Hidden Stories: Topic Modeling in Hydrologic Literature

Mashrekur Rahman[1], Jonathan M. Frame[1], Jimmy Lin[2], Grey S. Nearing[1,3]

[1]Department of Geological Sciences, University of Alabama
[2]David R. Cheriton School of Computer Science, University of Waterloo
[3]Upstream Tech Public Benefit Company, Natel Energy

Key Points:

- Topic Modeling is a form of unsupervised machine learning for Natural Language Processing (NLP)
- Topic Modeling can provide a high-level overview of topics and trends in hydrology literature
- This is a first step toward building a tool to help researchers navigate and synthesize a growing body of literature

Keywords:

- Hydrology Literature
- Science Communication
- Machine Learning
- Unsupervised Learning
- Natural Language Processing
- Topic Modeling

Corresponding author: Mashrekur Rahman, `mrahman18@crimson.ua.edu`

Abstract

Hydrologic research generates large volumes of peer-reviewed literature across a number of evolving sub-topics. It's becoming increasingly difficult for scientists and practitioners to synthesize this full body of literature. This study explores topic modeling with Latent Dirichlet Allocation (LDA) as a form of unsupervised learning applied to 42,154 article-abstracts from six high-impact (Impact Factor > 0.9) journals (Water Resources Research WRR, Hydrology and Earth System Sciences $HESS$, Journal of Hydrology JH, Hydrological Processes HP, Hydrological Sciences Journal HSJ, Journal of Hydrometeorology JHM) to provide a high-level contextual analyses of hydrologic science literature since 1991. We used a hybrid quantitative/qualitative approach to label a number of broad topics in this body of literature, and used these labeled topics to analyze topic trends, inter-topic relationships, and journal diversity. As an example of what we can learn from this type of analysis, results showed that data-driven research topics are gaining in popularity while some subsurface related topics lose popularity within our journal set and time period. While no journal in our sample was completely homogeneous, JHM and WRR exhibited the most notable preferences for certain topics over others. The methods and outcomes of this paper are potentially beneficial to scientists and researchers who aim to gain a contextual understanding of the existing state of hydrologic science literature. In the long term, we see topic modeling as a tool to help increase the efficiency of literature reviews, science communication, and science-informed policy and decision making.

1 Introduction

Hydrologic research generates large volumes of peer-reviewed literature (Figure 1) across a plethora of evolving topics and sub-topics. An increasing amount of effort is required for researchers and practitioners to synthesize this literature, and to track the state-of-the science in any particular topic area within the discipline. The dynamic nature of hydrology and the society (Montanari et al., 2013) means that the hydrologic community is increasingly required to advocate and advise sustainable development through water resources awareness and management (Rahaman & Varis, 2005). As a result of challenges like these, science communication is evolving rapidly, and there is a growing need for more sophisticated "scientific" ways to synthesize and communicate research findings (Nisbet & Scheufele, 2009).

Recent advances in computational linguistics, machine learning, and a variety of application-ready toolboxes for Natural Language Processing (NLP) can help facilitate analyses of vast electronic corpuses for a variety of objectives (Cambria & White, 2014). Information retrieval, text categorization, and other text mining techniques based on machine learning have been gaining popularity since the 1990s (Sebastiani, 2002). The ability to quickly synthesize large volumes of electronic text can help offer windows into trending topics and help scientists identify related efforts and research developments in a body of literature.

Topic modeling is one popular NLP technique that uses statistical algorithms to extract semantic information from a collection of texts in the form of thematic classes (Jiang, Qiang, & Lin, 2016). Topic models can be applied to massive collections of documents (Blei, 2012) and have been used to recommend scientific articles based on both content and user ratings (C. Wang & Blei, 2011). Topic modeling has also been used to cluster scientific documents (Yau, Porter, Newman, & Suominen, 2014), improve bibliographic search (Jardine & Teufel, 2014; Paul & Girju, 2009; Pham, Do, & Ta, 2018; Shu, Long, & Meng, 2009; Tang, Jin, & Zhang, 2008), and for a variety of other applications such as statistical modeling of the biomedical corpora (Blei, Franks, Jordan, & Mian, 2006), bibliometric exploration of hydropower research(Jiang et al., 2016), analyzing research trends in personal information privacy (Choi, Lee, & Sohn, 2017), meta-
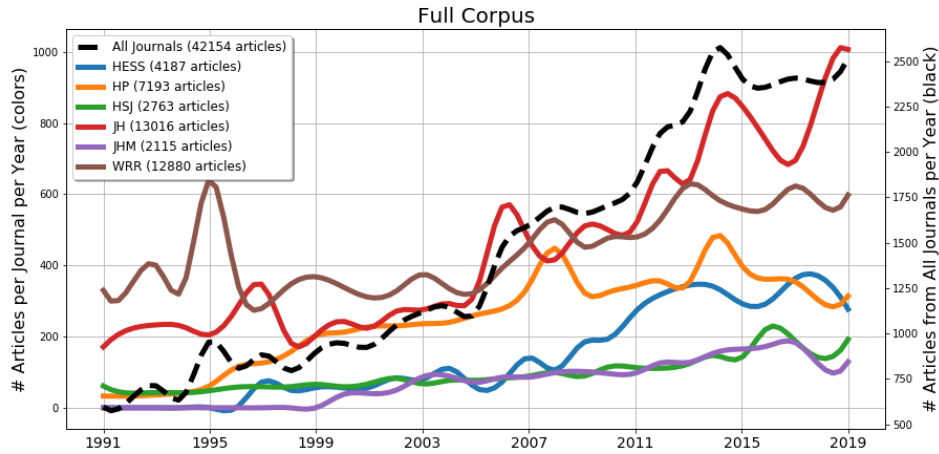
Figure 1. Number of articles published per year between 1991 and 2019 in six major hydrology journals (Source: Web of Science, Scimago Journal and Country Rank)

reviewing cloud computing literature (Upreti, Asatiani, & Malo, 2016), literature review of social science (Li & Liu, 2018), the technology-acceptance model (Mortenson & Vidgen, 2016), discovering themes and trends in transportation research (S. Sun, Luo, & Chen, 2017), Knowledge Management literature (Jussila et al., 2017), exploring the history of cognition (Priva & Austerweil, 2015) and exploring topic divergence and similarities in scientific conferences (Hall, Jurafsky, & Manning, 2008). Topic modeling algorithms allow for exploration of a broad range of data including non-English corpuses (Riddell, 2014), software engineering data (X. Sun et al., 2016), and even historical newspapers (Yang, Torget, & Mihalcea, 2011). Given the incremental popularity of topic models and its versatile applicability in a wide range of applications, we wish to explore the potential for topic modeling to aid bibliometric exploration of peer-reviewed hydrologic science literature.

Peer-reviewed abstracts offer snapshots of the historical and current trends and developments in both theoretical and applied research. Article abstracts are perceived as concise representations of full-texts and are used for bibliometric analyses (Gatti, Brooks, & Nurre, 2015; Griffiths & Steyvers, 2004). Although techniques such as *scientometrics* (Mingers & Leydesdorff, 2015) have been traditionally used for ranking articles and authors based on citation data, topic modeling allows for contextual understanding of particular scientific domains and disciplines. Hydrologic research articles encompass a wide range of research topics including flood prediction, climate change etc., all of which are consequential to global socioeconomic well-being. Water managers and policy makers, who ideally make decisions about water resources based on state of the knowledge of hydrologic science, depend on data, tools and predictions provided by scientists and practitioners in this field. It is therefore imperative for at least many stakeholders of hydrologic research to understand topics and trends in this discipline without having to read thousands of research articles.

In this study we applied topic modeling using unsupervised learning with Latent Dirichlet Allocation (LDA) on 42,154 article-abstracts from six high-impact (Impact Factor > 0.9) journals in hydrology (Water Resources Research WRR, Hydrology and Earth System Sciences $HESS$, Journal of Hydrology JH, Hydrological Processes HP, Hydrological Sciences Journal HSJ, Journal of Hydrometeorology JHM). LDA identifies groups

of words commonly found together, and produces relationships between these word clusters (topics) and individual documents. Using these topic-word distributions, we then relied on a hybrid quantitative/qualitative approach to label a number of broad topics. We analyzed how these topics relate to each other and change over time and between journals.

As an example of what can be learned from this type of analysis, results show that data-driven research topics have been gaining in popularity in recent years, while some subsurface related topics are declining in popularity within our journal sample set (although this may be due to the introduction of new groundwater journals during our study's time period). Significant statistical relationships were observed between topics - for example, research on anthropogenic interventions and effects is significantly correlated with research on climate change, hydromorphology, flooding, water quality, and extreme events.

We further analyzed topic distributions in individual journals to help understand the diversity of topics within journals and uniqueness of topics between journals. While no journal in our sample is completely homogeneous in terms of the topics of papers published, JHM and WRR exhibited the most notable preference for certain topics. A majority of the journals in our corpus appear to be broadening their scope over time.

The methods and outcomes of this type of literature analysis are potentially beneficial to scientists and researchers who aim to gain a high-level or contextual understanding of the existing state of hydrologic science. In the long term, we see NLP, and topic modeling in particular, as potentially useful for helping scientists navigate growing bodies of peer-review literature, and to help increase the efficiency of science communication and science-informed policy and decision making outside of the academic discipline itself.

2 Methods

Table 1 lists all notations used throughout this paper, including variables and indices related to the model and corpus.

2.1 Data acquisition and preprocessing

2.1.1 Repository of article-abstracts

We chose journals for this analysis based on Impact Factors coupled with our subjective perception of the journal's role within the hydrologic science community. Our corpus consists of the abstracts of all peer-reviewed articles published in six hydrologic journals with an Impact Factor (IF) of greater than 0.9, between 1991 and 2019 according to SciMago's Web of Science. The list of journals, journal abbreviations that we will use throughout the rest of this article, corresponding IF, years of available data, and total number of abstracts are listed in Table 2. Article-abstracts were acquired from SciMago's Web of Science in the form of bib files.

The corpus was restricted to the six journals listed in Table 2 because we previously performed the entire analysis reported in this paper on all journals with an Impact Factor greater than 0.9 in Scimago's 'Water Science and Technology' classification, but the results were too diverse for a meaningful analysis. We report here an analysis only using a sub-selection of journals, and specifically focus on multi-disciplinary hydrology journals (i.e., we did not include journals focused primarily on groundwater, regional studies, marine science, desalinization, cryosphere, etc.).

Table 1. List of notations for indices, parameters and variables

Notation	Meaning
Corpus Parameters	
M	Number of documents
N_d	Number of words in document d
t_d	Year of publication of document d
LDA Model Components	
K	Number of topics
K_{opt}	Optimal number of topics
α	Parameters of a Dirichlet prior on on the per-document topic distribution
β	Parameters of a Dirichlet prior on the per-topic word distribution
μ_d	Distribution of topics over document d
z	list of K topics
$\mathbf{z_d}$	Per-word topic vector for document d
$\mathbf{w_d}$	Word collection in document d
Derived Distributions	
μ_{kj}	Weight of a particular topic k over all documnets in journal j
μ_{kt}	Average weight for topic k over all documents at time t
$\hat{\mu_k}$	Mean weight of topic k over all documents
μ_{kj}^t	Weight of topic k in journal j at time t
μ_m	Topic distribution over entire corpus of M documents
Derived Metrics & Functions	
p	LDA perplexity score
c	LDA coherence score
JSD	Jensen-Shannon Divergence
KLD	Kullback-Leibler Divergence
I	Indicator function
$R_{k,j}$	Correlation coefficient between topics k and j
H_j	Topic entropy (diversity) of journal j
$d_{js}(j,i)$	Jensen-Shannon distance between journals j and i
$d_{js}^d(j)$	Jensen-Shannon distance of journal j from entire corpus
$d_{js}^t(j)$	Jensen-Shannon distance between journals j and i

Table 2. Repository of article-abstracts

Journal Name	Abbreviation	IF	Years Available	Total Abstracts
Journal of Hydrometeorology	JHM	2.410	2000-2019	2115
Water Resources Research	WRR	2.135	1991-2019	12880
Hydrology and Earth System Sciences	HESS	2.134	1997-2019	4187
Journal of Hydrology	JH	1.830	1991-2019	13016
Hydrological Processes	HP	1.417	1991-2019	7193
Hydrological Sciences Journal	HSJ	0.910	1991-2019	2736

2.1.2 Preprocessing the corpus

Performance of topic modeling is influenced by the quality of input training data. Data preprocessing in text mining involves converting acquired data into canonical format for efficacious feature extraction (Feldman, Sanger, et al., 2007). In addition, a portion of the raw data from any corpus is extraneous and may not add value to the analysis - as such, training data requires appropriate preprocessing, as described presently.

We used separate temporally-segregated dataframes for abstracts from each journal. All sets of data were processed through identical multi-layered cleaning routines. We initiated the process by first creating a dataframe of all article-abstracts and their corresponding metadata. We then filtered nonsensical elements such as stopwords, punctuation, and symbols, in addition to subjective manual identification and removal of unwanted elements.

In the next step, we formed bi-grams and tri-grams, and then segmented the texts by tokenizing with whitespaces as word boundaries, followed by lemmatization to normalize into a canonical format. The resultant output was converted into a term frequency-inverse document frequency (TF-IDF) format for ingesting by the LDA model implemented in *Gensim* - a Python library for NLP (Řehřek & Sojka, 2011).

2.2 Latent Dirichlet Allocation

LDA builds on more traditional Latent Semantic Analysis (Landauer, Foltz, & Laham, 1998), and captures the intuition that text documents exhibit multiple topics in different proportions. Documents are represented as mixtures of topics (per-document topic distributions) and each topic is characterized by a distribution over words (per-topic word distributions).

We can build an intuition of this model as follows. It is assumed that the per-document topic distributions of all documents in a corpus share a common Dirichlet prior parameterized by α, and that the per-topic word distributions also share a (different) common Dirichlet prior parameterized by β. The distribution over a particular word w in a document d with topic distribution μ_d can be understood as (Blei, Ng, & Jordan, 2003):

$$p(w|\mu_d, \beta) = \sum_{k=1}^{K} p(z_k|\mu_d)p(w|z_k, \beta), \tag{1}$$

where z_k is a particular topic from K total topics. Treating the per-document topic distribution as latent and integrating over all N_d words in each document d and over all M documents in corpus D gives:

$$p(D|\alpha, \beta) = \sum_{d=1}^{M} \int_{\mu_d} p(\mu_d|\alpha) \left(\prod_{n=1}^{N_d} p(w_{dn}|\mu_d, \beta) \right) d\mu_d \tag{2}$$

The above is an intuition only. In actuality, LDA assumes a generating model (i.e., a model of how the corpus was produced) that samples each μ_d once for each word in a corpus, which means that each document contains a mixture of topics, which is why each document has its own per-document topic distribution. This means that each document d can be associated with an N_d vector of topics, $\mathbf{z_d}$, - one topic assignment (out of K total topics) for each word in the document. This generating model is described in more detail by Blei et al. (2003) and others.

Training the LDA model involves estimating the per-document topic distributions, μ_d, and the per-document topic vectors, $\mathbf{z_d}$, given the words in a document, $\mathbf{w_d}$, and the Dirichlet priori parameters: $p(\mu_d, \mathbf{z_d}|\mathbf{w_d}, \alpha, \beta)$. This can be done using a variety of methods, including Gibbs Sampling (Griffiths & Steyvers, 2004), variational expectation-maximization

191 (VEM) (Blei et al., 2003), and others. Model overfitting is generally not a major issue
192 for unsupervised learning with LDA, which is a Bayesian model.

193 Here, we used an LDA implementation in the Python *Gensim* package with VEM.
194 We trained our models with the number of passes set to 3000 and chunksize (number
195 of documents in a batch) set to 100. We used a parallel implementation of LDA in *Gensim*
196 to train individual models with topic sizes ranging from $K = 2$ to $K = 40$; each model
197 trained using 40 shared-memory cores on a single node of a high performance cluster.
198 Using these settings it took on the order of a few hours to train a single model: between
199 3-15 hours per model on our particular machine, depending on K.

2.3 Choosing an optimal number of topics

201 Ideally we want to maximize the number of topics to increase their variety and "depth"
202 in terms of how the model partitions the article-abstracts. In practice, however, a num-
203 ber of topics, K, above some (unknown) optimal number of topics, K_{opt}, increases the
204 occurrence of common words among different topics, resulting in compromised quality
205 of topics (Lu, Mei, & Zhai, 2011). We therefore adopted a hybrid quantitative/qualitative
206 approach for deciding the optimal number of topics, K_{opt}.

2.3.1 Data-driven approach to choose optimal number of topics

208 We used a combination of perplexity p and coherence c scores as metrics to eval-
209 uate model performance over a range of numbers of topics Perplexity is a popular met-
210 ric for evaluating language models (Chen, Beeferman, & Rosenfeld, 1998). Perplexity is
211 an information theory metric that measures something like how surprised the model might
212 be on the introduction of new data (Zhao et al., 2015). Formally defined by Blei et al.
213 (2003), perplexity for a collection of M documents is:

$$p = exp \left\{ -\frac{\sum_{d=1}^{M} log \, p(w_d)}{\sum_{d=1}^{M} N_d} \right\} \qquad (3)$$

215 Perplexity is a decreasing function of the probability assigned to each per-document word
216 distribution. Lower perplexity indicates a better model.

217 Topic coherence c is a measure of similarity in semantics between the high prob-
218 ability words in a certain topic. We use $Gensim's$ built-in topic coherence model, which
219 is an implementation of the method described by (Röder, Both, & Hinneburg, 2015). Cal-
220 culating topic coherence is a four-stage process involving segmentation of word subsets,
221 probability calculation, confirmation measure, and aggregation.

222 Figure 2 (adapted from Röder et al., 2015) illustrates these four steps. t represents
223 an input collection of words, and the first stage creates a set of different kinds of seg-
224 mentation of words S from t, since coherence measures the fitting together of words or
225 a set of words. Secondly, probabilities of occurrence of words P are calculated based on
226 reference corpus. Confirmation measure ingests both P and S to yield the agreements
227 φ of pairs of S. In the final step, the aforementioned scores are aggregated to compute
228 coherence c.

229 We trained LDA models using identical hyperparameters for a range of topics num-
230 bers from $K = 2$ to $K = 40$, logging the coherence c and perplexity p scores for each
231 K. The resulting scores are plotted in Figure 3. To determine K_{opt}, we considered a range
232 of number of topics K for which coherence c peaks, accompanied by a decreasing trend
233 for perplexity p plot - i.e., $K = 20$ to $K = 32$.

Figure 2. Illustration of the four stages of the unified topic coherence framework. In stage 1, input words t are segmented into smaller sets S. Probabilities of occurrence P of words are calculated based on the reference corpus in the second stage. In the third stage, P and S are ingested to measure φ between pairs of words S. Coherence c is calculated in the final step.
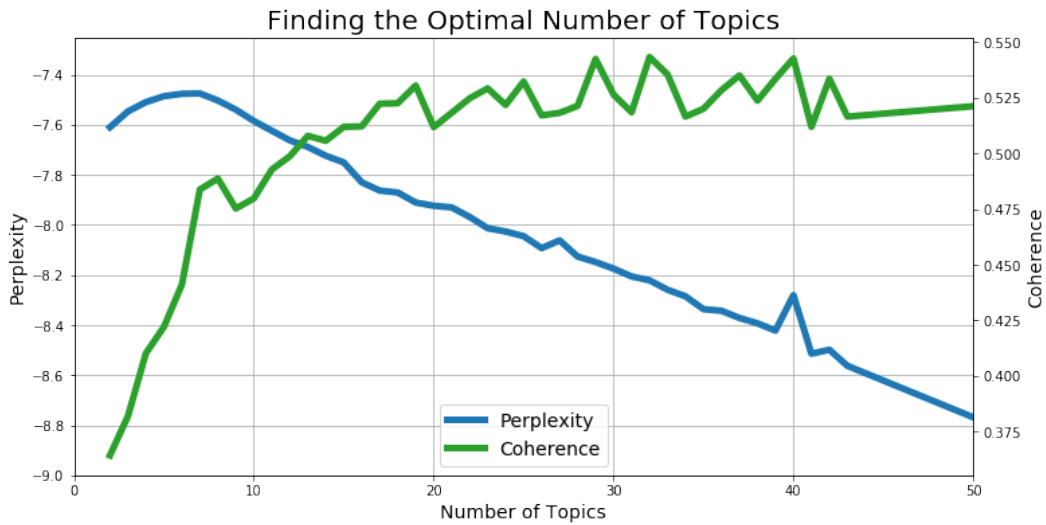

Figure 3. Variation of topic coherence c and perplexity p based on LDA models trained for a range of topic numbers ($K = 2$ to $K = 40$). Lower perplexity and higher coherence indicate a better model. These values guide our subjective analysis for choosing K_{opt}

2.3.2 Qualitative approach to choosing optimal number of topics

Qualitative perception of topics allows for data-driven evaluation metrics to be supported by manual validation. We subjectively assessed the quality of topics for various K, looking for increasing or decreasing occurrence of similar words within certain topics and backtracking into the dataframe to observe the titles of documents associated with each topic. We drew on our prior education and experience in hydrology to make these assessments, and also solicited input from several other professional hydrologists. Based on this and the aforementioned objective indicators, we chose $K_{opt} = 25$. This is where the coherence score had an inflection point (i.e., started to level off around its maximum value), and subjectively the topics at $K_{opt} = 25$ did not contain a significant amount of redundancy.

There was consistency between individual topics found with different values of K as K increased. Figure 4 is a partial illustration of the topic evolution with increasing

topic number. All of the topic names shown on this chart were chosen by researchers based on looking at the keywords that the model associated with each topic, as well as the 100 abstract titles that had the strongest association with each topic. With a low number of topics, $K = 2$, the model partitioned the dataset into categories that were (vaguely) related to surface hydrology and terrestrial processes vs. subsurface and hydraulics. With further increase in number of topics - e.g., $K = 5$ - the surface hydrology topic was partitioned into topics related primarily to climate change, terrestrial processes, and modeling, while the subsurface topic split into topics defined by keywords related to hydraulics and groundwater, with some papers splitting to join the more refined modeling and terrestrial processes topics. The LDA model partitioning became more refined with further increases in the number of topics, and the resulting topics became clearer and more well-defined. Increased topic refinement caused separation and merger of different closely related topics. As an example, at $K = 10$, a single modeling related topic split into hydraulic modeling and catchment modeling. Hydraulic modeling split further and combined with a flow and transport topic to form a topic based on flow and transport modeling. Simultaneously, catchment modeling split further and merged with specific subtopics such as climate change, water management and statistical hydrology. It's important to understand that especially at small topic numbers, these topics are fairly vague and the topic names that we assigned are indicators of broad themes.

3 Analysis Methods

This section describes the methods we used to analyze document-topic and topic-word distributions from the LDA model, as well as for computing topic trends, distributions over time, inter-topic correlations, and distributions of topics within journals.

3.1 Temporal distribution of topics

There are multiple methods of analyzing temporal trends and distribution of topics. Griffiths and Steyvers (2004) applied a disjointed time-blind topic model and rearranged documents according to their publication dates. Blei and Lafferty (2006) developed a sequential topic modeling approach that learns time-dynamic parameters for the document-topic and topic-word distributions constrained by linear filtering theory. X. Wang and McCallum (2006) introduced a non-Markov joint modeling framework where topics are associated with a continuous distribution over document timestamps. We adopted Griffiths and Steyvers's (2004) approach of time-unaware topic modeling and post-hoc aggregation of results according to their timestamps. We calculated temporal topic distributions for a given year μ_{kt} as the proportion of all topic weights over all papers from a given year, t:

$$\mu_{kt} = \frac{\sum_{d=1}^{M} \mu_{dk} \times I(t_d - t)}{\sum_{d=1}^{M} I(t_d - t)}. \tag{4}$$

μ_{dk} represents the weight for topic k assigned to document d, t_d is the year in which document d was published, and I is an indicator function such that $I(0) = 1$ and $I(x) = 0$ for $x \neq 0$. Henceforth, I will carry the same meaning.

3.2 Inter-topic correlations

We explored relationships between topics by looking at the correlation coefficient $R_{k,j}$ between the topic weights over the whole corpus M for each pair of topics:

$$R_{k,j} = \frac{\sum_{d=1}^{M} (\mu_{dk} - \hat{\mu_k})(\mu_{dj} - \hat{\mu_j})}{\sqrt{\sum_{d=1}^{M} (\mu_{dk} - \hat{\mu_k})^2} \sqrt{\sum_{d=1}^{M} (\mu_{dj} - \hat{\mu_j})^2}}, \tag{5}$$

Figure 4. Evolution of topics with increasing number of topics K. Lines in the Sankey diagram represent papers shared by each topic (at different topic numbers), where each paper is weighted by the relative proportion of inclusion in the sending topic (i.e., the topic at the smaller number of topics).

where μ_{dk} is the weight for topic k assigned to document d, and $\hat{\mu}_k$ is the mean weight for topic k assigned over all documents in the corpus. All correlations were tested for significance at $\alpha = 0.1$, and we report only correlations with significance at this level.

3.3 Journal diversity

the K-nomial distribution over topics in a particular journal j, μ_j, is:

$$\mu_{kj} = \frac{\sum_{d=1}^{M} \mu_{dk} \times I(j_d - j)}{\sum_{l=1}^{K} \sum_{d=1}^{M} \mu_{dk} \times I(j_d - j)}, \tag{6}$$

where μ_{kj} is the relative popularity of a particular topic in a particular journal as a fraction of popularity of all topics in the journal.

Entropy is a measure of the uncertainty in a probability distribution (Shannon, 1948). We calculated the total entropy of each μ_j, H_j, as a measure of the diversity of the per-journal topic distributions:

$$H_j = -\sum_{k=1}^{K} (\mu_{kj} log(\mu_{kj})), \tag{7}$$

The popularity of a particular topic in a particular journal for a particular year, μ_{kj}^t is a fraction of the popularity of all topics in a journal for a particular year:

$$\mu_{kj}^t = \frac{\sum_{d=1}^{M} \mu_{dk} \times I(|j_d - j| + |t_d - t|)}{\sum_{l=1}^{K} \sum_{d=1}^{M} \mu_{dk} \times I(|j_d - j| + |t_d - t|)}, \tag{8}$$

A timeseries of topic popularity in each journal, μ_{kj}^t, allows us to quantify the evolving relationship between topic distributions in different journals over time. To do this, we consider journal "uniqueness" as a measure of distance of a the topic distribution in a particular journal from the topic distribution over the entire corpus of all journals. This distance is quantifiable by Jensen Shannon distance d_{js} (Endres & Schindelin, 2003), a close relative of Jensen-Shannon divergence JSD (Osterreicher & Vajda, 2003). Jensen-Shannon divergence is a class of information-theoretic divergence based on Shannon entropy (Lin, 1991). It measures similarity between two probability distributions, where JSD=0 represents identical distributions. JSD is also a symmetrized and smoothed version of Kullback-Leibler divergence KLD.

For journal j, μ_j is the overall topic distribution across all articles in the journal. Considering the topic distributions from two journals, μ_a and μ_b, the JSD is:

$$JSD(\mu_a, \mu_b) = \frac{1}{2}KLD(\mu_a, \mu^*) + \frac{1}{2}KLD(\mu^*, \mu_b), \tag{9}$$

where

$$KLD(\mu, \mu^*) = \sum_{k=1}^{k} \mu_k log \frac{\mu_k}{\mu_k^*} \tag{10}$$

is the Kullback-Leibler divergence between the topic distributions μ and μ^*, and $\mu^* = \frac{1}{2}(\mu_a + \mu_b)$.

Hall et al. (2008) and X. Sun et al. (2016) explored the space of similarity and differences between journals with hierarchical clustering. However, X. Sun et al. (2016) used Jensen-Shannon distance d_{js} instead of JSD for this purpose. We also used Jensen-Shannon distance d_{js} as the metric for understanding the relationship dynamics between the different journals and demonstrate their divergence according to their corresponding popularity of topics:

$$d_{js}(i, j) = \sqrt{JSD(\mu_i, \mu_j)} \tag{11}$$

We estimated journal "Uniqueness" as the Jensen-Shannon distance d_{js} of each journal from the entire corpus:

$$d_{js}^d(j) = \sqrt{JSD(\mu_j, \mu_m)}, \tag{12}$$

where μ_m is the topic distribution over entire corpus of M abstracts. Temporal variation of this uniqueness was estimated by calculating the Jensen-Shannon distance on a per-year basis for each journal, d_{js}^t.

4 Results and Analysis

4.1 Naming the topics

The LDA model outputs a certain number of words in each topic and assigns weights to each of those words based on their likelihood of appearance within a particular topic. The topics from our $K = 25$ LDA model correspond strongly with research areas within hydrology. We identified and named the $K = 25$ topics by first looking at the topic-word distributions (the set of words most likely to appear within a particular topic), and the per-document topic distributions (from the titles of articles most closely associated with each topic). Here again, we draw on our prior training and education in hydrology. We reinforced our choices of names for these topics with an informal survey sent to four reputable hydrologists outside of our research group.

Figure 5 illustrates the topic-word distributions in the form of wordclouds. Again, the topic labels in this figure were assigned by the researchers using the procedure described above.

4.2 Temporal distribution of topics

The popularity of each identified topic changes with time, and these trends are also shown in Figure 5. Some topics, such as "Precipitation Variability Extremes", "Precipitation Observation", "Water Management", "Floods", "Climate Change", "Systems Hydrology" and "Modeling Forecasting" demonstrate a clear rising trend in popularity. These rising trends might be attributed to researchers increasingly leveraging the availability and accessibility of hydrology related data, both in terms of breadth and depth. Topics such as "Hydrogeochemistry", "Soil Moisture", "Statistical Hydrology", "Rainfall-Runoff", "Water Quality", "Channel Flow", "Sediment Erosion", "Subsurface Flow Transport", "Scaling Spatial Variability", "Land Surface Fluxes", "Hydrogeology", "Land Cover" and "Groundwater" have demonstrated explicit decreasing temporal trends. Such behaviors might be attributed to a multitude of intrinsic and extrinsic reasons, including an inflation of specialized journals and authors' preferences for such journals. The remainder of topics do not demonstrate any discernible increasing or decreasing trend.

We further coupled the individual temporal distributions of topics with a relative popularity of topics plot (Figure 6). Unlike Figure 5, this plot shows topic trends on the same scale. Although "Subsurface Flow Transport was the most popular topic in the 1990s, it steadily lost popularity within our corpus since then. However, "Uncertainty" rose from the second most popular topic in 1991 to become the current most popular topic. The other most popular topics currently are "Water Management", "Precipitation Variability", "Climate Change", "Modeling Calibration", and "Precipitation Observation".

4.3 Inter-topic correlations

An intuitive way to depict inter-topic correlations $R_{k,j}$ are chord-diagrams. Correlation coefficients measure correlations between per-paper topic weights, meaning that a higher $R_{k,j}$ indicates that papers that contain word groups that indicate a high degree of inclusion in topic k also tend to contain word groups that indicate a degree of inclusion in topic j. Positive correlation coefficient between pairs of topics indicate some
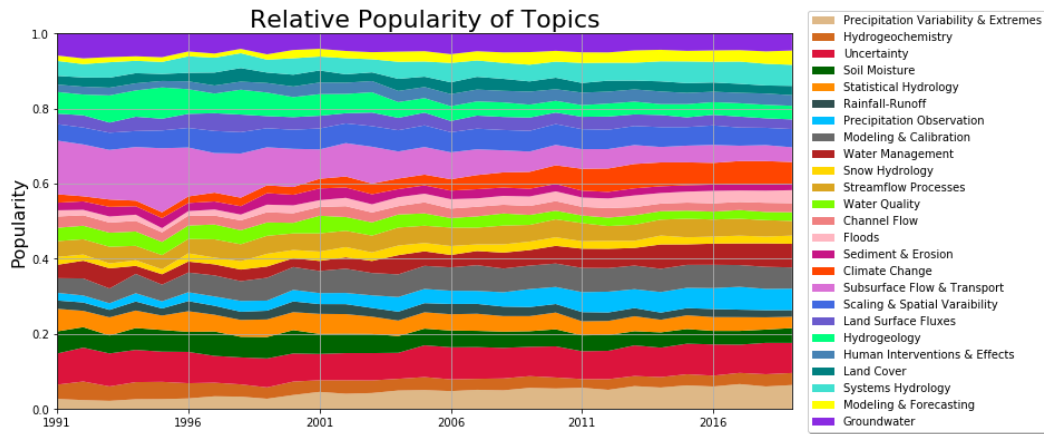
Figure 5. Wordclouds show the words most strongly associated with each topic, and the sizes of words within the wordclouds are proportional to their likelihood of appearance within individual topics. Topic trends are independent and not depicted relative to each other (see Figure 6). Colors representing each particular topic will be followed throughout the rest of this manuscript.

Figure 6. Temporal variation of topic popularity relative to each other.

degree of information exchange between these topics, and vice-versa. Positive and negative inter-topic correlations are shown in Figure 7, where the width of each chord represents the overall correlation between a pair of topics. For ease of viewing, positive correlations are only plotted for $R_{k,j} > 0.05$ and negative correlations $R_{k,j} < -0.05$.

4.3.1 Positive inter-topic correlations

Both modeling related topics - "Modeling Calibration" and "Modeling Forecasting" are (predictably) positively correlated with "Uncertainty" indicating uncertainty quantification research is a commonality in hydrological modeling communities. A distinctly significant correlation can be observed between "Scaling Spatial Variability" and "Rainfall-Runoff" topics, pertaining to the scale dependencies of rainfall-runoff models and studies (e.g., Chiew et al., 2010; Faurès, Goodrich, Woolhiser, & Sorooshian, 1995; Koren et al., 1999). "Systems Hydrology" demonstrates strong correlations with "Water Management" and "Floods". "Human Interventions Effects" is a topic about the impacts of anthropogenic interventions on natural hydrosystems. Research communities working within this domain clearly (and plausibly) exchange information with a number of other topics, including "Climate Change", "Sediment Erosion", "Floods", "Water Quality" and "Precipitation Variability Extremes". Multiple studies focus on the impacts of human interventions and climate change on natural hydrosystems (e.g., Gornitz, Rosenzweig, & Hillel, 1997; Haddeland et al., 2014; Mittal, Bhave, Mishra, & Singh, 2016). Studies also relate anthropogenic interventions with changing water quality and erosion (e.g. Nicholls et al., 2018; Rahman, Hassan, Islam, & Shamsad, 2000; Romanescu, 2013).

Subsurface and related research communities - e.g., "Groundwater", "Hydrogeochemistry", "Water Quality", "Hydrogeology" - also demonstrate significant relationships. We again observe such patterns between precipitation related topics, i.e. "Snow Hydrology" and "Precipitation Observation"; "Rainfall-Runoff","Precipitation Observation" and "Precipitation Variability Extremes". Again, as might be expected, "Land Cover" research demonstrates clear exchange with the "Soil Moisture" and "Land Surface Flux" topics.

4.3.2 Negative inter-topic correlations

One distinct narrative from the analysis of negative inter-topic correlations is the lack of papers associated with both surface and subsurface related topics. Both model-

Figure 7. Inter-topic correlations: positive correlations in the upper subplot and negative correlations in the lower subplot. Only correlations with significance at $\alpha = 0.10$ are shown.
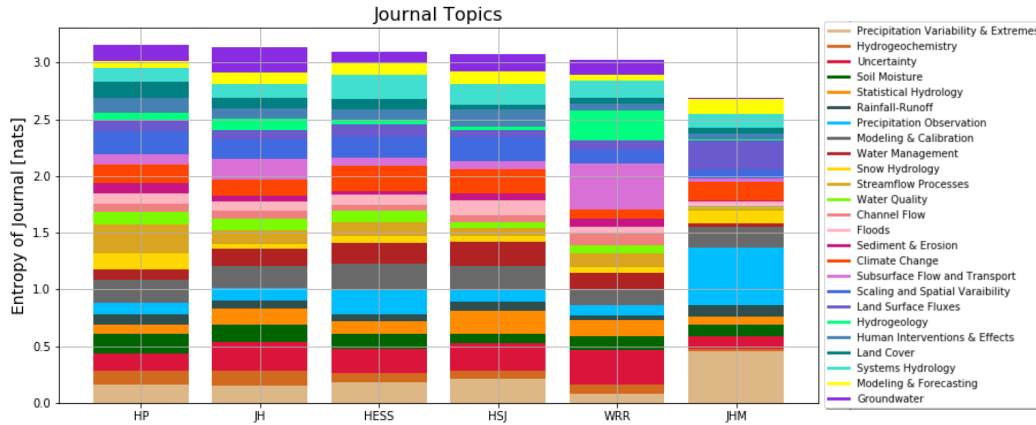
Figure 8. Total bar height represents the overall diversity of topic distributions of each journal for the whole study period. The stacked color bars represent the fraction of papers representing each individual topic in that journal.

ing related topics are negatively correlated with topics such as "Hydrogeology", "Hydrogeochemistry" and "Water Quality". Some perhaps unexpected absences of correlation are between "Groundwater" and "Systems Hydrology", "Modeling Forecasting", "Scaling Spatial Variability", "Soil Moisture", "Uncertainty", "Snow Hydrology" research communities. "Modeling Forecasting" topics lack correlation with "Snow Hydrology", "Water Quality", "Sediment Erosion", "Subsurface Flow Transport", "Hydrogeochemistry", and "Soil Moisture". These negative correlations indicate potential for expanding avenues of collaborative research.

4.4 Journal diversity

We leveraged the unique advantage of topic modeling to provide a contextual understanding of the six high-impact journals in hydrology sampled for this study. Total entropy , H_j, is a measure of the diversity of topics in each journal. The stacked bar plots in Figure 8 show the relative fraction of topic representation in each journal, with the total height of each bar representing the journal's topic entropy.

Most of the journals in this study had relatively similar diversity with HP being the most topic-diverse and JHM being the least. It could be plausibly argued that JHM is a specialty journal, dealing with only one aspect of hydrological research (hydrometeorology); precipitation-related topics dominate that journal. Of the other five journals, WRR is the least diverse, with more papers in the "Water Quality" and "Subsurface Flow and Transport" topics. These are both topics that have topic specific journals, and so it might be the case that if a larger sample of journals was analyzed that we might find that WRR has a more representative mixture of topics than the other journals analyzed here.

Figure 9 shows the temporal variability of topic entropy (diversity) over time. The overall diversity for our entire corpus rose from the 1990s and peaked around 2009. Since then, the overall entropy of the corpus has remained steady or slightly decreased. $HESS$

Figure 9. Temporal variation of the diversity of each journal, as measured by the entropy of that journal's topic distribution in a particular year.

and JHM started publishing in 1997 and 2000 respectively, and the diversity of this corpus rose steadily around this time. JHM again demonstrated lower overall diversity compared with the other five, and even a dip in diversity in 2019 that might be an anomaly. WRR rose steadily in topic diversity during this time period.

4.4.1 Topic Variation by Journals

We use Figure 10, a parallel coordinate plot, to depict the journal-wise variation of topic popularities μ_{kj}. Each topic in the plot is a coordinate and each journal is a line in the coordinate. This analysis offers insights into the distribution of topics within the six journals. As expected, JHM inclines significantly towards precipitation, climatological and forecasting related topics. Conversely, JHM steers away from subsurface related topics such as "Hydrogeochemistry" and "Groundwater". WRR distinctly publishes most "Subsurface Flow Trends" and "Hydrogeology" related papers among the journal set; demonstrating a clear preference of researchers in these topics for WRR. HP leads in publishing research in "Streamflow Processes" topics, while HSJ leads in "Statistical Hydrology" and "Floods" topics. For most of the other topics, including "Uncertainty", "Climate Change", "Scaling Spatial Variability", and "Systems Hydrology", popularity of topics is more homogeneously distributed. HESS and JH appear to be publishing relatively indistinctive, and therefore homogeneous mixture of topics.

4.4.2 Uniqueness and divergence of journals

Differences between journals, as measured by the Jensen-Shannon Distance, d_{js} between pairs of journals, are shown in Figure 11. Here again, we observe significant differences between JHM and the rest of the corpus. The highest degrees of topic similarity are between $HESS$ vs. HP and HJ. WRR is also similar to JH, but less so to $HESS$.

We used the Jensen-Shannon distance from the topic distribution of each journal to the topic distribution of the full corpus, $d_{js(j,m)}$, to represent journal uniqueness. A journal is more unique if this distance is greater, and vice-versa. The temporal variation of these distances for each journal d_{js}^t is demonstrated in Figure 12. This figure shows that the topic distributions in most of the journals are becoming less unique (i.e., the journals are generally becoming more similar). The exception to this HP, which has increased in uniqueness for the past six years (since 2012).
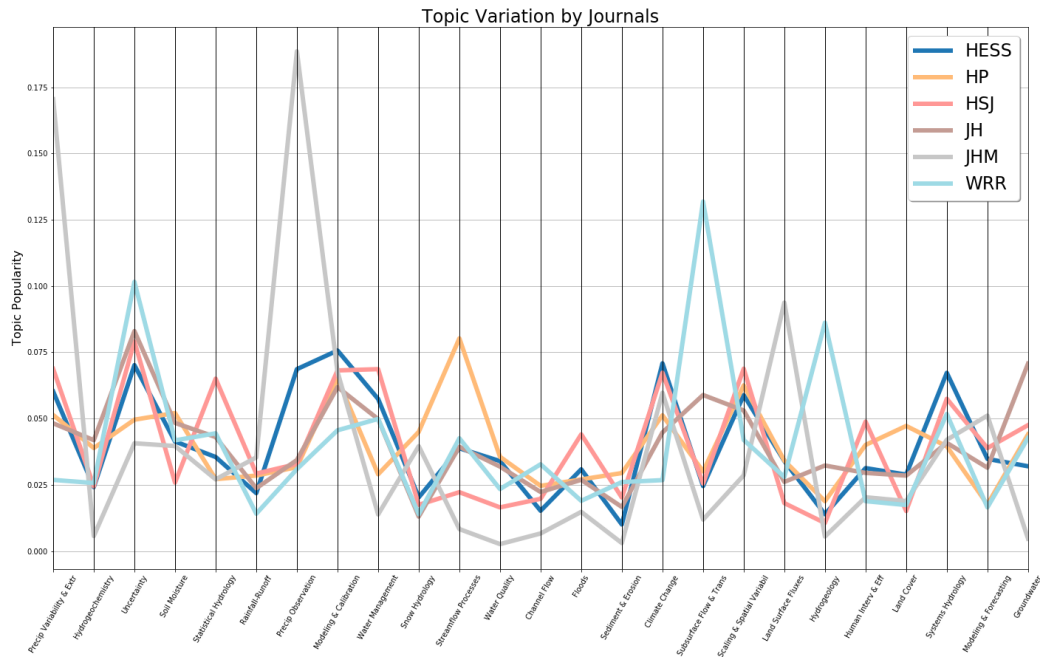
Figure 10. Parallel coordinate plot where each topic is a coordinate and each journal is a line in the coordinate

JHM again demonstrates the highest uniqueness among the six owing to its bias towards more meteorology related topics and papers. Although WRR had the most homogeneous mixture of topics in the early 1990s (Figure 9, and here had the lowest degree of uniqueness relative to the rest of the corpus during the same time period. While both journals increased in topic diversity steadily, JH has retained the most representative journal in this group.

5 Conclusion

5.1 Summary of findings

In this paper, we applied topic modeling using latent Dirchlet allocation (LDA) on the article-abstracts of six high-impact journals in hydrologic science. This yielded a contextual understanding of topic trends and diversity in this corpus of hydrologic science literature using unsupervised learning, without any a priori understanding of or labels on the dataset. Human understanding was used a posteriori to assign topic names. This method leverages commonly available computational resources - i.e., a small compute cluster - to train multiple parallelized LDA models. The resulting topics were carefully identified with the help of veteran hydrologists. Our intent with these experiments is to provide an example of and intuition about LDA to hydrologists, and to help develop a first-order, high-level picture of existing hydrological literature to aid researchers, practitioners, and stakeholders to understand broad themes in hydrological research. of this science, the results were further used to analyze the evolution of topics based on LDA's partitioning of abstract-words for different topics with increasing number of topics.

Posterior document-topic and topic-word distributions generated from the model were aggregated to analyze temporal trends in topic distributions, relative temporal distribution of topics, and inter-topic correlations. Significant inter-topic relationships were observed for data driven topics related to modeling, forecasting, and uncertainty. Some

Figure 11. Jensen-Shannon distance between the whole-period topic distributions in each journal. Low distances indicate similar distributions of topics between two journals.
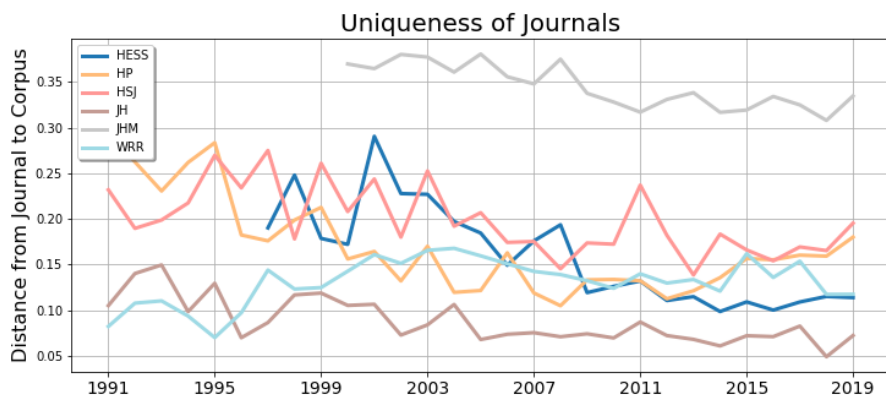

Figure 12. Temporal variation of individual journal uniqueness, measured as the Jensen-Shannon distance of each journal from the entire corpus

subsurface topics such as subsurface flow and transport, groundwater and hydrogeology lost significant popularity within the journals in our sample set. Notable relationships could be seen among research topics and communities concentrating on anthropogenic activities and their impacts on hydrosystems, climate and the environment. Such relationships could also be seen between data-driven research communities, indicating a broader exchange of big data and data-driven methods between them.

We further utilized topic distributions in specific journals to assess the total diversity of topics in individual journals, as well as temporal evolution of journal complexities, uniqueness of individual journals, and differences between topic distributions in pairs of journals. Overall, with increasing volume of publications, the journals in our dataset appear to be broadening their scopes and gradually including a more interdisciplinary variety of research topics.

5.2 Future outlook

The volume of scientific research in general is exploding. It is impossible for anyone to keep up, and practitioners are generally familiar with a very small slice of the literature even in their own field. This makes it difficult for researchers to be confident they fully understand the state of the science, and makes it challenging to expand into new research topics. We envision that in the long-term future, ML will be an integral part of the tool set available to help scientists synthesize the existing state-of-the-science. While this paper does not give us a tool for directly aiding literature review, but it is an early step in helping us understand how we might approach problems related to synthesizing diverse bodies of hydrological literature. There have been several biobliometric analyses of hydrology literature (e.g., Clark & Hanson, 2017; Koutsoyiannis & Kundzewicz, 2007; Rajaram et al., 2015; Zare, Elsawah, Iwanaga, Jakeman, & Pierce, 2017), however ML has the potential to allow for faster, and more contextual analyses of large corpuses. In the future, we envision an interactive website with tools that researchers can use to aid topic-based literature discovery.

References

Blei, D. M. (2012). Probabilistic topic models. *Communications of the ACM*, *55*(4), 77–84.

Blei, D. M., Franks, K., Jordan, M. I., & Mian, I. S. (2006). Statistical modeling of biomedical corpora: mining the Caenorhabditis Genetic Center Bibliography for genes related to life span. *Bmc Bioinformatics*, *7*(1), 250.

Blei, D. M., & Lafferty, J. D. (2006). Dynamic topic models. In *Proceedings of the 23rd international conference on machine learning* (pp. 113–120).

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, *3*(Jan), 993–1022.

Cambria, E., & White, B. (2014). Jumping NLP curves: A review of natural language processing research. *IEEE Computational intelligence magazine*, *9*(2), 48–57.

Chen, S. F., Beeferman, D., & Rosenfeld, R. (1998). Evaluation metrics for language models.

Chiew, F., Kirono, D., Kent, D., Frost, A., Charles, S., Timbal, B., ... Fu, G. (2010). Comparison of runoff modelled using rainfall from different downscaling methods for historical and future climates. *Journal of Hydrology*, *387*(1-2), 10–23.

Choi, H. S., Lee, W. S., & Sohn, S. Y. (2017). Analyzing research trends in personal information privacy using topic modeling. *Computers & Security*, *67*, 244–253. Retrieved from http://www.sciencedirect.com/science/article/pii/S0167404817300603 doi: https://doi.org/10.1016/j.cose.2017.03.007

Clark, M. P., & Hanson, R. B. (2017). The citation impact of hydrology journals. *Water Resources Research*, *53*(6), 4533–4541.

Endres, D. M., & Schindelin, J. E. (2003). A new metric for probability distributions. *IEEE Transactions on Information theory*, *49*(7), 1858–1860.

Faurès, J.-M., Goodrich, D., Woolhiser, D. A., & Sorooshian, S. (1995). Impact of small-scale spatial rainfall variability on runoff modeling. *Journal of hydrology*, *173*(1-4), 309–326.

Feldman, R., Sanger, J., et al. (2007). *The text mining handbook: advanced approaches in analyzing unstructured data.* Cambridge university press.

Gatti, C. J., Brooks, J. D., & Nurre, S. G. (2015). A historical analysis of the field of or/ms using topic models. *arXiv preprint arXiv:1510.05154*.

Gornitz, V., Rosenzweig, C., & Hillel, D. (1997). Effects of anthropogenic intervention in the land hydrologic cycle on global sea level rise. *Global and Planetary Change*, *14*(3-4), 147–161.

Griffiths, T. L., & Steyvers, M. (2004). Finding scientific topics. *Proceedings of the National academy of Sciences*, *101*(suppl 1), 5228–5235.

Haddeland, I., Heinke, J., Biemans, H., Eisner, S., Flörke, M., Hanasaki, N., ... others (2014). Global water resources affected by human interventions and climate change. *Proceedings of the National Academy of Sciences*, *111*(9), 3251–3256.

Hall, D., Jurafsky, D., & Manning, C. D. (2008). Studying the history of ideas using topic models. In *Proceedings of the 2008 conference on empirical methods in natural language processing* (pp. 363–371).

Jardine, J., & Teufel, S. (2014). Topical PageRank: A model of scientific expertise for bibliographic search. In *Proceedings of the 14th conference of the european chapter of the association for computational linguistics* (pp. 501–510).

Jiang, H., Qiang, M., & Lin, P. (2016). A topic modeling based bibliometric exploration of hydropower research. *Renewable and Sustainable Energy Reviews*, *57*, 226–237.

Jussila, J., Mustafee, N., Aramo-Immonen, H., Menon, K., Hajikhani, A., & Helander, N. (2017). A Bibliometric Study on Authorship Trends and Research Themes in Knowledge Management Literature. International Forum on Knowledge Asset Dynamics.

Koren, V., Finnerty, B., Schaake, J., Smith, M., Seo, D.-J., & Duan, Q.-Y. (1999). Scale dependencies of hydrologic models to spatial variability of precipitation. *Journal of Hydrology*, *217*(3-4), 285–302.

Koutsoyiannis, D., & Kundzewicz, Z. W. (2007). *Quantifying the impact of hydrological studies.* Taylor & Francis.

Landauer, T. K., Foltz, P. W., & Laham, D. (1998). An introduction to latent semantic analysis. *Discourse processes*, *25*(2-3), 259–284.

Li, C.-y., & Liu, F.-f. (2018). A Topic Modeling Method for Social Science Literature Based on LDA. *Computer Technology and Development*, *2*, 182–187.

Lin, J. (1991). Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*, *37*(1), 145–151.

Lu, Y., Mei, Q., & Zhai, C. (2011). Investigating task performance of probabilistic topic models: an empirical study of plsa and lda. *Information Retrieval*, *14*(2),

178–203.

Mingers, J., & Leydesdorff, L. (2015). A review of theory and practice in scientometrics. *European journal of operational research*, *246*(1), 1–19.

Mittal, N., Bhave, A. G., Mishra, A., & Singh, R. (2016). Impact of human intervention and climate change on natural flow regime. *Water resources management*, *30*(2), 685–699.

Montanari, A., Young, G., Savenije, H., Hughes, D., Wagener, T., Ren, L., . . . others (2013). "panta rhei—everything flows": change in hydrology and society—the iahs scientific decade 2013–2022. *Hydrological Sciences Journal*, *58*(6), 1256–1275.

Mortenson, M. J., & Vidgen, R. (2016). A computational literature review of the technology acceptance model. *International Journal of Information Management*, *36*(6, Part B), 1248–1259. Retrieved from `http://www.sciencedirect.com/science/article/pii/S0268401216300329` doi: https://doi.org/10.1016/j.ijinfomgt.2016.07.007

Nicholls, R. J., Hutton, C. W., Adger, W. N., Hanson, S. E., Rahman, M. M., & Salehin, M. (2018). *Ecosystem services for well-being in deltas*. Springer International Publishing.

Nisbet, M. C., & Scheufele, D. A. (2009). What's next for science communication? Promising directions and lingering distractions. *American journal of botany*, *96*(10), 1767–1778.

Osterreicher, F., & Vajda, I. (2003). A new class of metric divergences on probability spaces and its applicability in statistics. *Annals of the Institute of Statistical Mathematics*, *55*(3), 639–653.

Paul, M., & Girju, R. (2009). Topic modeling of research fields: An interdisciplinary perspective. In *Proceedings of the international conference ranlp-2009* (pp. 337–342).

Pham, P., Do, P., & Ta, C. D. C. (2018). W-PathSim: Novel approach of weighted similarity measure in content-based heterogeneous information networks by applying LDA topic modeling. In *Asian conference on intelligent information and database systems* (pp. 539–549). Springer.

Priva, U. C., & Austerweil, J. L. (2015). Analyzing the history of cognition using topic models. *Cognition*, *135*, 4–9.

Rahaman, M. M., & Varis, O. (2005). Integrated water resources management: evolution, prospects and future challenges. *Sustainability: science, practice and policy*, *1*(1), 15–21.

Rahman, M. M., Hassan, M. Q., Islam, M. S., & Shamsad, S. (2000). Environmental impact assessment on water quality deterioration caused by the decreased ganges outflow and saline water intrusion in south-western bangladesh. *Environmental Geology*, *40*(1-2), 31–40.

Rajaram, H., Bahr, J. M., Blöschl, G., Cai, X., Scott Mackay, D., Michalak, A. M., . . . Sander, G. (2015). A reflection on the first 50 years of water resources research. *Water Resources Research*, *51*(10), 7829–7837.

Řehřek, R., & Sojka, P. (2011). Gensim—statistical semantics in python. *Retrieved from genism. org*.

Riddell, A. B. (2014). How to read 22,198 journal articles: Studying the history of german studies with topic models. *Distant Readings: Topologies of German culture in the long nineteenth century*, 91–114.

Röder, M., Both, A., & Hinneburg, A. (2015). Exploring the Space of Topic Coherence Measures. In *Proceedings of the eighth acm international conference on web search and data mining - wsdm '15* (pp. 399–408). New York, New York, USA: ACM Press. Retrieved from `http://dl.acm.org/citation.cfm?doid=2684822.2685324` doi: 10.1145/2684822.2685324

Romanescu, G. (2013). Alluvial transport processes and the impact of anthropogenic intervention on the romanian littoral of the danube delta. *Ocean & coastal*

management, *73*, 31–43.

Sebastiani, F. (2002). Machine learning in automated text categorization. *ACM computing surveys (CSUR)*, *34*(1), 1–47.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell system technical journal*, *27*(3), 379–423.

Shu, L., Long, B., & Meng, W. (2009). A latent topic model for complete entity resolution. In *Proceedings of the 2009 ieee international conference on data engineering* (pp. 880–891). IEEE Computer Society.

Sun, S., Luo, C., & Chen, J. (2017). A review of natural language processing techniques for opinion mining systems. *Information Fusion*, *36*, 10–25. Retrieved from `http://www.sciencedirect.com/science/article/pii/S1566253516301117` doi: https://doi.org/10.1016/j.inffus.2016.10.004

Sun, X., Liu, X., Li, B., Duan, Y., Yang, H., & Hu, J. (2016). Exploring topic models in software engineering data analysis: A survey. In *2016 17th ieee/acis international conference on software engineering, artificial intelligence, networking and parallel/distributed computing (snpd)* (pp. 357–362).

Tang, J., Jin, R., & Zhang, J. (2008). A topic modeling approach and its integration into the random walk framework for academic search. In *2008 eighth ieee international conference on data mining* (pp. 1055–1060). IEEE.

Upreti, B. R., Asatiani, A., & Malo, P. (2016). To Reach the Clouds: Application of Topic Models to the Meta-Review on Cloud Computing Literature. In *2016 49th hawaii international conference on system sciences (hicss)* (pp. 3979–3988). doi: 10.1109/HICSS.2016.493

Wang, C., & Blei, D. M. (2011). Collaborative topic modeling for recommending scientific articles. In *Proceedings of the 17th acm sigkdd international conference on knowledge discovery and data mining* (pp. 448–456). ACM.

Wang, X., & McCallum, A. (2006). Topics over time: a non-markov continuous-time model of topical trends. In *Proceedings of the 12th acm sigkdd international conference on knowledge discovery and data mining* (pp. 424–433).

Yang, T.-I., Torget, A., & Mihalcea, R. (2011). Topic modeling on historical newspapers. In *Proceedings of the 5th acl-hlt workshop on language technology for cultural heritage, social sciences, and humanities* (pp. 96–104).

Yau, C.-K., Porter, A., Newman, N., & Suominen, A. (2014). Clustering scientific documents with topic modeling. *Scientometrics*, *100*(3), 767–786.

Zare, F., Elsawah, S., Iwanaga, T., Jakeman, A. J., & Pierce, S. A. (2017). Integrated water assessment and modelling: A bibliometric analysis of trends in the water resource sector. *Journal of Hydrology*, *552*, 765–778.

Zhao, W., Chen, J. J., Perkins, R., Liu, Z., Ge, W., Ding, Y., & Zou, W. (2015). A heuristic approach to determine an appropriate number of topics in topic modeling. In *Bmc bioinformatics* (Vol. 16, p. S8).