rise and fall of literary genres. Maps are spatial diagrams that order components to reveal structures in the space of the narrative universe. Trees are morphological diagrams that emphasize the passage of time, revealing divergence and drift of literary styles. We believe that these models can also guide abstraction over scientific texts, which in our case are large numbers of abstracts from publications and grant proposals. Our application appears to be less ambitious, but deals with the same type of raw data: natural language texts following some conventions of style.

While texts are the "real objects" of literature, they may not be the right "objects of knowledge for literary history" (Moretti, 2005). By analogy, in the institutional review of a research unit, available data about publications and grants are the real objects of science, but are not the right focus for understanding it. Currently, institutional reviews focus on productivity metrics, which do not address this problem. Abstraction is necessary to support an understanding of broader patterns and trends latent in research documents, such as the emergence and evolution of shared topics.

3   Methods

We develop methods for producing maps that support the distant reading of research activities at distinct levels of detail. These maps show research topics and their evolution over time. To produce the maps, we first retrieve research documents and model the topics of their titles and abstracts. We then feed the output of topic models at two distinct levels of detail into spatialization algorithms to produce maps of the research topics. The data and code for our experiments with topic modeling and spatialization, which inform our approach, are publicly available in our project's Github repository[3].

3.1   Data Sources

We analyze publications and funded projects from ERI's 240 researchers active from 2009 - 2019. These data from the past decade cover research activities and financial aspects including: publications issued by the unit, funded project proposals, grants, and contracts. We gather publication metadata with the Dimensions API, which is freely available for non-commercial use. ERI maintains records of active researchers and their funded projects. We only consider publications and funded projects with titles and abstracts, leaving 3,770 research documents for analysis (3,108 publication abstracts with titles and 662 funded project abstracts with titles).

---

[3]https://github.com/saralafia/ERI-maps

3.2   Text Pre-processing

We prepare the research document titles and abstracts for topic modeling by removing records with identical identifiers (DOIs), removing HTML tags, and reformatting ASCII extended characters. Figure 2 shows that the research documents are relatively concise with an average length of 1,678 characters; they range in length from 128 to 7,083 characters.
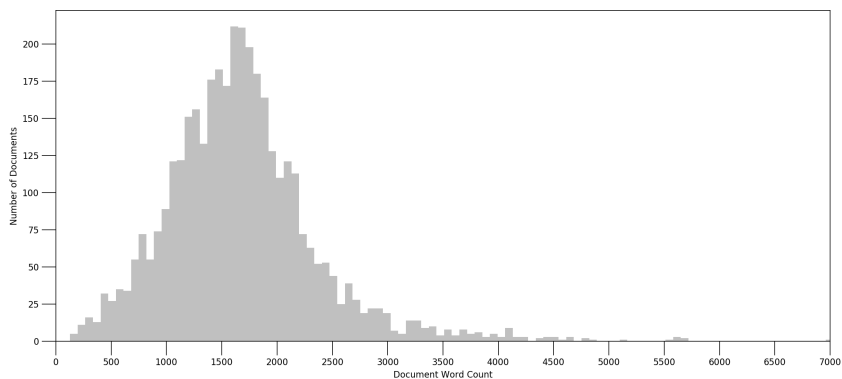

Figure 2: Distribution of research documents by length (word count)

Next, we reformat the research documents following a standard natural language processing pipeline (Boyd-Graber et al., 2014). We first determine distinct document terms using term frequency–inverse document frequency, or tf-idf (Salton, Wong, & Yang, 1975). This reflects the relative importance of a term to a document in a corpus and is often used as a weighting factor in information retrieval applications; this measure seeks to balance specific terms that show up frequently in relatively few documents (e.g. "polymerase") with those that show up frequently across many documents (e.g. "sample"). Frequent terms, scaled by how often they appear in the documents, are shown in Figure 3. We find that many of these frequent terms describe research methods (e.g. "estimate") rather than subject matter (e.g. "snow").

We remove common articles from the documents, as well as frequent and generic terms identified with tf-idf that qualify as stopwords. These terms are 'data', 'study', 'project', 'research', 'collaborative', 'include', 'result', 'increase', 'high', 'low', 'large', 'include', 'based'. We then tokenize the terms, convert them to lowercase forms, and construct n-gram models (bigrams, trigrams) to preserve contiguous sequences of words (e.g. 'climate change'). Finally, we lemmatize terms to resolve them to their base forms (e.g. 'changing' to 'change'). These steps produce a dictionary of word identifiers and word frequencies along with a document corpus to use in topic modeling.

Figure 3: Word cloud of frequent terms from the processed research documents

3.3 Topic Modeling

We consider two main kinds of topic modeling approaches: probabilistic (Latent Dirichlet allocation, hierarchical LDA) and matrix factorization (non-negative matrix factorization). In probabilistic approaches, each document is treated as a mixture of a small number of topics; words and documents get a probability score for each topic. The LDA algorithm is a generative probabilistic model that represents documents as mixtures of topics composed of words with certain probabilities (Blei et al., 2003). It assumes that similar words occur in similar contexts and aims to discover latent topics underlying the documents. Given our interest in modeling topics at multiple levels of detail, we also consider extensions of LDA proposed for learning topic hierarchies like hierarchical LDA, or hLDA. It estimates the structure of a topic hierarchy and partitions documents nonparametrically (Griffiths, Jordan, Tenenbaum, & Blei, 2004).

In matrix factorization approaches, dimensionality reduction methods from linear algebra are used to decompose a document-term matrix into a smaller set of matrices, which can be interpreted as a topic model (Lee & Seung, 1999). Non-negative matrix factorization (NMF) is a dimensionality reduction technique for decomposing samples, which in the case of topic modeling are documents. Like in LDA and hLDA, documents are represented as term vectors, which can be combined into a document term matrix. However, documents are represented as combinations of topics rather than likelihoods. In NMF, term weighting can also be used to boost distintinctive terms by term frequency inverse document frequency, or tf-idf (Salton et al., 1975). LDA has been positioned as an alternative to tf-idf, offering improved insights "into inter- or intra-document statistical structure" (Blei et al., 2003).

To determine a suitable number of topics to model, we consider the cognitive

9

implications of Miller's Law, which proposes that the average person can hold approximately 7±2 "chunks" of information in working memory (e.g. 7 digits, 6 letters, 5 words), limiting the simultaneous perception and processing of information by humans (Miller, 1956). Miller's Law, applied to our topic models, suggests a coarse level of detail (7±2 topics) that reviewers should be able to consider at once. For a suitable number of topics at a more detailed level, we reapply Miller's Law to each chunk of the coarse level, resulting in bounds of (5 x 5) and (9 x 9), or a range of 25 - 81 chunks.

We use these heuristics to guide the selection of models in combination with coherence as an interpretability measure. Coherence is a measure of the extent to which top terms representing a topic are semantically related, relative to how they relate to other terms in the corpus (Mimno, Wallach, Talley, Leenders, & McCallum, 2011). Coherence is considered to be more human interpretable for evaluating topic model quality than other measures, including perplexity and log-likelihood (Greene et al., 2014). It is based on the fundamental idea in classification that the members of a class should be more similar to each other than to members of other classes. We generate topic models with high coherence scores using each technique, which allows us to select and compare topic models at different levels of detail (Mimno et al., 2011).

Our topic modeling experiments are detailed in reproducible Jupyter notebooks in our project repository[4]. We first produce LDA models, selecting those with high coherence scores. We find topics related to methods like "modeling", "analysis", and "task" that are not indicative of ERI's research subjects. This may be an artifact of higher frequencies of methods-related terms like "modeling", which were prominent in the word cloud (Figure 3). In LDA, such frequent terms are not weighted or normalized as they are in other topic modeling approaches like NMF.

Next, we explore hierarchical LDA, or hLDA. We find that the nonparametric partitioning of the corpus produces broad, generic topics at each level. hLDA also overemphasizes research outliers by distinguishing them into separate and comparatively small categories. It seemed to be a promising technique, given its ability to model topics at multiple levels of detail within a single model, but we decided against using it due to its unsatisfactory interpretation of ERI's broad field of multidisciplinary research. hLDA may prove to be useful for modeling topics in a more constrained field of research, however.

Several factors suggest that non-negative matrix factorization (NMF) is a more suitable topic modeling approach for our purposes than LDA or its variants. We find that NMF produces topic models with higher coherence scores than our LDA models by about 17 percent on average. This may be because NMF is better suited to model smaller or sparser datasets, like titles and abstracts, rather than full text (Arora et al., 2013). We also find that NMF produces topics that are more indicative of subject matter, rather than methods. This may be due to term weighting with tf-idf, unlike LDA, which operates on raw term frequency (Greene et al., 2014).

---

[4]https://github.com/saralafia/ERI-maps

To find coherent topic models, we generate NMF models across a range of topic numbers (2 - 100) and calculate their coherence scores. Figure 4 shows a rapid increase in model coherence between 2 - 10 topics, a more gradual increase in coherence between 20 - 40 topics, and a continual, but modest increase in coherence up to 100 topics. While the model with 100 topics has the highest coherence score, it does not obey our usability constraints; its large number of topics exceeds the number of chunks of information that reviewers would be able to keep in mind. Instead, we select two distinct levels of topic models: a coarse level with 9 topics (falling in the 5-9 topic range) and a more detailed level of 36 topics (falling in the 25-81 topic range), which are optimally coherent.
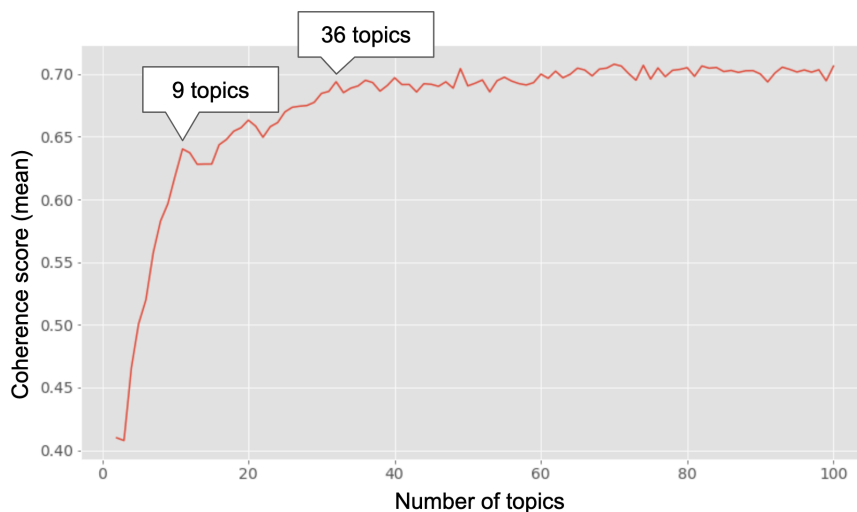
Figure 4: Coherence scores for NMF topic models with 2 - 100 topics

We generate the NMF models using Scikit-learn decomposition[5] as follows. Non-negative matrix factorization decomposes a document-term matrix into a smaller set of matrices, which can be interpreted as a topic model (Lee and Seung, 1999). We normalize the document-term matrix with tf-idf and decompose it into a product of two matrices: a topic-term matrix and a document-topic matrix. Figure 5[6] shows how terms and topics can be combined to approximately reconstruct documents. The topic-term matrix shows the topics associated with particular terms; the highest ranked terms are used to obtain a descriptor for each topic based on the 10 top terms for each. The document-topic matrix describes each document by a mixture of topics and shows how strongly it is associated with a topic or set of topics. The document-topic matrix forms the input to the subsequent spatialization while the topic-term matrix is used to reference topics and terms from their descriptors.

---

[5]https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.NMF.html
[6]Adapted from Wikimedia (CC BY-SA) https://upload.wikimedia.org/wikipedia/commons/7/7c/NMF.svg
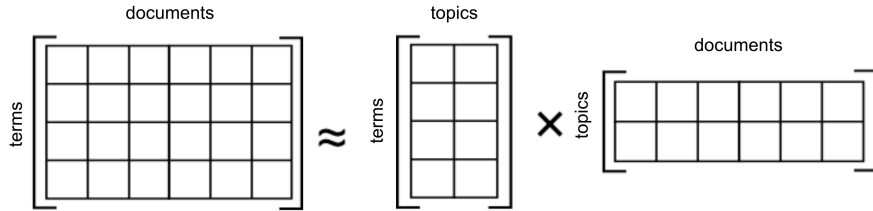
Figure 5: Conceptual overview of NMF decomposition of the document-term matrix

3.4 Spatialization

In our approach, the topic models serve as inputs to spatialization. Latent topics and their proximity can be represented in a number of ways, making the view of space underlying a spatialization a key choice (Kuhn, 2012). Space can be seen 1) as a field, or a continuous surface, with a topic value at each position; 2) as a collection of objects, which are located in an abstract space; or as 3) a network, which links pairs of objects to show a relation of interest like connectivity, hierarchy, degrees of separation, or flow.

To present the high dimensional topic mixtures of research documents in two-dimensional spatial configurations, we consider graphs, maps, and trees as distant reading methods (Moretti, 2005). The methods to produce Moretti's graphs, such as histograms and signatures, are well-understood but we decide not to pursue them because we do not have enough longitudinal data to support significant insights from graphs; our research documents span only 10 years. We develop a non-temporal variation on Moretti's trees as part of our evaluation (Section 5.3) to study the relationship between topics at two levels of detail; instead of showing topic evolution, we use these variations on trees to show topic clustering. Our spatializations produce maps, which handle time and events in the way maps traditionally do through a sequence of annual snapshots.

To produce maps, we conceptualize research documents as objects located in a continuous two-dimensional topic space. The topic mixture of each research document assigns it to a neighborhood of related research topics. To embed the research documents in a continuous topic space, we considered two dimensionality reduction techniques: t-distributed stochastic neighbor embedding, or t-SNE (Maaten & Hinton, 2008) and uniform manifold approximation and projection for dimension reduction, or UMAP (McInnes, Healy, & Melville, 2018). Both t-SNE and UMAP are suitable for our purposes, as they model high-dimensional objects as points in a low dimensional space, clustering similar objects and spacing apart dissimilar ones.

We considered each of these techniques, evaluating how well they convey the similarity of research documents; these experiments are detailed in repro-

ducible Jupyter notebooks in our project repository[7]. The inputs to the spatializations are the document-topic matrices resulting from the coarse (9) and detailed (36) NMF topic models. We first map research documents with t-SNE using manifold learning in Scikit-learn[8]. The t-SNE algorithm transforms the high dimensional document-topic matrix into a low dimensional coordinate representation. Each document is assigned a position based on its topic mixture, resulting in the placement of topically similar documents near each other and dissimilar documents farther apart. The UMAP process for assigning locations to research documents is nearly identical to that of t-SNE; a key difference is the assumption that documents are uniformly distributed on a complex surface, resulting in a distinct spatial configuration. We produced these with UMAP learn[9]. The axes in both t-SNE and UMAP are left unlabeled, as they describe complex curved paths in the original high-dimensional space and do not have human-interpretable meaning (Maaten & Hinton, 2008; McInnes et al., 2018).

We interactively explored the maps to interpret the effects of the map parameters, which balance local, pairwise similarity with global, inter-cluster similarity (Wattenberg, Viégas, & Johnson, 2016). The first parameter influencing the size, distance, and shape of clusters is perplexity, which controls the number of nearest neighbors. Perplexity describes how well a probability distribution predicts a sample. In our maps, low perplexity values produce clearly delineated clusters while high ones allow for more global connectivity and less clearly delineated clusters. We select a perplexity of 7; typical values fall between 5 - 50 (Maaten & Hinton, 2008). The second parameter is early exaggeration, which determines the compactness of clusters. This is an optimization method to create empty space in the map between clusters so they can achieve better global, in addition to local, organization (Maaten & Hinton, 2008). We select an early exaggeration of 5; the default value in the Scikit-learn implementation is 12.

Both mapping approaches appear to produce similar results in terms of the spatial layout of documents and clusters with a few notable differences. Figure 6 shows that t-SNE produces local clusters of similar objects that are visually distinct, while UMAP allows for more outliers and preserves compact clusters; for instance, all red documents clustered and labeled with "fault (seismic motion)" are concentrated in UMAP, while they are split into three distinct clusters in t-SNE. The effects of uniform spacing are also visible in UMAP; the red and blue clusters are disjoint in UMAP but are partial neighbors in t-SNE. The arrangement of individual documents and clusters of documents in t-SNE conveys topical similarity well. Based on these observations, we deem t-SNE to be a more compelling technique for our purposes.

---

[7]https://github.com/saralafia/ERI-maps
[8] https://scikit-learn.org/stable/modules/generated/sklearn.manifold.TSNE.html
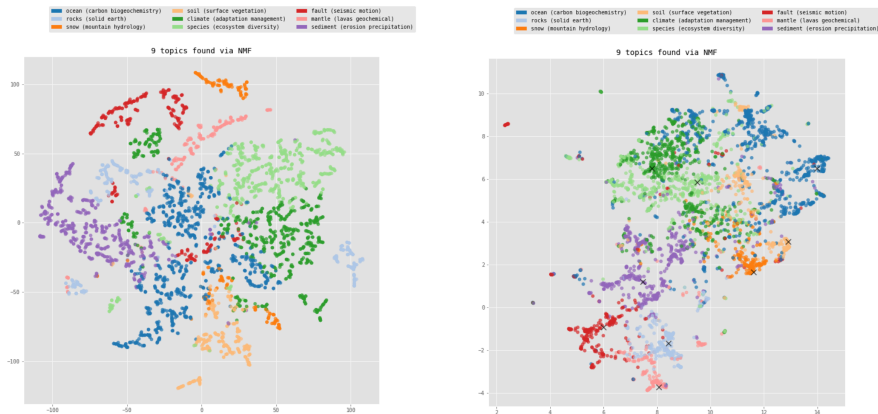[9]https://umap-learn.readthedocs.io/en/latest/

Figure 6: ERI research documents clustered by 9 topics with t-SNE (left) and UMAP (right); each color corresponds to the document's main topic, labeled with three term descriptors

4 Results

Our methods address the question of how to systematically elicit and represent the major topics of a complex, multidisciplinary body of research at multiple levels of detail that show their similarities and evolution over time. They result in maps of research documents located in a continuous topic space, which exhibit topical proximity in neighborhoods and capture multiple levels of detail over periods of time. We explore whether and how these maps of research topics support the qualitative assessment of a multidisciplinary body of research.

4.1 Reading Maps of Research Documents

The maps produced with t-SNE show research documents with similar topics forming neighborhoods at two distinct levels of detail. Documents are assigned to topic clusters, which are labeled with the first three terms from their topic descriptor. Topic modeling does not produce labels for the resulting topics, so this is a pragmatic choice that allows us to reference and interpret the topic clusters. The categorical colormap[10] offers perceptually distinct categories for visualizing the relatively large number of topics in the detailed topic model.

In the coarse map with 9 topics shown in Figure 7, we observe patterns related to the centrality, size, contiguity, and proximity of clusters. Documents assigned to the large "ocean" cluster are in the center of the map while smaller clusters like "snow", are on the periphery. This suggests that the documents described by the "ocean" topic are similar to more documents in the corpus than those assigned to the "snow" cluster, which may be more niche.

---

[10]https://colorcet.holoviz.org/
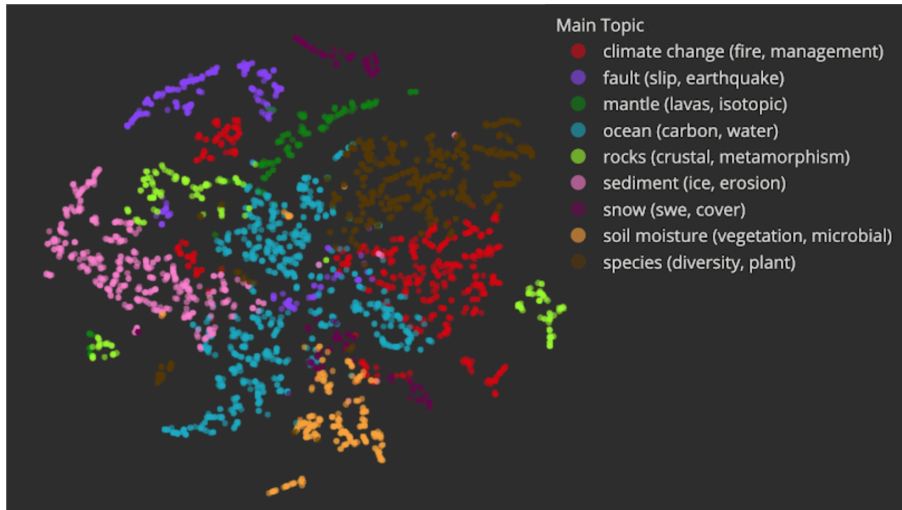
Figure 7: Coarse (9 topic) map of research documents (2009 - 2019)

The cluster labeled "rocks" is small and discrete compared with the "species" cluster, suggesting that more of ERI's research is ecological rather than geological in nature; however, these disciplinary identities are not mutually exclusive. Documents can be characterized by more than one research topic in the map.

Documents in the "soil moisture" cluster are uniformly located in a similar region of the map, while others like those in the "climate change" cluster, are dispersed and non-contiguous. This suggests a lack of internal conformity within this cluster. Lower document dispersion in the "soil moisture" cluster suggests topical homogeneity, while higher dispersion in the "climate change" cluster suggests more heterogeneous documents.

The adjacency of the "sediment" cluster with the "rocks", "climate change", and "ocean" clusters suggests that its documents straddle, and sometimes bridge, these research areas, particularly those on the clusters' edges. Clusters located farther apart are also dissimilar. The "snow" and the "soil moisture" clusters are found on opposite sides of the map; however, other documents described by these topics are neighboring at the bottom of the map, converging around an edge of the "climate change" cluster. Indeed, the documents found there bridge these areas; they address snowmelt, surface temperature in forests, biomass accumulation, streamflow changes, and other related ideas.

While the coarse map presents a distant overview of ERI's research topics, the detailed map shown in Figure 8 reveals intricate patterns. The center "population" cluster borders other research areas including the "species", "ocean", and "fisheries" clusters. Another multi-topic cluster found at the bottom map periphery gathers similar public policy research from different topics, like mitigating climate change impacts on fisheries and earth system science in Canada.

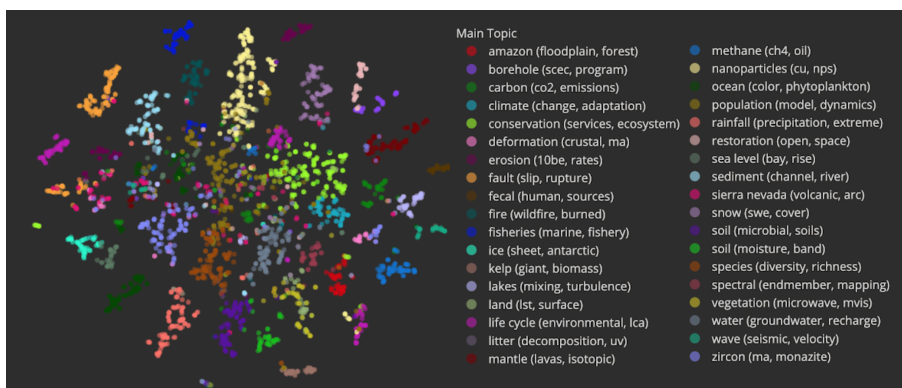The detailed map is made up of relatively even distributions of topic clusters.

15

Figure 8: Detailed (36 topic) map of research documents (2009 - 2019)

One exception is the "fecal" cluster on the right edge of the map, which is small and separated; its nearest neighbor is the "lakes" cluster below it. A larger "nanoparticles" cluster at the top of the map is associated with ERI's productive Center for Environmental Implications of Nanotechnology.

Central clusters tend to be less uniform than those at the edges. The "water", "conservation", and "methane" topic clusters are interspersed with documents addressing marine isotopes, stream mapping at a battlefield conservation site, and stream nitrate concentrations in mountainous watersheds. This is contrasted with the homogeneous clusters found at the edges, such as the "ice" cluster on the left edge dominated by documents addressing glaciers.

In the detailed map, we see that there are distinct, yet adjacent, areas of research involving similar researchers and shared ideas, such as integrating wildfire risk with the study of agricultural encroachment. The "conservation" and "fire" clusters are adjacent in the detailed map; in the coarse map, these documents fall under the "climate change" topic. In the detailed map, most "fire" research documents border the "sediment" and "fisheries" clusters, suggesting that documents about wildfire recovery and river restoration share similarities.

We have presented maps at two selected levels of detail - coarse (9 topics) and detailed (36 topics) - as described in our methods (Section 3). Yet, our method can reproduce and generate maps of research topics at any level of detail. The maps are systematically produced with the goal of improving upon the ad-hoc definition and interpretation of research thrusts in the institutional review process. "Reading" these data-driven maps generates qualitative insights; they represent topics extracted from the text of research documents that require viewers to perceive patterns in them. They also possess emergent qualities, revealing more than the sum of their parts (Moretti, 2005); they show patterns in ERI's research that were previously difficult or impossible to see when inspecting single documents, publication and project lists, or the work of individual researchers.

4.2 Deploying a Map Dashboard

To distribute and evaluate our maps, we deploy a public-facing dashboard[11] using Plotly, Dash for Python, and Heroku. The dashboard's "About" panel describes the map and allows users to select a level of detail, topics to map, and a year range. Figure 9 shows the "Search" panel, which allows users to filter data by ERI researcher or by academic department, returning metadata for a selected document, including its DOI when available. We make time explicit by showing a map snapshot for each year, which can be filtered by a range of years. This provides a backdrop for the interpretation of events, such as the acquisition of major grants or the hiring of new faculty in growing research areas. We provide evidence supporting these interpretations in our evaluation (Section 5).
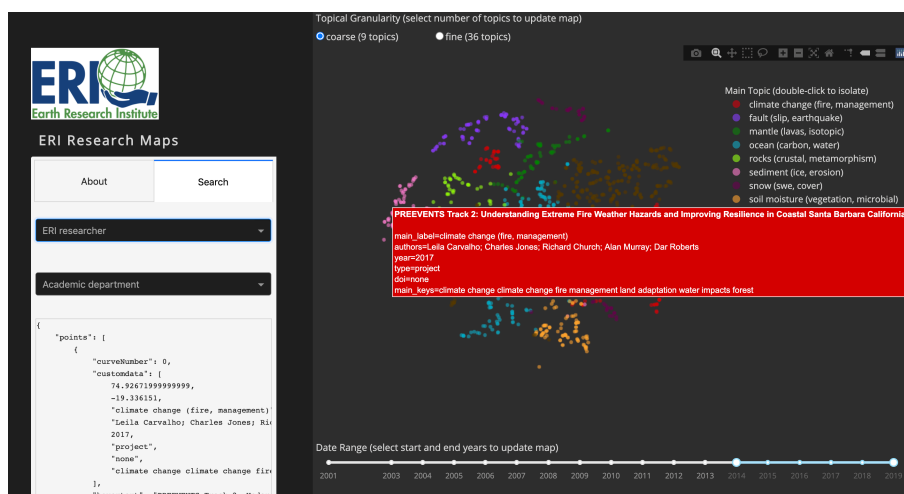
Figure 9: Search panel of the research map dashboard showing interactive features

5 Evaluation

Do the maps we developed support "distant reading" of research documents in the context of an institutional review? To answer this question more systematically, we evaluate the maps in three main ways. First, we use the maps to interpret and answer standard questions asked in the institutional review process. Second, we evaluate the maps in action, considering how they are used by the researchers whose work is being assessed (Solomon, 2002). We survey leading ERI researchers who determine if and how they think the maps support "reading at a distance". Finally, we assess the relationship between the

---

[11]https://eri-research-dashboard.herokuapp.com/

two maps by clustering the detailed topics and comparing the resulting clusters with the coarse topics.

5.1 Institutional Review Questions

How do maps of research topics support questions commonly posed to reviewers? Here, we consider the six institutional review questions about research accomplishment that UCSB's organized research units must regularly address[12]. They are currently answered using quantitative evidence, for example, numbers of publications by field of research and amounts of funding per researcher. While these benchmark questions are particular to UCSB, the concerns they address seem fairly representative of similar contexts elsewhere:

- **Research Quality and Significance:** Describe the quality and significance of research accomplished and in progress.

- **Trends and Research Specialties:** Comment on significant trends within the disciplines represented in the unit and relate these to current research specialties in your organized research unit (ORU).

- **Benefits to Campus and Departments:** Comment on how the ORU benefits the campus in general and academic departments in particular.

- **Participant Productivity, Influence, and Prominence:** Comment on the continuing productivity and influence of unit participants, locally as well as nationally. Comment on evidence of prominence in the fields represented in the ORU.

- **Collaborations and Interdisciplinarity:** Comment on the unit's collaborative/interdisciplinary work, its quality, and its impact on ORU research efforts and the campus.

- **Extramural Funding:** Describe the possible sources and availability of extramural funds to support the unit's-research. Are your participants sufficiently active in the pursuit of extramural funds in light of funding possibilities? How does the extent of annual extramural research funding compare with similar units nationwide?

We have claimed that maps of research topics can complement current evaluation metrics by supporting qualitative narratives. Here, we show how each of these questions can be addressed with maps of research topics:

Research Quality and Significance. Journal impact factors are a typical quantitative metric. Our maps complement this by generating a broader picture of cross-disciplinary topics from research publications. They highlight researchers' and departments' main topics and topical reach (diffuse or tightly

---

[12]https://www.research.ucsb.edu/organized-research-unit-oru-administration

clustered). Researchers in the Bren School of Environmental Science and Management are represented across all major topics, while those affiliated with Biology concentrate mainly in the "species" and "oceans" topic clusters.

Trends and Research Specialties. Funding agency priorities (e.g. NSF's "10 big ideas") and publisher classification schemes (e.g. fields of research) are typical sources of evidence. Our maps define research topics emerging from publications and projects that are not constrained by external classification schemes or historic disciplinary boundaries. The detailed map captures the topical diversity of research across affiliations while the coarse map emphasizes earth and environmental science topics unifying ERI's researchers.

Benefits to Campus and Departments. Evidence includes faculty recruitment, research computing infrastructure, and educational outreach programs. Temporal sequencing in our maps can be used to assess the impact of events, like the inception of educational programs (e.g. the Kids in Nature Program in 2012) or influential funding (e.g. a 2017 NSF award to upgrade campus computing resources). While causality cannot be determined, it is interesting to note growth in certain topic areas following these events (e.g. a rise in ecological restoration projects following the start of educational programming and community outreach). These insights provide concrete and solid support over anecdotal discussions in institutional reviews.

Participant Productivity, Influence, and Prominence. The professional accolades of individual participants, such as awards, are often reported as evidence. Our maps provide a more objective picture of the topics that each researcher addresses by showing the topical distribution of each researcher's documents. For example, geographer David Siegel's work is concentrated mainly in the "ocean" and "species" topic clusters while geographer Dar Roberts' work is more broadly dispersed across "species", "climate", "ocean", "snow", "sediment", and "soil moisture". While both accomplished researchers work extensively with remotely sensed imagery, differentiating their areas of expertise supports institutional management and reporting.

Collaborations and Interdisciplinarity. The affiliations of collaborators on funded projects are typically offered as evidence of interdisciplinarity. Our maps currently annotate each project by a single researcher and do not emphasize projects that have collaborators from multiple departments. This functionality could be added if ERI's leadership were interested to see who drives collaborations, not just what common topics they address.

Extramural Funding. This is currently based on award amounts. Our maps do not incorporate this kind of information, as these indicators already serve their purpose. The projects currently shown in the map have all been funded, but it could be valuable to also show the topics of unfunded projects, for example revealing changes to topics prioritized over time by funding agencies.

5.2   Researcher Survey

How do ERI's leading researchers interpret their own role in ERI's evolving research? We seek to understand researchers' interpretations of topics and re-

lationships shown in the maps. To gather feedback, we administered an online survey to researchers on ERI's advisory board. This survey also served as a rehearsal and internal review for the imminent 5 year review in which the primary map users will be external reviewers in leadership positions at similar institutes. The survey was kept intentionally short and contained the following items:

- ERI Topics: Take a minute to explore the first map, both at the coarse and fine granularity. How well do you think these topics represent ERI's research overall?

- PI Topics: Next, find yourself in the Search panel. Your publications or projects will be highlighted. How well do you think this map represents your research?

- Topic Evolution: Finally, try filtering the research documents using the time slider. Do you notice any trends and do these coincide with any events in ERI's history that you can recall?

- Other: Do you have any other comments or ideas for improving this tool?

We received responses from 5/13 members of the ERI advisory board in the 4 days that were allocated to the survey. The main ideas that emerged from the responses can be separated into observations made from the maps and comments about map design. These responses provide suggestive evidence, which is summarized as follows:

ERI Topics. A majority (3/5) of respondents felt that the coarse map adequately described ERI's research while the remainder had some objections. One noted that the coarse map "lacks several important categories (e.g, biogeochemistry, inland waters, carbon cycle)" but that "the detailed map represents the range of research". Another felt that the topics reduced all of ERI's research to "physical entities" that made it seem like a geology department. These concerns may relate to the design decision to label and color the documents by main topics; the labels include the first term from the topic descriptor with the second and third included in parentheses. Since topic modeling does not produce labels for the resulting topics, any succinct labeling in support of readability and verbalization skews the presentation. This feedback suggests that alternative approaches to labeling the topics could help because the objections raised were related to category names rather than the clustering of documents.

Researcher Topics. Respondents (3/5) felt that they understood the positions of their documents relative to ERI's research landscape. Several mentioned that their "assignments" aligned with their identities as researchers; one noted "I was largely in the species topic group and I do identify as a species-based researcher". Another felt that their work was categorized "imperfectly at best" as they work mainly on carbon cycling but had been associated with soils. These observations raise interesting challenges for visualizing perceived differences between researchers' self-assigned specialties and positions assigned to their work based on a relatively short period of time.

Topic Evolution. One researcher stated that trends in the map pointed to the founding of the UC Center for Environmental Implications of Nanotechnology at UCSB in 2013. Another noted that the map "appears to start out along the edges then fills in the middle... maybe it is selective hiring of people to bridge gaps?" These interpretations speak to the utility of the spatialization approach; researchers are able to associate patterns in the map with probable events in which interdisciplinary research topics emerge, bridging traditional clusters. Changes in topical "coverage" following a faculty hire or large funding awards were observable to the respondents when they used the maps in combination with the time slider. Their observations demonstrate the kinds of insights that we envisioned the temporally sequenced maps might offer.

Other. Most of the comments about map functionality address click interactions, background color, alphabetization of lists, and other details that are readily changed. Suggestions for additional functionality included ways to browse lists of related documents based on shared topics, to "visualize closely linked topics", and to search based on grants and papers. We expect to incorporate respondents' suggestions in preparation for the upcoming institutional review. We take the leading researchers' responses as a qualified endorsement of the generalization and visual presentation of work done at their institute.

5.3 Relationship Between Levels of Research Topics

Does the relationship between the detailed and coarse topics form a hierarchical tree or a lattice? This consideration is related to the granularity of the representation; many real-world phenomena, like administrative areas - cities in states within countries - can be represented by a tree but may function more like a lattice - city metropolitan regions spanning state boundaries (Hirtle, 1995). To study the relationship between topics in the detailed and coarse maps, we generate a dendrogram, which is a tree configuration to cluster the topics. It is different from the trees in distant reading (Moretti, 2005), which capture the evolution of topics over time. Rather, the tree shown in Figure 10 coarsens the detailed topic model by clustering the 36 detailed topics into 9 clusters based on their similarity.

We use SciPy[13] to cluster the topics and render them in a dendrogram. We experimented with different linkage methods for clustering the topics and selected Ward's minimum variance (Ward Jr, 1963). It is agglomerative, beginning with each topic in a separate cluster; at each step, it merges the two most similar clusters, resulting in compact, even-sized clusters (Aldenderfer & Blashfield, 1984). The vertical axis of the dendrogram represents the distance, or dissimilarity, between topics and the horizontal axis shows the labeled topics. Each of the nine clusters are assigned a color; for example, cluster A in the dendrogram is colored green and consists of three topics: 'conservation' and 'population, which are the most similar, and 'ocean'. The smallest clusters (A, D, F, G, and I) consist of three topics each and seem to be the most coherent,

---

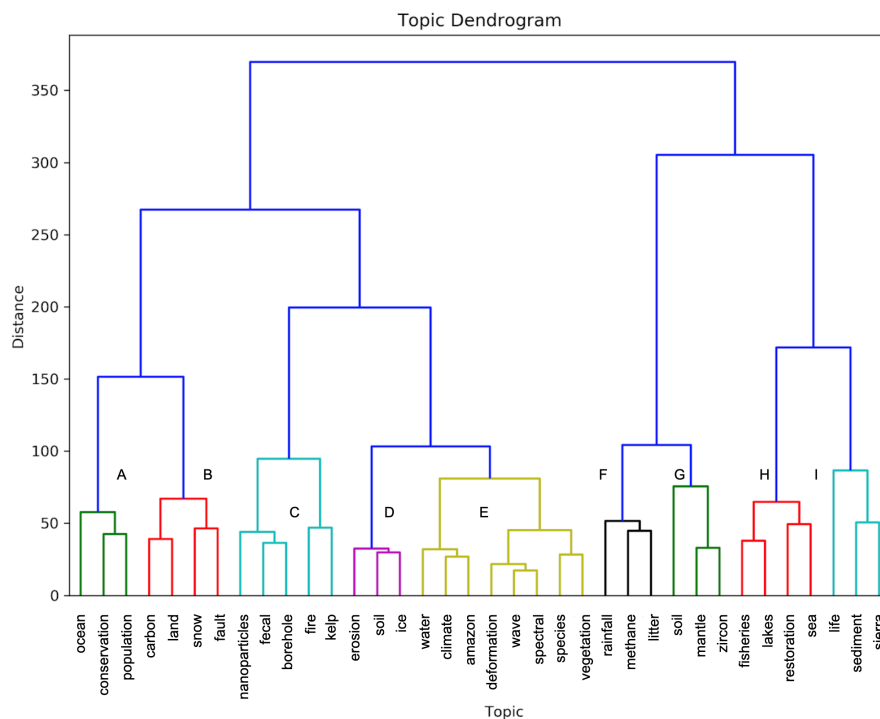[13]https://docs.scipy.org/doc/scipy/reference/cluster.hierarchy.html

Figure 10: Dendrogram of 36 clustered topics labeled by the first term in their descriptors

as to be expected. The largest cluster (E) consists of 8 topics and appears to be the most general, combining 'water', 'climate', 'vegetation', and 'species'.

To study the relationship between the topic models, we compare the resulting 9 topic clusters to the 9 coarse topics (Table 2). The "clusters" column records if the topics in a cluster share at least one term in common with a coarse topic. Terms associated with the clustered topics are based on the union of each member topic's ten most representative terms.

Topics 2 and 8 share terms with the topics in clusters 5 and 7. The other coarse topics seem to only partially correspond with the clusters. Topic 9 in the coarse model has the terms "sediment", "ice", and "erosion". In the detailed model, these terms are split into distinct topics: "sediment" is a term in Topic 28, "ice" is a term in Topic 10, and "erosion" is a term in Topic 15. Each detailed topic is associated with a limited number of coarse topics, suggesting that the detailed topics are not hierarchical refinements of coarse ones, but that moving between levels (i.e. "more" or "less distant") actually reconfigures the terms into a lattice of related topics.

Topic	Keywords	Clusters
1	ocean carbon water co2 phytoplankton global surface organic color emissions	A, B, D, E, F, G
2	ma rocks crustal metamorphism monazite crust zircon metamorphic deformation exhumation	E, G
3	snow swe snow cover cover water snow water modis snowmelt model water equivalent	A, B, E
4	soil soil moisture moisture vegetation microbial microwave surface soils band plant	B, D, E, G
5	climate change climate change fire management land adaptation impacts water forest	A, B, C, D, E, H
6	species diversity plant richness native biomass communities biodiversity effects ecosystem	A, C, E, G
7	fault slip earthquake rupture seismic motion ground faults earthquakes ground motion	B, C, E, I
8	mantle lavas isotopic crust 3he melt samoan geochemical 4he 3he 4he	E, G
9	sediment ice erosion rates 10be river sea glacial erosion rates precipitation	D, E, F, H, I

Table 2: Topics, descriptors with ten top terms, and corresponding topic clusters

6   Conclusions

The evaluation of the research maps demonstrates their potential to complement current approaches to institutional reviews. Our data-driven approach uses topic modeling and spatialization to derive high level topics from ERI's body of research and make them discoverable by researchers and departments.

In the next phase of this research, we envision a dynamic system for generating maps of research topics based on a continually updating corpus of research documents, which would require supervised classification of unseen documents based on our current data. We are also interested in highlighting relationships among topics, departments, researchers, or other units by presenting the research documents as a network. While we successfully developed field and network views of research documents in our previous work (Lafia et al., 2019), we did not exploit the potential of networks in this research for several reasons. For one, networks based on quantitative information (e.g. co-citation frequency) are already established scientometric techniques (Börner et al., 2003) and we are interested in developing complementary approaches to systematically model research topics. As ERI is a multidisciplinary institution, we also did not want to draw imbalanced comparisons between the contributions of individual researchers from different disciplines; to avoid this, we would prefer to produce network views of other significant units like academic departments. Finally, our approach embeds research documents as nodes; however, linking embedded ob-

jects would not allow us to see emerging patterns like documents that are highly connected and central.

A limitation of our approach is that while spatial information has three components (space, time, theme) we primarily take advantage of the thematic dimensions of data. While we incorporate temporal views in our maps, allowing for document subsetting by timespan and event detection, we could also consider making time the primary dimension. Previous work developing semantic signatures shows that time and space offer two ways to order knowledge (Janowicz, 2010). Views ordered primarily by time could be thought of as temporalizations, rather than the spatializations we develop, and could take the form of graphs from distant reading, where the evolution of specific topics could be tracked.

We are also interested in extending the evaluation of our work to coincide with ERI's external review. This would help us understand how external reviewers who do not have a personal connection to ERI's research interpret and evaluate its research topics. To determine the applicability and maturity of our approach for adoption in a broader context, we would also be interested in expanding our survey to include researchers or leaders affiliated with similar organized research units. This would allow us to build consensus around strategies for adopting maps of research as robust spatial decision support tools.

At the outset of this article, we argued for maps of the research "territory" to provide decision support. The maps we have produced give an impression of the underlying thematic structure of the research in the form of "research neighborhoods" that are meaningful within, and possibly across, institutions. Just as land use maps are used to manage resources and forecast growth in a regional planning context, maps of research can be used to do the same in an institutional setting. We envision maps of research topics being used internally (as part of the organized research unit's self-assessment) and externally (as a communication tool describing research trends and developments) which may be of interest to external reviewers, other research units, and the public.

7   Acknowledgements

References

Aldenderfer, M. S., & Blashfield, R. K. (1984). Validation Techniques. In *Cluster Analysis* (pp. 62–73). Sage Publications London.

---

[14]http://spatial.ucsb.edu/research/spatial-discovery

Allen, C., Hervey, T., Lafia, S., Phillips, D. W., Vahedi, B., & Kuhn, W. (2016). Exploring the Notion of Spatial Lenses. In *The Annual International Conference on Geographic Information Science* (pp. 259–274).

Andrienko, G., Andrienko, N., Jankowski, P., Keim, D., Kraak, M.-J., MacEachren, A., & Wrobel, S. (2007). Geovisual Analytics for Spatial Decision Support: Setting the Research Agenda. *International Journal of Geographical Information Science*, *21*(8), 839–857.

Arias, S., & Warf, B. (2009). Introduction: the Reinsertion of Space into the Social Sciences and Humanities. In *The Spatial Turn: Interdisciplinary Perspectives*. Routledge, London.

Arora, S., Ge, R., Halpern, Y., Mimno, D., Moitra, A., Sontag, D., ... Zhu, M. (2013). A Practical Algorithm for Topic Modeling with Provable Guarantees. In *International Conference on Machine Learning* (pp. 280–288).

Bawden, D., & Robinson, L. (2015). Informetrics. In *Introduction to Information Science* (pp. 165–181). Facet Publishing.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, *3*(Jan), 993–1022.

Börner, K., Chen, C., & Boyack, K. W. (2003). Visualizing Knowledge Domains. *Annual Review of Information Science and Technology*, *37*(1), 179–255.

Boyd-Graber, J., Mimno, D., & Newman, D. (2014). Care and Feeding of Topic Models: Problems, Diagnostics, and Improvements. In *Handbook of Mixed Membership Models and their applications* (pp. 225–255). CRC Press Boca Raton, FL.

Fabrikant, S. I., Ruocco, M., Middleton, R., Montello, D. R., & Jörgensen, C. (2002). The First Law of Cognitive Geography: Distance and Similarity in Semantic Space. *Proceedings of GIScience 2002*, 31–33.

Gärdenfors, P. (2000). Semantics. In *Conceptual Spaces: The Geometry of Thought* (pp. 151–201). MIT Press.

Greene, D., O'Callaghan, D., & Cunningham, P. (2014). How Many Topics? Stability Analysis for Topic Models. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases* (pp. 498–513). Springer.

Griffiths, T. L., Jordan, M. I., Tenenbaum, J. B., & Blei, D. M. (2004). Hierarchical Topic Models and the Nested Chinese Restaurant Process. In *Advances in Neural Information Processing Systems* (pp. 17–24).

Hermann, M., & Leuthold, H. (2003). *Atlas der politischen Landschaften: Ein weltanschauliches Porträt der Schweiz*. Vdf Hochschulverlag an der ETH.

Hirtle, S. C. (1995). Representational Structures for Cognitive Space: Trees, Ordered Trees and Semi-Lattices. In *International Conference on Spatial Information Theory* (pp. 327–340). Springer.

Ioannidis, J. P. (2018). Meta-Research: Why Research on Research Matters. *PLoS Biology*, *16*(3), e2005468.

Janowicz, K. (2010). The Role of Space and Time for Knowledge Organization on the Semantic Web. *Semantic Web*, *1*(1, 2), 25–32.

Jappe, A., Pithan, D., & Heinze, T. (2018). Does Bibliometric Research Confer Legitimacy to Research Assessment Practice? A Sociological Study of Reputational Control, 1972-2016. *PloS One*, *13*(6), e0199031.

Kuhn, W. (1996). Handling Data Spatially: Spatializating User Interfaces. In *Advances in GIS Research II: Proceedings of the 7th International Symposium on Spatial Data Handling* (Vol. 2, p. 13B).

Kuhn, W. (2012). Core Concepts of Spatial Information for Transdisciplinary Research. *International Journal of Geographical Information Science*, *26*(12), 2267–2276.

Lafia, S., Last, C., & Kuhn, W. (2019). Enabling the Discovery of Thematically Related Research Objects with Systematic Spatializations. In *14th International Conference on Spatial Information Theory (COSIT 2019).* Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Lee, D., & Seung, H. S. (1999). Learning the Parts of Objects by Non-negative Matrix Factorization. *Nature*, *401*, 788–791.

Maaten, L. v. d., & Hinton, G. (2008). Visualizing Data Using t-SNE. *Journal of Machine Learning Research*, *9*(Nov), 2579–2605.

McGuire, P. (2019, Oct). How Higher Education's Data Obsession Leads Us Astray. *The Chronicle of Higher Education*. Retrieved from https://www.chronicle.com/article/How-Higher-Education-s-Data/247409

McInnes, L., Healy, J., & Melville, J. (2018). UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv preprint arXiv:1802.03426*.

Miller, G. A. (1956). The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. *Psychological Review*, *63*(2), 81.

Mimno, D., Wallach, H. M., Talley, E., Leenders, M., & McCallum, A. (2011). Optimizing Semantic Coherence in Topic Models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (pp. 262–272).

Moher, D., Naudet, F., Cristea, I. A., Miedema, F., Ioannidis, J. P., & Goodman, S. N. (2018). New Principles for Assessing Scientists. *Issues in Science and Technology*, *35*(1), 20–23.

Moretti, F. (2005). *Graphs, Maps, Trees: Abstract Models for a Literary History*. Verso.

Polonioli, A. (2020). In Search of Better Science: On the Epistemic Costs of Systematic Reviews and the Need for a Pluralistic Stance to Literature Search. *Scientometrics*, *122*(2), 1267–1274.

Radicchi, F., Fortunato, S., & Castellano, C. (2008). Universality of Citation Distributions: Toward an Objective Measure of Scientific Impact. *Proceedings of the National Academy of Sciences*, *105*(45), 17268–17272.

Salton, G., Wong, A., & Yang, C. (1975). A Vector Space Model for Automatic Indexing. *Communications of the ACM*, *18*(6), 391–398.

Skupin, A., & Fabrikant, S. I. (2003). Spatialization Methods: A Cartographic Research Agenda for Non-geographic Information Visualization. *Cartog-*

raphy and Geographic Information Science, *30*(2), 99–119.

Solomon, P. (2002). Discovering Information in Context. *Annual Review of Information Science and Technology*, *36*(1), 229–264.

Tobler, W. R. (1970). A Computer Movie Simulating Urban Growth in the Detroit Region. *Economic Geography*, *46*(sup1), 234–240.

Underwood, T. (2016). Distant Reading and Recent Intellectual History. *Debates in the Digital Humanities 2016*, 530–533.

Vinkler, P. (2010). Indicators are the Essence of Scientometrics and Bibliometrics. *Scientometrics*, *85*(3), 861–866.

Ward Jr, J. H. (1963). Hierarchical Grouping to Optimize an Objective Function. *Journal of the American Statistical Association*, *58*(301), 236–244.

Wattenberg, M., Viégas, F., & Johnson, I. (2016). How to Use t-SNE Effectively. *Distill*, *1*(10), e2. doi: 10.23915/distill.00002

Wise, J. A., Thomas, J. J., Pennock, K., Lantrip, D., Pottier, M., Schur, A., & Crow, V. (1995). Visualizing the Non-Visual: Spatial Analysis and Interaction with Information from Text Documents. In *Proceedings of Visualization 1995 Conference* (pp. 51–58).