

UC Santa Barbara

UC Santa Barbara Previously Published Works

Title

Mapping research topics at multiple levels of detail.

Permalink

<https://escholarship.org/uc/item/6xx6z4b7>

Journal

Patterns, 2(3)

Authors

Lafia, Sara

Kuhn, Werner

Caylor, Kelly

et al.

Publication Date

2021-03-12

DOI

10.1016/j.patter.2021.100210

Peer reviewed

Patterns

Mapping research topics at multiple levels of detail

Highlights

- A method for modeling and mapping topics from bibliometric data
- A web application dashboard with interactive, multilevel topic maps
- A user evaluation of topic mapping's utility in institutional reviews

Authors

Sara Lafia, Werner Kuhn, Kelly Caylor, Libby Hemphill

Correspondence

slafia@umich.edu

In Brief

Many institutions use metrics to evaluate their research productivity; however, it is challenging to effectively summarize and evaluate research across academic disciplines. We use topic modeling to develop maps of science at multiple levels of detail. The maps were used in an institutional review and evaluated by leading researchers at an earth science institute. We demonstrate that mapping research topics supports the review process by offering insights into interdisciplinary research collaborations and areas of expertise at the institute.



Article

Mapping research topics at multiple levels of detail

Sara Lafia,^{1,3,*} Werner Kuhn,² Kelly Caylor,² and Libby Hemphill¹¹ICPSR, University of Michigan, Ann Arbor, MI 48109, USA²Department of Geography, University of California, Santa Barbara, CA 93106, USA³Lead contact*Correspondence: slafia@umich.edu<https://doi.org/10.1016/j.patter.2021.100210>

THE BIGGER PICTURE Research institutes and organizations are interested in communicating the impact of their work and its value to a broader audience. However, quantifying impact and providing high-level views of interdisciplinary research trends are challenging. To address this, we leverage distant reading methods from the digital humanities to model the topics of a large body of interdisciplinary research products and visualize them in maps. We analyze 3,770 academic publications and grants affiliated with an interdisciplinary earth science research institute over a 10-year period and model its research topics. We then map the topics at two distinct levels of detail and evaluate the interpretation of the maps through a survey of leading researchers. We show that the topic maps reveal insights including the emergence of interdisciplinary collaboration areas and evolving areas of expertise over time.



Proof-of-Concept: Data science output has been formulated, implemented, and tested for one domain/problem

SUMMARY

The institutional review of interdisciplinary bodies of research lacks methods to systematically produce higher-level abstractions. Abstraction methods, like the “distant reading” of corpora, are increasingly important for knowledge discovery in the sciences and humanities. We demonstrate how abstraction methods complement the metrics on which research reviews currently rely. We model cross-disciplinary topics of research publications and projects emerging at multiple levels of detail in the context of an institutional review of the Earth Research Institute (ERI) at the University of California at Santa Barbara. From these, we design science maps that reveal the latent thematic structure of ERI’s interdisciplinary research and enable reviewers to “read” a body of research at multiple levels of detail. We find that our approach provides decision support and reveals trends that strengthen the institutional review process by exposing regions of thematic expertise, distributions and clusters of work, and the evolution of these aspects.

INTRODUCTION

Universities and funding agencies request that organized research units (ORUs) summarize and report on their research, collaboration, and growth as part of periodic institutional reviews. These reviews typically ask questions about trends in research quality, significance, research specialties, areas of influence or prominence, and interdisciplinarity collaborations. The review process is not unique to universities or research institutes; many kinds of organizations, including those in non-governmental, governmental, and industry settings, regularly conduct “meta-research”¹ on their activities in order to provide

a high-level view of their impact and productivity. Yet, it remains unclear how best to summarize and present interdisciplinary bodies of work in ways that generate useful insights and can support effective reviews.

Bibliometrics and scientometrics support the quantitative study of published documentation and academic disciplines;² they have 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 collaborators or competitors.³ 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.⁴

Quantitative metrics like impact factors, however, have been recognized as poor choices for assessing or comparing research output of scholars and journals. They are often not comparable across academic disciplines⁵ and have been found to be vulnerable to manipulation.⁶ A study of the relationship between journals and citation rates has demonstrated evidence of a cumulative advantage for publications in “high-impact” journals.⁷ The single numbers these metrics produce also obscure differences between disciplines and outlets over time.

Alternative quantitative metrics have been developed in response to these limitations. The Eigenfactor metrics⁸ consider author centrality in citation networks, while the SCImago index⁹ considers the flow of prestige between thematically related journals. Altmetrics¹⁰ capture a more comprehensive picture of the ecosystem of scientific products and activity, like discourse about scientific software, that goes beyond the partial view from formal citations. These metrics are reshaping how scientists value research products and assess impact.¹¹

In this vein, there is a growing desire for interdisciplinary research evaluation that can more adequately capture impact and quality. One strategy has been to complement quantitative metrics with high-level characterizations and narratives.¹² Another has been to develop maps that chart the structure of knowledge domains and show the development of research areas, their interconnections, and evolution within them.¹³ These approaches offer more contextual information than single-measure quantitative metrics. Science maps are examples of spatializations,¹⁴ which use space as a metaphor to map abstract domains to thematic spaces in which nearby elements are similar. They can help evaluators process more information¹⁵ than can be effectively communicated by a single quantitative metric; they also make patterns and trends more apparent.

In this article, we examine the utility and benefits of spatialization to produce maps of research that support an institutional review, specifically by revealing trends and providing decision support. To develop and test our ideas, we situate our study in the context of an ORU at the University of California at Santa Barbara (UCSB): the Earth Research Institute (ERI) (<https://www.eri.ucsb.edu/>). ERI’s stated mission is to “support research and education in the sciences of 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’s faculty and researchers are supported by 145 different funding agencies covering the full breadth of earth and environmental sciences.

To date, ERI has taken an *ad hoc* approach to characterizing its research. For example, anecdotal observations based on faculty hires from ERI’s last institutional review indicated that its expertise had broadened from traditional earth science and crustal studies to include conservation and biodiversity topics. To formally capture and verify this kind of institutional knowledge about ERI’s evolving research expertise, we propose a data-driven approach for eliciting cross-cutting research topics. Our approach demonstrates how science mapping can complement current quantitative or *ad hoc* approaches to institutional reviews by uncovering

trends and relationships obscured by other metrics. We produced research maps that capture the latent thematic structure of an interdisciplinary body of research at multiple levels of detail. To do this, we analyzed research publications and funded projects from 240 researchers spanning 24 academic departments affiliated with ERI between 2009 and 2019. We then evaluated the insights that the maps can support by surveying researchers within the institution whose work is represented in the maps.

In the remainder of this article, we situate our work in relation to existing approaches for abstracting and mapping information. Specifically, we discuss science mapping as a method for domain analysis and knowledge representation. We then describe our approach to produce maps of a body of research at two levels of detail. Finally, we report how leading ERI researchers evaluate the potential for our maps to support an institutional review. A delay in the actual institutional review (resulting from COVID-19) precluded feedback from external reviewers in time for our research project. We find that our approach complements the review process by exposing and relating thematic expertise, highlighting relationships between academic departments or teams of authors, analyzing topical distributions and clusters of work, and tracking the evolution of these aspects over time.

Background

The interpretation of interdisciplinary research trends and impact is an important task for many research institutions, and single-value quantitative metrics are insufficient. We review methods that facilitate trend and impact analysis by abstracting and visually summarizing large collections of research documents. To situate our contribution, we first review science mapping applications in scientometrics and knowledge domain visualizations. We then describe dimensionality reduction and data visualization techniques used to design science maps, namely topic modeling and clustering techniques.

Science mapping

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.¹⁶

Science mapping charts the structure and evolution of knowledge in a domain or discipline by using maps as visual communication metaphors.¹³ Science maps are based on bodies of scientific literature analyzed using computational tools and visualized to highlight trends, which can be interpreted using theories of scientific change.¹⁷ Scientometric applications use quantitative metrics, including author co-citation,¹⁸ document or journal co-citation,¹⁹ co-word analysis,²⁰ and other bibliometrics extracted from documents. Many applications configure bibliometric elements using multidimensional scaling, network analysis, tree maps, or other visualization techniques.²¹ Similarity measures are constructed and applied along with dimensionality reduction to visualize scientific documents.¹³

A number of recent applications combine topic modeling with interactive visualizations to provide decision support. A visual

topic modeling system called UTOPIAN²² combines several dimensionality reduction techniques, including topic modeling and clustering, to merge or split topics based on user input. A related system called Termite²³ presents salient terms discovered from each topic, which can be used to explore documents. Other systems for visualizing and interpreting topics include LDAvis,²⁴ TopicLens,²⁵ and VISTopic.²⁶ Like Termite, LDAvis supports interpretation of relevant relationships between terms and discovered topics; topics are presented in a low-dimensional view, showing their correspondence with terms. Like UTOPIAN, TopicLens responds dynamically to user input by regenerating multilevel topic models and embeddings based on user specifications. Similarly, VISTopic supports multilevel topic representation but partitions the corpus of input documents hierarchically.

Although our work bears similarities to these systems, we distinguish our contribution as follows. First, several of these existing systems allow users to adjust the level of detail in the visualizations, which is handled hierarchically. Strict hierarchies may not offer the best knowledge representation, however, especially in applications like institutional reviews where topical overlap is of interest. For example, a coarse representation of a corpus may have a topic about “ecology,” while a more detailed representation may have topics about “nutrient cycling” and “predation”; while related, these topics can also be independent of the more general “ecology” topic. Alternative tree-like structures, like semilattices or sets of partially overlapping concepts, might be more adequate for knowledge organization.²⁷ We chose not to take a hierarchical approach when modeling topics. Instead, we handle level of detail by selecting numbers of topics in advance.

Second, we chose not to exploit the potential of network visualizations based on quantitative metrics like co-citation. Network-based measures are well established¹³ and support specific kinds of questions; in previous work, we found that embedding research objects based on their topical similarity revealed their distribution and the coverage of their corpus, while linking them revealed their topical connectivity and centrality.²⁸ As ERI is an interdisciplinary institution, however, we did not want to use metrics or create visual representations that would draw imbalanced comparisons between the contributions of individual researchers from different disciplines. Instead, we treat research documents as objects embedded in a continuous topic space, which form regions of research that change over time and vary by level of detail.

Finally, while many prior systems offer use cases with real data, few involve usability testing. We demonstrate the utility of our application, which is situated in a real institutional review. This allows us to collect valuable insights about science map interactions and interpretations as reported under [Evaluation](#).

In the following sections, we focus on dimensionality reduction and data visualization techniques that underpin science mapping and support the exploration and discovery of research documents at multiple levels of detail.

Dimensionality reduction

Dimensionality reduction is a key step in producing science maps, as it addresses the problem of displaying complex, high-dimensional data in a low-dimensional space like a two-dimensional map.¹³ This is analogous to cartographic general-

ization, where computational and cognitive issues of complexity are addressed by deliberately reducing the level of detail in the representation.²¹ To reduce the level of detail in our corpus of research documents, we use topic modeling to identify major themes shared by research documents. 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.

Topic models are statistical machine learning techniques that can uncover structures in collections of documents, for example, by grouping documents in which similar terms co-occur.²⁹ Topic models have been applied to classify and summarize large collections of documents, as well as solving similarity judgment problems.³⁰ Topics themselves can also be of interest; for example, the National Institutes of Health and the National Science Foundation have developed topic-based search interfaces to explore trends across related research projects.²⁹

We consider two main kinds of topic modeling approaches: latent Dirichlet allocation (LDA) and non-negative matrix factorization (NMF). LDA represents documents as mixtures of topics composed of words with certain probabilities.³⁰ It assumes that similar words occur in similar contexts and aims to discover latent topics in the documents. LDA offers insights “into inter- or intra-document statistical structure”³⁰ and has been positioned as an improvement over other measures used in information-retrieval applications like term frequency-inverse document frequency, or tf-idf,³¹ which is used to determine the relative importance of terms in a given document or corpus.

In matrix factorization approaches, a document-term matrix is decomposed into a smaller set of matrices, which can be interpreted as a topic model.³² NMF is a dimensionality reduction technique for decomposing samples, which are documents in topic modeling. Similar to LDA, documents are represented as term vectors, which can be combined into a document term matrix. However, documents are represented as combinations of co-occurring terms rather than likelihoods. In NMF, term weighting using tf-idf, for example,³¹ can also be used to boost distinctive terms.

A central challenge in topic modeling is the selection of an appropriate number of topics; selecting too few leads to overly broad topics, while selecting too many leads to redundancy.³³ Best practices recommend a combination of human evaluation strategies and topic coherence measures.²⁹ Coherence measures quantify the degree to which statements in a set support one another; in topic modeling, coherence measures evaluate sets of words that compose topics.³⁴

Data visualization

Data visualization controls the transformation and layout of data into a map.¹³ To visualize research documents, we use clustering methods to further abstract the topic models and give a visual impression of their underlying structure, in particular, the similarity between concepts. Broadly, the outputs of these clustering methods can be interpreted as spatializations, which offer high-level views of content through the familiar visual modality of maps.¹⁵

In general, space and time are fundamental ordering relations for knowledge representation.³⁵ The “spatial turn” observed in the social sciences and humanities has exploited the idea of spatial organization to facilitate cross-disciplinary exchange,

Table 1. Summary of core concepts of spatial information

Concept	Definition	Intuition	Question	Examples
Location	a description of where things are by spatial relations	things are located relative to one another 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 8:00 a.m. 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

allowing many lines of thought to converge.³⁶ In cognitive science, it has been claimed that conceptual spaces in which nearby concepts are similar underlie human thinking and learning.³⁷ 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.”³⁸ The distance-similarity metaphor treats distance in abstract spaces as metaphorically equivalent to dissimilarity.³⁹ These powers of spatial representation underpin the idea of spatialization, which maps abstract domains to spaces in which nearby elements are similar.¹⁴ Spatialization has been applied to organize multidimensional and thematically diverse collections.

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.³⁹ 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 representations and labels for the features at each level of detail.²¹

Spatialization methods are related to a broader suite of “macroscopic research” devices,⁴⁰ including science maps¹³ and “distant reading” diagrams⁴¹ that enable the study of patterns at multiple levels of detail over time. Distant reading,⁴¹ in

the digital humanities, provides methods for deliberately abstracting and visualizing text; 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. Distant reading uses graphs, maps, and trees to spatially configure units, like genres and novels, and reveal latent structures in their source material. These methods are generic enough to guide abstraction over many kinds of texts, which in our case are large numbers of abstracts from publications and grant proposals. They support a broader understanding of latent trends, such as the emergence and evolution of shared research topics.

To further systematize our spatialization methods, we apply the theory of core concepts of spatial information.⁴² The concepts summarized in Table 1 provide a high-level vocabulary with which to ask and answer questions about phenomena in space and time. They capture distinct ways of computing with spatial information; thus, they are applicable to geographic as well as other kinds of spaces. They provide us with a set of interchangeable lenses through which research data can be spatialized and viewed.⁴³

To produce maps, we first 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* are then *located* in this continuous two-dimensional topic space according to their topic mixtures using two embedding techniques: t-distributed stochastic neighbor embedding, or t-SNE,⁴⁴ and uniform manifold approximation and projection, or UMAP.⁴⁵ Both t-SNE and UMAP model high-dimensional research objects as points in a low-dimensional map space while clustering similar objects and spacing apart dissimilar ones. This embedding results in regions of documents in which *events*, like changes in the configurations of individual or departmental research, can be detected over time.

Methods

We produce maps that support the distant reading of the ERI's activities at distinct levels of detail. These maps show research topics and their evolution over time. The input to these maps are the descriptions of two kinds of research documents: publications and funded projects. We take the titles and abstracts from their metadata and model topics from them at two distinct levels of detail. We then feed the resulting document topic models into spatialization algorithms to output maps of the research topics.

Data sources

We analyzed publications and funded projects from ERI's 240 researchers active from 2009 to 2019. We gathered publication metadata using the Dimensions API, which is available for non-commercial use. We retrieved publications for each active researcher at ERI during the study period. These publications were then hand-curated by ERI staff to verify that they were associated with the correct researcher and sponsored by ERI during the period of analysis. This yielded 3,108 publications. We retained the title, abstract, year, digital object identifier (DOI) (if available), and authors. Examples of publication outlets include *PLOS One*, *Proceedings of the National Academy of Sciences*, and *Environmental Science & Technology*. Field-of-research




Figure 1. Distribution of research documents per year for analysis period

codes assigned to publications in Dimensions (<https://app.dimensions.ai/browse/categories/publication/for>) characterize major research areas and include earth sciences, biological sciences, environmental sciences, and engineering.

We also used ERI’s internal data on funded proposals, grants, and contracts. Similarly, we retained only the title, abstract, year, and identifier (if available). This yielded 662 funded projects. The majority of funding for projects came from federal agencies like the National Science Foundation, National Aeronautics and Space Administration, and National Oceanic and Atmospheric Administration. Partnerships with municipal and state agencies along with other universities also provided substantial funding. [Figure 1](#) summarizes the numbers of research publications and projects per year over the period of analysis.

Text pre-processing

We combined the metadata of the 3,770 research documents and performed text pre-processing by removing records with identical identifiers (DOIs or grant numbers), removing HTML tags, and reformatting ASCII extended characters. To determine whether to set a document length threshold, we checked the document distribution. [Figure 2](#) shows a normal distribution of lengths, which are relatively concise; the average document is 1,678 characters long.

Next, we followed a standard natural language processing pipeline to reformat the titles and abstracts of the research documents.⁴⁶ We first determined distinct document terms using tf-idf.³¹ This measure 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; we use this measure 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”). Many frequent terms describe research methods (e.g., “estimate”) rather than subject matter (e.g., “snow”).

We removed the following frequent and generic terms, which had low tf-idf scores: “data,” “study,” “project,” “research,” “collaborative,” “include,” “result,” “increase,” “high,” “low,” “large,” “include,” and “based.” We then constructed unigram and bigram models to preserve contiguous sequences of terms (e.g., “climate_change”). We did not lemmatize the input text because we did not want to lose the variation of domain-specific terms (e.g., “hydrology” and “hydrological”). We created a

normalized document term matrix composed of 3,770 documents and 80,152 distinct terms. We set the minimum document frequency to 2 and we considered both unigrams and bigrams. This resulted in a corpus of documents and term frequencies to use in topic modeling.

Topic modeling

We applied LDA³⁰ and NMF³² to the normalized document term matrix. Our goal was to model a range of topics for the documents and to generate coherent

topics at multiple levels of detail that describe major research themes at ERI.

To determine a range of topic values to model, we used Miller’s law⁴⁷ as a heuristic. It 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’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 reapplied Miller’s law to each chunk of the coarse level, resulting in bounds of (5×5) and (9×9) , or a range of 25–81 chunks, or in our case topics, to generate.

To compare the models and evaluate their quality, we use coherence as an interpretability measure. It is based on the fundamental idea in classification that the members of a class should be more similar to one another than to members of other classes and measures the extent to which top terms representing a topic are semantically related, relative to other terms in the corpus.⁴⁸ Coherence is considered to be more human interpretable for evaluating topic model quality than other measures, including perplexity and log likelihood.³³ Specifically, we use the topic coherence Word2Vec metric, which generates word embeddings to evaluate the similarity of term level descriptors from topics.⁴⁹

We generated LDA and NMF models across a range of topic numbers (2–100) and calculated their coherence scores. [Figure 3](#) shows a comparison between coherence scores for the LDA and NMF topic models. We generated LDA models using Gensim’s Mallet wrapper (<https://radimrehurek.com/gensim/models/wrappers/ldamallet.html>) and NMF models using Scikit-learn (<https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.NMF.html>). The NMF model was initialized with non-negative double-singular value decomposition (“nndsvd”), which is optimized for sparse data.

We found NMF to be a more suitable topic modeling approach for our purposes than LDA. It produced topic models with higher coherence scores than our LDA models by about 17% on average. This may be because NMF is better suited to modeling smaller or sparser datasets, like titles and abstracts, rather than full text.⁵⁰ We also found that NMF produced topics that were more indicative of subject matter, rather than methods. This




Figure 2. Distribution of research documents by word count

may be due to term weighting with tf-idf, unlike LDA, which operates on raw term frequency.³³

Although the addition of topics increases the coherence of the models, we wanted to select models that followed the Miller’s law heuristic we previously established; the NMF model with 100 topics has the highest coherence score, but this value is out of range. To select topic models, we relied on human evaluation⁵¹ of the most coherent models within a first range of 5–9 topics and a second range of 25–81 topics. Specifically, ERI’s director, Kelly Caylor, evaluated the topic descriptors for models within each range and selected two topic models to develop into maps: a coarse-grained model with 9 topics and a fine-grained model with 36 topics. This choice was important because we wanted to ensure that the themes emerging from the topic models were interpretable, in addition to being coherent, and could support institutional reporting.

Table 2 shows samples of topics and topic descriptors as a list of top terms for each of the NMF models we generated. Whereas most of the terms are unigrams, some bigrams, like “species_richness,” also capture scientific concepts that are compound terms. NMF results in a document-topic matrix in which each document is described by a mixture of topics with different strengths of association. The document-topic

matrix forms the input to the subsequent spatializations, while the topic-term matrix is used to reference topics and term descriptors.

Spatialization

The inputs to the spatializations are the document-topic matrices resulting from the coarse (9) and detailed (36) NMF topic models. We first mapped research documents with t-SNE using manifold learning in Scikit-learn (<https://scikit-learn.org/stable/modules/generated/sklearn.manifold.TSNE.html>). 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 one another and dissimilar documents farther apart. The UMAP process for assigning locations to research documents is similar 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 (<https://umap-learn.readthedocs.io/en/latest/>). 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.^{44,45}

We interactively explored the maps to interpret the effects of the map parameters, which balance local, pairwise similarity with global, intercluster similarity.⁵² 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 values allow for more global connectivity and less clearly delineated clusters. Typical values fall between 5 and




Figure 3. Coherence scores for NMF and LDA topic models with 2–100 topics

Table 2. Topics and descriptors discovered for NMF topic models

Topic	Coarse-grained descriptors	Fine-grained descriptors
1	ocean carbon water co2 phytoplankton global surface organic color emissions	wave seismic velocity rayleigh pressure surface structure wind noise
2	ma rocks crustal metamorphism monazite crust zircon metamorphic_ deformation exhumation	deformation crustal ma pamir shear himalayan rocks exhumation tibetan_himalaya
3	snow swe snow_cover cover water snow_ water modis snowmelt model water_equivalent	snow swe snow_cover cover snow_water water_ equivalent snowmelt_ equivalent snowpack snow_depth
4	soil soil_moisture moisture vegetation microbial microwave surface_ soils band plant	soil_moisture moisture soil band smap surface roughness surface_soil algorithm m3
5	climate change climate_ change fire management land adaptation impacts water forest	climate climate_change change adaptation future impacts models responses global species
6	species diversity plant richness native biomass communities biodiversity effects ecosystem	species diversity richness plant native species_ richness biodiversity communities biomass abundance
7	fault slip earthquake rupture seismic motion ground faults ground motion	fault slip rupture earthquake faults motion ground ground_motion seismic
8	mantle lavas isotopic crust 3he melt samoan geochemical 4he	mantle lavas isotopic crust 3he samoan 4he geochemical melt
9	sediment ice erosion rates 10be river sea glacial erosion_rates precipitation	ice ice_sheet sheet antarctic greenland glacial ka retreat antarctic_ice holocene

50³⁴. The second parameter is early exaggeration, which determines the compactness of clusters. This optimization method creates empty space between clusters so they can achieve better global and local organization.⁴⁴

To select the map parameters, we used a visual inspection method. The director of ERI, Kelly Caylor, evaluated the topic regions resulting from the t-SNE and UMAP configurations against the benchmark of the previous institutional review report. Based on his familiarity with the institute’s research, the director confirmed that the results from t-SNE with an early exaggeration value of 5 and a perplexity value of 7 were easiest to interpret and supported his reporting needs. Furthermore, Figure 4 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 regions in the t-SNE map. 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, the director deemed t-SNE to be a more compelling technique for reporting purposes.

RESULTS

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

Reading maps of research documents

The maps produced with t-SNE show research documents with similar topics forming regions 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 assigning labels is a pragmatic choice that allows us to reference and interpret the topic clusters. The categorical colormap (<https://colorcet.holoviz.org/>) 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 5, 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.

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”




Figure 4. 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

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.

Whereas the coarse map presents a distant overview of ERI's research topics, the detailed map shown in Figure 6 reveals intricate patterns. The center “population” cluster borders other research areas, including the “species,” “ocean,” and “fisheries” clusters. Another multitopic 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. 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.

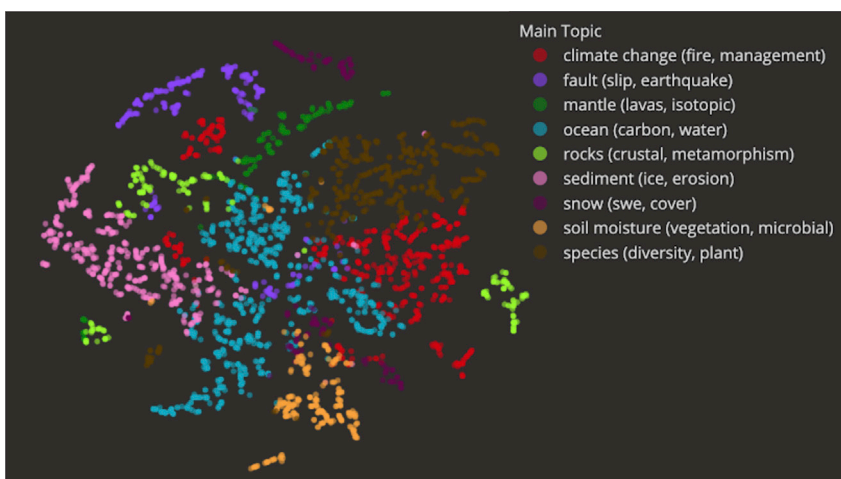


Figure 5. Coarse (9 topic) map of research documents (2009–2019)




Figure 6. Detailed (36 topic) map of research documents (2009–2019)

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). 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, as they represent topics extracted from the text of research documents. The maps also possess emergent qualities, revealing more than the sum of their parts;⁴¹ 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.

Deploying a map dashboard

To distribute and evaluate our maps, we deploy a public-facing dashboard (<https://eri-research-dashboard.herokuapp.com/>) 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 7 shows the “Search” panel, which allows users to filter data by ERI researcher or by academic department and return 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 the next section, [Evaluation](#).

Evaluation

Do the maps we developed support “distant reading of research documents in the context of an institutional review”? To answer this question, we evaluate the maps in two 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.⁵³ We surveyed lead-

ing ERI researchers who determined if and how they think the maps support “reading at a distance.”

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 accomplish-

ment that UCSB’s ORUs must regularly address (<https://www.research.ucsb.edu/organized-research-unit-oru-administration>). They are currently answered using quantitative evidence, for example, numbers of publications by field of research and amounts of funding per researcher. Although these benchmark questions are particular to UCSB, the concerns they address are representative of similar contexts elsewhere:

- **Research quality and significance:** describe the quality and significance of the 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 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

JIFs 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 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.

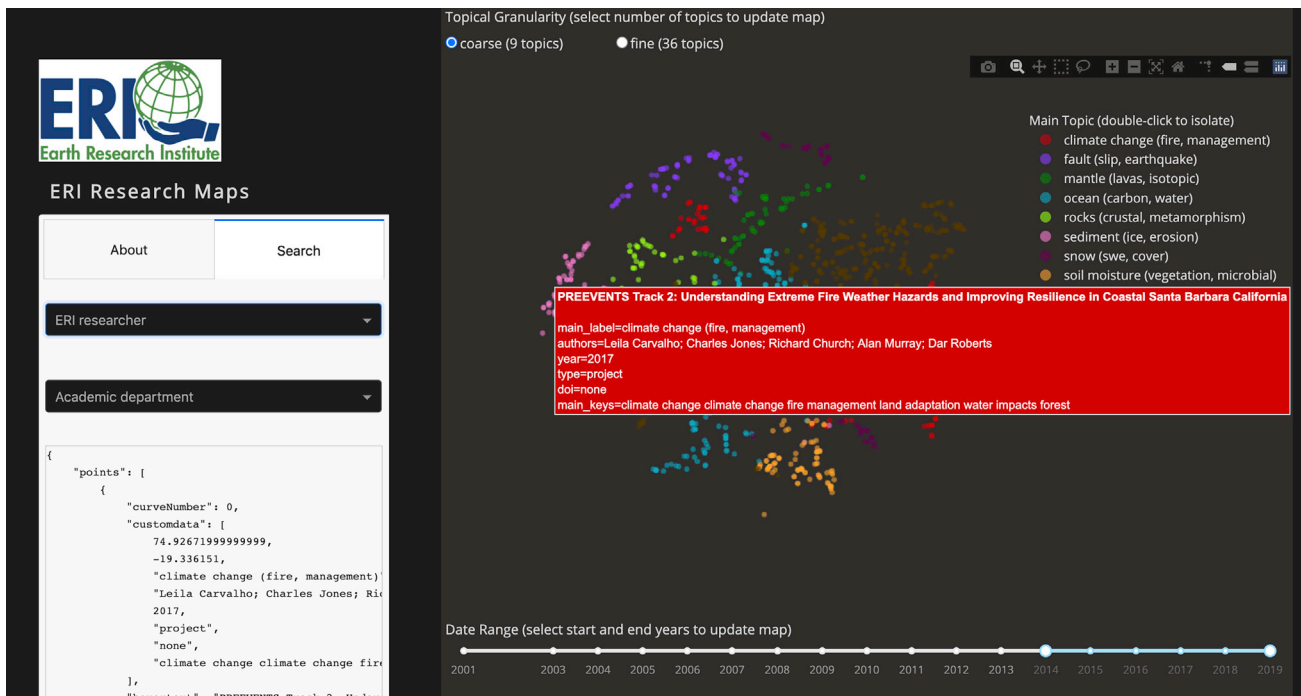


Figure 7. Search panel of the interactive research map dashboard

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). Although 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’s work is more broadly dispersed across “species,” “climate,” “ocean,” “snow,” “sediment,” and “soil moisture.” Although both accomplished researchers work exten-

sively 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 because existing indicators are effective. 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, to reveal changes to topics prioritized over time by funding agencies.

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 relationships 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, at both the coarse and the fine levels of detail. How well do you think these topics represent ERI’s research overall?

- Principal investigator (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. 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. Because 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 easily 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.

DISCUSSION

We applied science mapping, dimensionality reduction, and visualization techniques to uncover research relationships and temporal trends in a corpus of research documents. To confirm the utility of this approach, we surveyed researchers represented within the maps. Our research has immediate benefits for ERI as they prepare for their external review. It facilitates ERI's efforts to identify research trends and areas of expertise, determine the impact of various investments on ERI's productivity, and differentiate scholars' unique areas of contribution. Similar systems would be useful for other research enterprises and funders interested in understanding their own trends and productivity.

One limitation of our approach is that it primarily takes advantage of the thematic dimension of data and treats the spatial and temporal components of the data as secondary. Although temporal views are incorporated in our maps, allowing for document subsetting by time span and event detection, making time a primary dimension could prove valuable. Previous work on semantic signatures has shown that time and space offer two complementary ways to order knowledge.³⁵ Views ordered primarily by time could be thought of as temporalizations, rather than the spatializations we develop, tracking the evolution of topics in the form of graphs from distant reading.⁴¹

Another limitation is that our approach does not take advantage of all of the core concepts of spatial information presented in [Table 1](#). This interpretation suggests technical ways in which our work can be extended. Currently, we embed research documents (*objects*) in a continuous topic space (*field*), which forms regions of research topics. The number of research topics selected (*granularity*) influences the configurations of the topic regions; in our maps, these configurations (detailed and coarse) are independent and are not linked. Time is also handled as a series of annual snapshots over a decade, where change is depicted as the reconfiguration of topic regions between these intervals (*event*).

First, adopting additional topic modeling approaches, such as hierarchical⁵⁴ and dynamic topic models,⁵⁵ would account for multiple levels of thematic and temporal detail within a single model rather than producing separate models at different levels of detail. Second, adopting other visualization methods to depict network information from the documents¹⁷ would convey additional relationships holding among the documents, such as co-authorship or funding patterns. Future modeling and visualization choices should be guided by the priorities of the institute in order to ensure they support the review process.

In terms of evaluation, we are also interested in expanding the survey we conducted to coincide with ERI's external review. This

would give us further insights into how external reviewers who do not have a personal connection to ERI's research interpret and evaluate the research topics. To determine the applicability and maturity of our approach for adoption in a broader context, we would also be interested in surveying researchers or leaders affiliated with similar ORUs. This would allow us to build consensus around strategies for adopting maps of research as robust decision support tools.

At the outset of this article, we proposed that maps of the research "territory" could provide actionable decision support. The maps we have produced give an impression of the underlying thematic structure of the research in the form of research regions 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 ORU's self-assessment and externally as a communication tool describing research trends and developments, which are likely of interest to external reviewers, other research units, and the public.

EXPERIMENTAL PROCEDURES

Resource availability

Lead contact

Sara Lafia is the lead contact of this study and can be reached at slafia@umich.edu.

Materials availability

The code developed for the topic models and data visualizations reported in this article are available in our public Github repository: <https://github.com/saralafia/ERI-maps>. The code developed for the reporting dashboard is available in our public Github repository: <https://github.com/saralafia/ERI-dashboard>.

Data and code availability

The data and code supporting our analysis for the institutional review is available in our public Github repository: <https://github.com/saralafia/ERI-5-year-review>. More information about ERI's review process is available on its website: <https://www.eri.ucsb.edu/2014-external-review>.

ACKNOWLEDGMENTS

We thank the members of ERI's advisory board, along with Daniel R. Montello and James Frew at UCSB, for supporting and guiding this study. We also acknowledge support from an anonymous private grant (<http://spatial.ucsb.edu/research/spatial-discovery>) awarded to the UCSB Center for Spatial Studies and UCSB Library to study challenges and strategies that libraries and researchers face when trying to discover research data on diverse platforms. This material is based upon work supported by the National Science Foundation under grant 1930645.

AUTHOR CONTRIBUTIONS

Conceptualization, S.L., W.K., and K.C.; methodology, S.L.; data curation, S.L. and K.C.; writing – original draft, S.L., W.K., and K.C.; writing – review & editing, S.L. and L.H.; visualization, S.L.; supervision, W.K. and L.H.; funding acquisition, K.C. and L.H.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: August 31, 2020

Revised: December 9, 2020

Accepted: January 20, 2021

Published: February 15, 2021

REFERENCES

- Ioannidis, J.P. (2018). Meta-research: why research on research matters. *PLoS Biol.* 16, e2005468.
- Bawden, D., and Robinson, L. (2015). Informetrics. In *Introduction to Information Science*, D. Bawden and L. Robinson, eds. (Facet Publishing), pp. 165–181.
- Jappe, A., Pithan, D., and Heinze, T. (2018). Does bibliometric research confer legitimacy to research assessment practice? A sociological study of reputational control, 1972–2016. *PLoS One* 13, e0199031.
- Vinkler, P. (2010). Indicators are the essence of scientometrics and bibliometrics. *Scientometrics* 85, 861–866.
- Radicchi, F., Fortunato, S., and Castellano, C. (2008). Universality of citation distributions: toward an objective measure of scientific impact. *Proc. Natl. Acad. Sci. U S A* 105, 17268–17272.
- Moher, D., Naudet, F., Cristea, I.A., Miedema, F., Ioannidis, J.P., and Goodman, S.N. (2018). New principles for assessing scientists. *Issues Sci.* 35, 20–23.
- Larivière, V., and Gingras, Y. (2010). The impact factor's Matthew Effect: a natural experiment in bibliometrics. *J. Am. Soc. Inf. Sci. Technol.* 61, 424–427.
- Bergstrom, C.T., West, J.D., and Wiseman, M.A. (2008). The Eigenfactor metrics. *J. Neurosci.* 28, 11433–11434.
- Guerrero-Bote, V.P., and Moya-Aneón, F. (2012). A further step forward in measuring journals' scientific prestige: the SJR2 indicator. *J. Informetrics* 6, 674–688.
- Priem, J., Groth, P., and Taraborelli, D. (2012). The altmetrics collection. *PLoS One* 7, e48753.
- Piwowar, H. (2013). Value all research products. *Nature* 493, 159.
- 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, 1267–1274.
- Börner, K., Chen, C., and Boyack, K.W. (2003). Visualizing knowledge domains. *Annu. Rev. Inf. Sci. Technol.* 37, 179–255.
- Kuhn, W. (1996). Handling data spatially: spatializing user interfaces. In *Advances in GIS Research II: Proceedings of the 7th International Symposium on Spatial Data Handling* (p. 13B). volume 2.
- 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). IEEE.
- Andrienko, G., Andrienko, N., Jankowski, P., Keim, D., Kraak, M.-J., MacEachren, A., and Wrobel, S. (2007). Geovisual analytics for spatial decision support: setting the research agenda. *Int. J. Geograph. Inf. Sci.* 21, 839–857.
- Chen, C. (2017). Science mapping: a systematic review of the literature. *J. Data Inf. Sci.* 2, 1–40.
- Chen, C. (1999). Visualising semantic spaces and author co-citation networks in digital libraries. *Inf. Process. Manage.* 35, 401–420.
- Chen, C. (2006). Citespace ii: Detecting and visualizing emerging trends and transient patterns in scientific literature. *J. Am. Soc. Inf. Technol.* 57, 359–377.
- Callon, M., Courtial, J.-P., Turner, W.A., and Bauin, S. (1983). From translations to problematic networks: an introduction to co-word analysis. *Information (Int. Soc. Sci. Council)* 22, 191–235.
- Skupin, A., and Fabrikant, S.I. (2003). Spatialization methods: a cartographic research agenda for non-geographic information visualization. *Cartography Geogr. Inf. Sci.* 30, 99–119.
- Choo, J., Lee, C., Reddy, C.K., and Park, H. (2013). Utopian: user-driven topic modeling based on interactive nonnegative matrix factorization. *IEEE Trans. Visual. Comput. Graph.* 19, 1992–2001.
- Chuang, J., Manning, C.D., and Heer, J. (2012). Termite: Visualization techniques for assessing textual topic models. In *Proceedings of the*

- International Working Conference on Advanced Visual Interfaces (pp. 74–77).
24. Sievert, C., and Shirley, K. (2014). LDAvis: A method for visualizing and interpreting topics. In Proceedings of the Workshop on Interactive Language Learning, Visualization, and Interfaces (pp. 63–70).
 25. Kim, M., Kang, K., Park, D., Choo, J., and Elmqvist, N. (2016). TopicLens: efficient multi-level visual topic exploration of large-scale document collections. *IEEE Trans. Visual. Comput. Graph.* *23*, 151–160.
 26. Yang, Y., Yao, Q., and Qu, H. (2017). Vistic: a visual analytics system for making sense of large document collections using hierarchical topic modeling. *Vis. Inform.* *1*, 40–47.
 27. 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.
 28. Lafia, S., Last, C., and 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.
 29. Boyd-Graber, J., Mimno, D., and Newman, D. (2014). Care and feeding of topic models: problems, diagnostics, and improvements. In *Handbook of Mixed Membership Models and Their Applications* (CRC Press), pp. 225–255.
 30. Blei, D.M., Ng, A.Y., and Jordan, M.I. (2003). Latent Dirichlet allocation. *J. Machine Learn. Res.* *3*, 993–1022.
 31. Salton, G., Wong, A., and Yang, C. (1975). A vector space model for automatic indexing. *Commun. ACM* *18*, 391–398.
 32. Lee, D., and Seung, H.S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature* *401*, 788–791.
 33. Greene, D., O’Callaghan, D., and 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.
 34. Röder, M., Both, A., and Hinneburg, A. (2015). Exploring the space of topic coherence measures. In Proceedings of the Eighth ACM International Conference on Web Search and Data Mining (pp. 399–408).
 35. Janowicz, K. (2010). The role of space and time for knowledge organization on the semantic web. *Semantic Web* *1*, 25–32.
 36. Arias, S., and Warf, B. (2009). Introduction: the reinsertion of space into the social sciences and humanities. In *The Spatial Turn: Interdisciplinary Perspectives*, B. Warf and S. Arias, eds. (Routledge), pp. 1–10.
 37. Gärdenfors, P. (2000). Semantics. In *Conceptual Spaces: The Geometry of Thought*, P. Gärdenfors and F. Zenker, eds. (MIT Press), pp. 151–201.
 38. Tobler, W.R. (1970). A computer movie simulating urban growth in the detroit region. *Econ. Geogr.* *46*, 234–240.
 39. Fabrikant, S.I., Ruocco, M., Middleton, R., Montello, D.R., and Jörgensen, C. (2002). The first law of cognitive geography: distance and similarity in semantic space. *Proc. GIScience 2002*, 31–33.
 40. Underwood, T. (2016). Distant reading and recent intellectual history. *Debates Digital Human.* *2016*, 530–533.
 41. Moretti, F. (2005). *Graphs, Maps, Trees: Abstract Models for a Literary History* (Verso).
 42. Kuhn, W. (2012). Core concepts of spatial information for transdisciplinary research. *Int. J. Geograph. Inf. Sci.* *26*, 2267–2276.
 43. Allen, C., Hervey, T., Lafia, S., Phillips, D.W., Vahedi, B., and Kuhn, W. (2016). Exploring the Notion of Spatial Lenses. In The Annual International Conference on Geographic Information Science (pp. 259–274). Springer.
 44. Maaten, L.v. d., and Hinton, G. (2008). Visualizing data using t-SNE. *J. Machine Learn. Res.* *9*, 2579–2605.
 45. McInnes, L., Healy, J., Saul, N., and Großberger, L. (2018). UMAP: uniform manifold approximation and projection. *J. Open Source Softw.* *3*, 861.
 46. Bird, S., Klein, E., and Loper, E. (2009). *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit* (O’Reilly Media, Inc.).
 47. Miller, G.A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychol. Rev.* *63*, 81.
 48. Mimno, D., Wallach, H.M., Talley, E., Leenders, M., and McCallum, A. (2011). Optimizing Semantic Coherence in Topic Models. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (pp. 262–272). Association for Computational Linguistics.
 49. O’Callaghan, D., Greene, D., Carthy, J., and Cunningham, P. (2015). An analysis of the coherence of descriptors in topic modeling. *Expert Syst. Appl.* *42*, 5645–5657.
 50. Arora, S., Ge, R., Halpern, Y., Mimno, D., Moitra, A., Sontag, D., Wu, Y., and Zhu, M. (2013). A Practical Algorithm for Topic Modeling with Provable Guarantees. In International Conference on Machine Learning (pp. 280–288).
 51. Chang, J., Gerrish, S., Wang, C., Boyd-Graber, J., and Blei, D. (2009). Reading tea leaves: how humans interpret topic models. *Adv. Neural Inf. Process. Syst.* *22*, 288–296.
 52. Wattenberg, M., Viégas, F., and Johnson, I. (2016). How to use t-SNE effectively. *Distill* *1*, e2, <https://doi.org/10.23915/distill.00002>.
 53. Solomon, P. (2002). Discovering information in context. *Annu. Rev. Inf. Sci. Technol.* *36*, 229–264.
 54. Griffiths, T., Jordan, M., Tenenbaum, J., and Blei, D. (2003). Hierarchical topic models and the nested Chinese restaurant process. *Adv. Neural Inf. Process. Syst.* *16*, 17–24.
 55. Blei, D.M., and Lafferty, J.D. (2006). Dynamic topic models. In Proceedings of the 23rd International Conference on Machine Learning (pp. 113–120).