Thematic analysis of 18 years of physics education research conference proceedings using natural language processing

Tor Ole B. Odden^{1,*}, Alessandro Marin^{1,†} and Marcos D. Caballero^{1,2}

¹*Center for Computing in Science Education, University of Oslo, 0316 Oslo, Norway*

²*Department of Physics and Astronomy & CREATE for STEM Institute, Michigan State University, East Lansing, 48824 Michigan, USA*

(Received 30 January 2020; accepted 4 June 2020; published 29 June 2020)

We have used an unsupervised machine learning method called latent Dirichlet allocation (LDA) to thematically analyze all papers published in the Physics Education Research Conference Proceedings between 2001 and 2018. By looking at co-occurrences of words across the data corpus, this technique has allowed us to identify ten distinct themes or “topics” that have seen varying levels of prevalence in physics education research (PER) over time and to rate the distribution of these topics within each paper. Our analysis suggests that although all identified topics have seen sustained interest over time, PER has also seen several waves of increased interest in certain topics, beginning with initial interest in qualitative, theory-building studies of student understanding, which gave way to a focus on problem solving in the late 2000s. Since 2010 the field has seen a shift toward more sociocultural views of teaching and learning with a particular focus on communities of practice, student identities, and institutional change. Based on these results, we suggest that unsupervised text analysis techniques like LDA may hold promise for providing quantitative, independent, and replicable analyses of educational research literature.

DOI: [10.1103/PhysRevPhysEducRes.16.010142](https://doi.org/10.1103/PhysRevPhysEducRes.16.010142)

I. INTRODUCTION: THE CHALLENGES AND OPPORTUNITIES OF A GROWING LITERATURE BASE

There is now an overwhelming amount of scientific research literature, even in fields as relatively young as physics education research (PER). For example, since 2005 Physical Review Physics Education Research (formerly Physical Review Special Topics Physics Education Research) has published over 700 articles. The Physics Education Research Conference (PERC) Proceedings have published over 1300 papers since 2001. Journals like *The Physics Teacher* and *The American Journal of Physics* have published many thousands of articles on education, starting in 1963 and 1933, respectively. There are almost certainly thousands more relevant articles in the science education and general education research literature.

This magnitude of literature presents a significant challenge for researchers in that they will never have enough time to read every new article that comes out. Additionally, other than through targeted literature reviews, it is difficult for

researchers to find all of the papers that might be relevant to them. This is doubly the case with new researchers, who have an increasingly long history of research and scholarship with which to familiarize themselves. For this reason, we are more and more reliant on recommendations, whether by colleagues, automated tools such as Google Scholar, or periodic newsletters like those delivered by the American Physical Society.

However, this volume of literature also presents several opportunities. Most scholars have experienced the joy of stumbling on a useful article right when one needs it the most—with so many published articles, there are always likely to be a few undiscovered gems for every reader as long as one knows where to look. At a broader level, this amount of literature indicates that our field has now matured beyond its nascent stages—we have progressed to the point of having sustained several waves of methodological and theoretical development in our decades of existence. Thus, this body of literature can give us insight into the development of our field—where PER has come from and where it is going. However, to perform this kind of analysis we need new methods, as it is an unreasonable proposition to ask a single researcher (or even a research team) to read and summarize all of the articles within the field.

The subfield of machine learning known as natural language processing has, in recent years, developed techniques for this kind of large-scale data mining and text analysis. In this study, we use one such technique, an

*t.o.odden@fys.uio.no

†These authors contributed equally to this work.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

unsupervised machine learning method known as latent Dirichlet allocation (LDA) [1,2], to analyze all of the PERC papers published from the start of the conference proceedings in 2001 through 2018. Using this analysis, we are able to see which themes have been prevalent in the published PER literature and how those themes have changed over time.

The remainder of the paper is organized as follows. In Sec. II, we present an overview of the LDA technique, including its theoretical underpinnings, assumptions, and evaluation methods. We situate this description within a discussion of natural language processing methods intended for a general audience. In Sec. III we describe our specific analysis process start to finish. In Sec. IV we present the results of our LDA model and describe the implications it has for understanding the development of the PER literature over time. Finally, in Secs. V and VI, we compare these results to other attempts at characterizing the PER literature as a whole [3], unpack the limitations of automated text analysis techniques like LDA, and suggest future areas of study that it might be used to address.

II. LITERATURE REVIEW: NATURAL LANGUAGE PROCESSING AND LATENT DIRICHLET ALLOCATION

A. General overview of natural language processing techniques and their utility

We begin with an overview of automated text analysis methods aimed at a general readership. The field of machine learning that focuses on analysis of language and text is known as *natural language processing* (NLP). Broadly speaking, there are two kinds of natural language processing techniques: supervised and unsupervised. *Supervised techniques* take in data that is paired with a set of outcomes and aim to “train” an algorithm to fit, predict, or classify that data. For example, if one has already sorted a set of documents into several different categories, supervised text analysis techniques can be used to “learn” the features of each category in order to predict the categorization of a new, unseen document. The accuracy of this classification can then be tested by running the algorithm on a “holdout set”—a segment of the original dataset that was removed before the algorithm was trained.

In contrast, *unsupervised techniques* do not have a predetermined result—that is to say, users do not specify a desired set of outcomes ahead of time. Instead, they aim to extract latent structure within data that is not known *a priori*: for example, unsupervised NLP can be used for analyzing large amount of texts to determine key themes or features, group texts into thematic clusters, or detect nuances in different styles of speech [4]. Such analyses can, for example, be used to extract insights about trends in social media posts [5], marketing data [6], and scientific research literature [7,8]. These insights and algorithms can

then be applied to create classification and recommendation engines such as those used by shopping websites and news publications [9].

Based on the exploratory aims of our research and our data source for this project (articles from the physics education research literature), we have chosen to use an unsupervised natural language processing method. However, because unsupervised NLP methods do not produce particular, pre-determined outcomes, they can be difficult to evaluate. This is especially the case in applications such as textual data mining, in which multiple possible interpretations of a single text can exist simultaneously.

In their introduction to automated text analysis methods, Grimmer and Stewart [4] address these difficulties by presenting four guiding principles for using such methods:

1. Text analysis techniques are always based on imprecise models of language

More concisely, “all quantitative models of language are wrong—but some are useful” (p. 269). In other words, as with all modeling techniques, many assumptions and simplifications must be made in order to even begin the automated text analysis process. Common examples include removing punctuation or certain words, disregarding word order, and simplifying the many nuances of language. In this way, automated text analysis methods are epistemologically similar to the simplified “toy models” that are regularly used in both physics and physics education research [10].

2. Text analysis is used to get insights into data, but cannot replace careful analysis by humans

Unlike supervised learning methods, in which one can directly compute the accuracy of one’s results, unsupervised methods have no “ground truth” to check against. And, as previously mentioned, it is common for a single text to have multiple possible interpretations for different readers. At the end of the day, humans are always the most reliable evaluators of the results and meaning of automated text analyses. Thus, these methods can amplify what humans are able to do, but cannot replace human evaluation and interpretation.

3. There is no single best method for text analysis

There are many different approaches to automated text analysis, both supervised and unsupervised. Each method has certain strengths and weaknesses—that is, each reveals certain aspects of text data, while hiding others. Particular methods, or models, must be chosen based on the research questions under investigation.

4. Validation is very important

Because unsupervised machine learning methods cannot be directly evaluated on their accuracy, it is important that researchers spend the time evaluating and validating their

results using a variety of different methods. Grimmer and Stuart suggest combining “experimental, substantive, and statistical evidence to demonstrate that the measures are as conceptually valid as measures from an equivalent supervised model” (p. 271). We describe several approaches for validation of our chosen technique, latent Dirichlet allocation, in Sec. II. D.

With these guidelines in mind, we are aiming to use automated text analysis to understand how PER has conceptualized and studied teaching and learning over time, based on the publications produced over its history. To perform this analysis, we have chosen to use an unsupervised text analysis technique called latent Dirichlet allocation.

B. Latent Dirichlet allocation

Latent Dirichlet allocation is a generative probabilistic modeling technique that tries to extract the latent or hidden themes from a set of texts or documents [1]. It assumes that documents are mixtures of a certain number of latent topics, where each topic is a mixture of words. The goal of LDA is to take a group of texts and extract these latent “topics” from them, in the form of groups or sets frequently co-occurring words. The model requires a user to specify the number of topics (K) and a numeric parameter for how “mixed” these topics should be (α) in advance. Thereafter the algorithm iteratively “learns” the distribution of words in each topic, and the distribution of those topics in each document. Finally, the user can take these groups of words and, assuming they are recognizably distinct from one

another, apply meaning to them and give them a topic name or label.

As an illustrative example, Fig. 1 shows an LDA model that we trained on roughly 9000 Huffington Post headlines [11] from 2012 to 2018. We chose to train the model on three distinct subsets of headlines—those that had been previously classified (by a human coder) under the headings of “sports,” “religion,” or “technology”—in order to see whether the model could produce topics that fit with these three pre-established themes. Based on the results, it seems to have been successful; when given $K = 3$ topics, the model was able to produce recognizable sets of words that fit with each of the expected themes—for example, *sports* includes the words “game,” “nfl,” “win,” “player,” and “fan,” while *religion* includes words such as “pope_francis,” “christian,” “church,” and “god.” Notice, however, that the model is not completely coherent from a human-interpretability standpoint, as it has associated the word “muslim” (which clearly belongs with the religion topic) most strongly with the *tech* topic.

The model was also able to classify headlines according to their topic mixtures, and examples are shown that were classified (to within a tolerance of 5%) as purely one topic, 50/50 mixtures of two topics, and equal mixtures of all three topics: for example, one headline that was classified as being approximately 50% religion and 50% technology was the 2016 headline “Mark Zuckerberg Met with Pope Francis and Gave Him a Drone” [12], shown at the bottom of the figure.

FIG. 1. Example results from an LDA model trained on a dataset of roughly 9000 Huffington Post headlines from 2012–2018 [11].

In order to understand how LDA “learns” the topics from a set of documents, it is important to understand the probabilistic way it approaches words, topics, and documents.¹ As previously mentioned, LDA assumes that every document is a mixture of an integer number of topics—mathematically, this means that LDA treats each document as a probabilistic distribution across all topics present in the corpus, where the exact distribution within a single document is a latent variable (not known in advance). Each topic is itself a mixture of words, with some words more likely to be used than others; mathematically, this means that LDA conceptualizes topics as weighted probabilistic distributions of words, with the weight on each word determining how central that word is to the topic (that is, how likely it is to be used in texts about that topic). For example, in the model shown in Fig. 1 words such as “apple,” “facebook,” and “google” would be weighted very highly in the tech topic, while nontech related words like “pope_francis” would be weighted much lower.² However, these word distributions are also latent (not specified in advance).

Using this probabilistic model of topics and words, LDA is then able to reverse engineer the topics in a set of texts, estimating these topic distributions based on the co-occurrence of words in each document. That is, starting with a random initialization of document-topic and word-topic distributions, the LDA algorithm³ iteratively examines the words that occur in each document and updates these distributions until they reach stability, at which point (if all goes well) each topic is primarily composed of words that frequently occur together. Note, however, that both the model of text assumed by LDA and this iterative algorithm disregard word order, which is known in the field of NLP as the “bag of words” approach [4].

Although it is based on a very simplified model of text, LDA has been successfully used in numerous applications. Since its proposal by Blei *et al.* in 2003 [1], LDA has been used in the business world for spam filtering [13], analyzing large software systems [14], and analyzing brand data for marketing [6]. The New York Times has used LDA to create a recommendation engine for its articles which constantly updates as new articles are written and which takes into account reader histories and preferences [9]. In the academic sphere, LDA has been used to analyze political texts [4], dissertations in library science [15], journalistic trends [16], and scientific research articles

[8,17]. Over time the technique has been refined, leading to the development of an online algorithm based on variational inference that make it possible to analyze streams of data [2], incorporation of other latent variables like sentiment [6], and development of a supervised variant of the technique [18]. However, to our knowledge it has not yet been widely deployed within educational research domains, much less discipline-based educational research communities like PER.

We feel that this technique holds promise because it has the potential to reveal interesting patterns in how our field has conceptualized and reconceptualized approaches or theories central to the study of teaching and learning. Thus, our research questions are as follows:

1. Based on an LDA analysis of a selection of the physics education research literature, what approaches and theories have been prevalent over the last two decades?
2. How has the prevalence of these approaches and theories evolved over time?

C. LDA inputs and outputs

In practice, the inputs and outputs of the LDA algorithm are matrices in which each row is a probability distribution, as shown in Fig. 2. The input is a matrix where the rows correspond to documents, and the columns are all the words across the entire corpus (size $D \times V$). Each entry corresponds to the word count, or the number of times a word occurs in the corresponding document. The outputs are two factorized matrices, θ_D and β_K , that split the input using a set of topics $T_{1:K}$. θ_D is the *document-topic matrix*, which has rows corresponding to documents and columns corresponding to topics (dimensions $D \times K$). The entries are the amount of each topic the model predicts is present in each document. β_K is the *topic-word matrix*, where each row is a topic and each column is a single word from the corpus (dimensions $K \times V$). The entries are the likelihood of a particular word being drawn, given a particular topic. Thus, LDA can also be thought of as a kind of probabilistic matrix factorization [2].

Note, however, that in both of the output matrices, each row is a probability distribution across different discrete outcomes (choice of topics or choice of words) so each row will add up to 1. The values in a particular entry will also vary considerably. For example, in the topic-word matrix β_K , if one had a corpus containing 4000 words and every word in a topic T_k was equally likely to be chosen, the values in that row would uniformly be $1/4000 = 2.5 \times 10^{-4}$. However, in actual models, most topics have a large number of words that are rated very low, on the order of 1×10^{-6} or 1×10^{-7} , while a select set of words will be rated highly, on the order of 1×10^{-2} or even higher. It is therefore common, when presenting topics, to pick the 5–10 most likely words for a particular topic and use those as the descriptor for that topic. This is how we chose the words for the topics shown in Fig. 1.

¹Although we have described this model in qualitative terms here, we provide a more mathematically detailed description of this model in Appendix A. Interested readers are also encouraged to consult the descriptions in Refs. [1,8,20].

²Thus, every topic is assumed to include every possible word in the corpus, but weighted differently depending on the nature of the topic.

³Those interested in how the LDA algorithm derives these distributions are encouraged to consult the descriptions of the algorithms in Blei *et al.* [1] and Hoffman *et al.* [2].

FIG. 2. Probabilistic matrix factorization interpretation of LDA.

As can be seen from this matrix factorization view, LDA is based on several assumptions about the input dataset. First, as previously mentioned, it assumes that word order does not matter, only word counts in each document. Although this is a major assumption, which disregards much of the nuance of language, it is commonly made in natural language processing studies. It is also in line with Grimmer and Stewart's [4] first principle of text analysis, "all quantitative models of language are wrong, but some are useful." Second, LDA assumes that all texts will necessarily be composed of a mixture of several topics—in other words, this is a "mixed membership" model, rather than a model that requires that each document only be classified into a single category. So, for example, given the 3 topics from Fig. 1, sports, religion, and tech, one document might be classified as 20% sports, 50% religion, and 30% tech; another document might be 90% religion, 6% sports, and 4% tech; and so on. However, we feel that this is a major strength of the model when trying to analyze a community like physics education, which studies a fairly narrow set of phenomena and frequently mixes, borrows, and exchanges ideas, theories, and methods (as opposed to, for example, the diverse set of disciplines and methods present in a publication venue like *Science*). Third, LDA assumes that words that are central to a specific topic are more likely to be used more frequently in documents about that topic, and that the most highly weighted words within that topic will tend to frequently co-occur within documents. This is known as the distributional hypothesis

of linguistics [19]. So, for example, if the primary topic of a text is dog ownership, the words "dog," "bark," "leash," and "kennel" are likely to be used more often and co-occur more frequently than the words "truck," "laptop," "spaghetti," or "particle accelerator."

D. LDA model evaluation

In the end, the goal of this type of analysis is to produce a single, trained model that can be used to analyze and make claims about both the textual data it was trained on and future data not yet seen. However, as previously mentioned, because LDA is an unsupervised method it requires multiple levels and types of validation in order to ensure that the model results make sense and that one is not arbitrarily choosing a model whose results deviate significantly from the norm. In other words, due to the unsupervised nature of LDA it can be difficult to gauge how one's choice of input parameters will affect the eventual model results. So, in practice, researchers frequently use a variety of evaluation methods, individually or in combination, to compare the quality of LDA models generated under different conditions [4,20,21]. There are several such approaches discussed in the literature:

1. Coherence of topics

Coherence is essentially a measure of how well a topic "hangs together" based on the tendency of the top words in

the topic to co-occur. More technically, it can be thought of as a measure of the degree of semantic similarity between high-scoring words in a topic—that is, are those words being used together and/or in a similar way? This is based on the “distributional hypothesis” of linguistics, that says that words with similar meanings tend to occur in similar contexts [19,20]. Although there are many ways of evaluating coherence, the most commonly used approach creates a “sliding window” across documents, capturing sections of text and looking at the co-occurrences of words from a particular topic within that window [22]. Words that tend to occur together produce a higher-scoring (that is, more coherent) topic. This score, called the “coherence value” or C_V , is then normalized to a scale of 0–1, with a higher value indicating a more coherent topic.

2. Perplexity of topic models

Perplexity is a statistical measure of the log probability of new data (unseen test documents) fitting into the model, based on the geometric mean of the per-word likelihood. Essentially, perplexity is a measure of how well the model fits, represents, or reproduces the statistics of a new set of unseen data, such as Wikipedia articles. However, perplexity has been found to correlate poorly with human interpretability [23].

3. Face validity

Face validity involves using human domain experts to evaluate the quality of produced topics and to make judgments on their consistency and representativeness [24]. In other words, checks of face validity require someone who is familiar with the body of text under analysis to make a judgment about how well the topics fit with their knowledge of the target domain. This is especially helpful when the researchers doing the analysis are not themselves experts in the analyzed literature base [20].

4. Comparisons with qualitative methods

One can also use standard qualitative research methods to evaluate the distinctiveness and interpretability of topics. This can, for example, involve comparing model outputs to the results of human coding using researcher-generated coding schemes [4]. In contrast to face validity checks, these kinds of methods do not necessarily have to involve domain experts—for example, they can involve “intrusion tasks” [23] in which extraneous “intruder” words are inserted into topics and human raters must try to determine which of the words is the intruder.

For our study, we used two of these validation methods. Our primary metric for evaluating our models was to compute the average coherence score based on the top 20 words in each topic. We supplemented this metric with occasional checks on the face validity of our results, as two

of the authors of this study are members of the PER community who are familiar with the general literature and so were able to provide some judgments as to the face validity of the topics. We also showed preliminary results at various points in the analysis to several other members of the community, from a variety of backgrounds, to make sure that our emerging results made sense. We did not use either perplexity or qualitative methods in this study, but we see this as a rich area of research for future projects.

III. METHODS

Our primary tool in this analysis was Gensim, an open-source and widely used Python package for doing natural language processing and topic modeling. Gensim includes an LDA algorithm based on variational inference [2], along with other useful functionality for processing and filtering text data, and certain additional packages for visualization. Along with Gensim we also used the data analysis package pandas, the Natural Language Tool Kit (NLTK), and standard scientific computing libraries such as numpy, scipy, and matplotlib.

Our research approach consisted of 4 stages, as shown in Fig. 3. We elaborate on each of these stages below.

A. Data collection and preprocessing

Our initial data collection involved downloading all available PDFs of the PERC proceedings from 2001 to 2018, which are available online through an open-access license. In total, this resulted in 1302 individual papers. We chose to use PERC proceedings because they provided a useful snapshot of the research being done in the PER community at an intermediate level of development. In other words, PERC proceedings are peer reviewed, which introduces a minimum threshold of research progress in order to be accepted. However, PERC papers are also understood to reflect work-in-progress, which may or may not lead to publication in a full-length format later on. At 4 pages each, PERC papers are significantly easier to write than a full-length research article, and there is no charge to publish in the proceedings. This, we would argue, makes them a more representative sample of the work that has been done than analysis of full-length papers, for example from PRPER.

After downloading all of the papers, we used the package PyPDF2 to scrape the text into a file. We then performed a series of data-cleaning steps, as is standard for text-based machine learning research [25]. First, we applied minor corrections and filtering for cases where papers were included twice (for example, due to corrections) or where the text of a paper had been corrupted. We then removed the references and acknowledgments sections from the texts. We then filtered the text itself by removing numbers, symbols, and punctuation,

FIG. 3. Flowchart diagram of methods: (a) Data collection and preprocessing; (b) Repeated topic modeling across filtering parameters, number of topics, and α values; (c) Aggregation and clustering of topics; (d) Analysis of literature with final model.

lower-casing all words, reconnecting words that had been broken across lines, removing section headers (i.e., “Introduction,” “Conclusion”), and removing *stop words* (common words that do not carry any specialized meaning, such as “the,” “a,” “and,” “which,” or “of”). Finally, we split the resulting blocks of text into lists of individual words, or *tokens*.

After this initial round of data cleaning we used a method from NLTK to lemmatize the words in our dataset. *Lemmatization* is the process of matching words with the same or similar meaning in order to reduce the size and complexity of a dataset. For example, the words “test,” “tested,” “testing,” and “tests” all refer to the same essential activity, but without additional processing LDA would treat them as separate words. A lemmatizer reduces such words to their bases or “stems”—for example, it would reduce all instances of the above-mentioned test-related words to the same base word, test. This kind of data reduction can have the unintended side-effect of reducing the amount of granularity and detail in the data—for example, “physics” and “physical” might be reduced to the same base word, “physic,” even though they tend to be used in different ways. To offset this limitation, we included a part-of-speech tagger (taken from the NLTK library) in our lemmatization function, which helps to distinguish between cases in which a word was used as a noun, verb, adjective, or adverb and match words across different parts of speech. Thus, our final set of words includes, for example, both “gravitational” and “gravity” (distinguishing similar words used in different parts of speech) but has reduced all instances of “answer,” “answers,” and “answering” to the base-word “answer.”

Next, we used an algorithm to find and create bigrams, which are pairs of words that are frequently associated and may carry additional meaning together—for example, “problem solving,” “sense making,” and “high school.” These were then joined together into a single word with an underscore, such as “high_school.”

Finally, we took the resulting lists of words, and reduced them into a “bag of words” matrix, as described in Sec. II.D. This left us with a matrix in which each row represented a document, each column a single word, and each entry the number of times that word appeared in that document. After preprocessing, we were left with 29 026 words and bigrams.

B. LDA modeling and grid search across hyperparameters

The next step involved filtering out the words that occurred very frequently and very rarely in the data. This is a critical step to topic modeling techniques like LDA, because words that occur too frequently provide little insight into the data and can wash out interesting trends, while those that occur seldom (such as names of people or institutions) can significantly increase processing time. To determine the optimal values of these filtering parameters, we created a variety of LDA models with different combinations of the filtering parameters and evaluated them based on the spread and averages of the models’ coherence scores, aiming to find the combination that maximized coherence. Based on these tests, we chose the following filtering parameters:

FIG. 4. Frequency of 20 most common words in the data corpus (29 026 words total). For each word, both the overall number of occurrences and the fraction of documents containing that word are shown.

1. Removal of all words that occurred in more than 55% of the documents (frequent words)

The 5 most frequent words in our dataset were “student,” “physic,” “question,” “course,” and “learn.” After creating a variety of models with different filtering thresholds for the most frequent words and evaluating their average coherence scores, we set this threshold to 55%. This removed the previously mentioned top 5 words, plus 97 others, for a total of 103 words. These words, along with a graph of the frequency of the top 20 words in the corpus, are shown in Figs. 4 and 5.

We have chosen to remove these words because our goal in this study is to reveal interesting differences and structures in the PERC proceedings data. However, in order to see such differences, we needed to remove words that are widely used within the community but may carry different meanings for different research groups based on their methods and theoretical commitments. Most research groups study topics that involve “physics,” “learning,” and “courses” to some degree (see Fig. 4), so it would be no great finding to see that these topics are widespread within PER. Additionally, words such as “model” can be used in multiple different contexts—for example by researchers studying how students build mathematical models [26], computational models [27], generalized mental models [28], use model-based reasoning [29], or who are defining “toy models” of student cognition [10]. Words with such wide-ranging meanings can cause problems for

topic analysis methods like LDA. Our choice of the 55% threshold was intended to remove such words, while leaving in words that had narrower theoretical or methodological significance. For example, the most common words in the dataset after filtering were “score,” “activity,” “teacher,” “interview,” and “energy”; although relatively common, we argue that each of these words is associated with a specific research method (score, interview), context (teacher), or topic (energy).

2. Removal of all words that occurred less than 15 times across all papers in the dataset (rare words)

The filtering for rare words removed many more, on the order of 24 500 words. Although this is a significant amount of the dataset, we made an assumption (common to work in NLP, e.g., Ref. [25]) that these words were unlikely to be central to any specific research topic, because they occurred so rarely. This cut allowed us to increase processing speed, due to the reduced dimensionality of the resulting dataset, and also eliminate words that would likely only contribute noise and blur interesting results.

After applying these two filters, we had reduced our dataset from 29 026 unique words and bigrams to 4645.

In addition to testing the effects of different combinations of these two filtering thresholds, we also computed a variety of models to test the effects of varying our two primary hyperparameters, the numbers of topics (K) and

<i>able</i>	<i>design</i>	<i>include</i>	<i>order</i>	<i>similar</i>
<i>analysis</i>	<i>determine</i>	<i>indicate</i>	<i>paper</i>	<i>small</i>
<i>analyze</i>	<i>develop</i>	<i>individual</i>	<i>particular</i>	<i>student</i>
<i>answer</i>	<i>developed</i>	<i>instruction</i>	<i>physic</i>	<i>study</i>
<i>approach</i>	<i>difference</i>	<i>instructor</i>	<i>point</i>	<i>suggest</i>
<i>ask</i>	<i>different</i>	<i>introductory</i>	<i>possible</i>	<i>support</i>
<i>base</i>	<i>discussion</i>	<i>involve</i>	<i>present</i>	<i>table</i>
<i>begin</i>	<i>education</i>	<i>keywords</i>	<i>problem</i>	<i>take</i>
<i>case</i>	<i>example</i>	<i>know</i>	<i>process</i>	<i>teach</i>
<i>change</i>	<i>experience</i>	<i>knowledge</i>	<i>provide</i>	<i>term</i>
<i>class</i>	<i>explain</i>	<i>large</i>	<i>question</i>	<i>test</i>
<i>common</i>	<i>focus</i>	<i>learn</i>	<i>reason</i>	<i>think</i>
<i>compare</i>	<i>follow</i>	<i>level</i>	<i>related</i>	<i>time</i>
<i>complete</i>	<i>give</i>	<i>like</i>	<i>require</i>	<i>type</i>
<i>concept</i>	<i>goal</i>	<i>make</i>	<i>research</i>	<i>understand</i>
<i>conceptual</i>	<i>group</i>	<i>mean</i>	<i>response</i>	<i>university</i>
<i>consider</i>	<i>help</i>	<i>model</i>	<i>result</i>	<i>use</i>
<i>context</i>	<i>high</i>	<i>need</i>	<i>science</i>	<i>way</i>
<i>course</i>	<i>idea</i>	<i>new</i>	<i>second</i>	<i>work</i>
<i>data</i>	<i>identify</i>	<i>note</i>	<i>set</i>	
<i>department</i>	<i>important</i>	<i>number</i>	<i>show</i>	

FIG. 5. Alphabetical list of most frequent words in data corpus, removed using frequency filter of 55%.

relative mixtures of topics⁴ (α), evaluating the results based on the spread and average value of the models' coherence scores. After these tests, we were unable to distinguish any major differences based on α values, so in the next stage of the project (clustering and aggregation) we created datasets with a mixture of α values. We were, however, able to see differences in model performance by varying the number of topics K that the model produced. In order to choose the number of topics to include in our final model, we used an “elbow method” [20], which involves training a large number of models with varying numbers of topics, plotting the average coherence score vs number of topics, and then choosing topic number that appears at or above a leveling-off point in the graph. Since we used a mixture of α values, we generated elbow plots for each of $\alpha = 1, 5, 7.5, 10$, and 12.5 , and examined each of these plots separately. Three of these plots shown in Fig. 6, with α values of $1, 10$, and 12.5 ; one can see that the average coherence scores in each graph rises with increasing numbers of topics until the model includes 8–10 topics, at which point they level off, indicating a point of diminishing returns. Based on these “elbows” for the different plots and several face validity checks of sample models, we chose a value of $K = 10$ topics for our final model.

⁴For more information on the meaning of the α value, see Appendix A.

It is worth noting that coherence scores, which we used as our primary metric for evaluating topics, vary between 0 and 1, with higher scores indicating more coherent topics. In most of our models we observed coherence scores between 0.4 and 0.52. Although these scores are slightly lower than values achieved by other studies of scientific literature through LDA (e.g., Ref. [20]) we suspect this is a function of the relatively small amount of text included in our analysis. Syed and Spruit [20] have shown that average coherence scores seem to increase with additional documents (up to a point) in their analysis of 15 000 articles from the field of fisheries and aquatic science. Thus, considering our corpus consisted of ~ 1300 4-page papers, we find these coherence values to be acceptable.

C. Aggregation of topics, cluster analysis, and choice of final model

After determining our model parameters, we encountered a significant challenge with the modeling process: because the LDA algorithm randomly initializes both the topic-word matrix β_K and the document-topic matrix θ_D , topics tend to change from run to run depending on the random initial state of the model. In other words, a single model is replicable, as long as it uses the same initialization each time (which is usually provided in the form of an integer-valued random seed). However, if one were to train two models, both using the same dataset and parameters but with different random initializations, the models could

FIG. 6. Graphs of elbow plot with different document-topic mixtures (α values).

return significantly different sets of topics and document categorizations. One can see evidence of this issue from the graphs shown in Fig. 6: each data point in that figure is a model with a different random initialization, resulting in a slightly different set of topics and varying coherence scores. One can think of this variation as being analogous to the LDA function finding local minima in the complex, multidimensional word space defined by our corpus. However, such random effects are a major obstacle if one is trying to choose a particular model to use, because different models could potentially produce different sets of claims.

These issues are to our knowledge not well addressed in the LDA literature. Some studies acknowledge them, while many others ignore them. For example, Roberts *et al.* [24] state that “Blei (2012) provides an excellent overview of LDA and related models but does not mention the issue of local optima at all. The original paper introducing LDA mentions local optima only in passing to warn against degenerate initializations” (p. 9). Those studies that do confront the issue have used techniques such as aggregating coherence scores across multiple runs and then choosing the single model with the highest coherence score [20], comparing the topics of multiple generated models to those of a single “reference model” [24], clustering and aggregating individual topics across multiple runs [5], and using other methods such as non-negative matrix factorization that can be easily aggregated [30]. However, although these techniques are sufficient for generating stable sets of topics we found that they did not allow us to perform our desired analyses, such as computing topic prevalence over time.

For this reason, in our study we stabilized our model using a clustering and aggregation technique inspired by Mantyla *et al.* [31] and Mehta *et al.* [32]. After creating and examining multiple models with different random initializations, we noted that although topics tended to change from run to run, they seemed to circulate around certain recurring themes. For example, one model might have a topic focused around the term “labs” and another topic based on the word “tutorial,” while another model would have a blended topic incorporating both labs and tutorial. Based on this observation, and in line with other

documented approaches in the literature [32], we made the assumption that our topics could be viewed as variations on certain central themes, which were being obscured by the inherent stochastic nature of the LDA technique. In order to determine what these central themes were, we created an aggregate dataset composed of 500 runs of our chosen LDA model, each using the parameters we had defined above (10 topics, multiple α values, and removal of words that appeared less than 15 times or in greater than 55% of the dataset), appending the resulting topics into a matrix. We then ran a clustering analysis on that matrix, using a Gaussian mixture modeling technique, which can be thought of as a more advanced variant of clustering techniques like K -means that allows for nonspherically shaped clusters. We set our clustering coefficient equal to the number of topics K produced by each model (10), and extracted the centroids of each cluster. Finally, we used a cosine similarity metric between these centroids and the 10 topics for each model in our dataset to search for the single model that has the greatest degree of overlap with the cluster centroids. This model was then chosen for our final analysis.

Although somewhat abstract, this aggregation technique has a fairly intuitive interpretation, analogous to the data collection approaches taught in introductory physics. Our original observation, that each model produced slightly varying topics, can be thought of as similar to the commonly encountered problem of finding random variance in a set of measurements. In such cases, one technique is to take many measurements, assuming that they will cluster together, and then compute the mean and standard deviation of the resulting cluster. Each of our 500 models provided one set of 10 “measurements” (that is, topics) which could be thought of as points in the high-dimensional “word space” of our data corpus. Our clustering technique allowed us to find the mean of each cluster of measurements. We then chose a single model whose topics fell closest to each of these mean values, which was used for our final analysis of the literature.

We realize that this model selection process calls into question the objectivity of our analysis. However, as previously mentioned, due to its unsupervised nature,

objectivity (in the form of producing some kind of “ground truth”) is not the goal of this type of analysis. Rather, our goal is to produce a replicable, validated, independent model of the data that reveals interesting trends and structures. By *replicable*, we mean that our methods are transparent and our code can be re-run to produce the same output by any researcher with the requisite computational tools.⁵ By *independent*, we mean that we do not specify any specific desired results or theoretical commitments in advance—although we made various decisions along the line when processing our data, determining model parameters, and selecting our final model, the actual model results arise from an unsupervised algorithm that we do not directly control. By *validated*, we mean that the process of model selection must be informed by multiple levels of validation, as discussed above. By *model*, we mean that these results should be evaluated with all of the previously discussed simplifications and assumptions in mind—that is, as a simplified perspective on an incredibly diverse and complex set of literature.

D. Analysis of literature using the final model

Once we had identified our single representative model, we used it to perform several analyses on our document corpus to both validate that model and answer our research questions. Our first analysis involved finding the papers that were most aligned with each topic (the “most representative papers”) to provide face validity for our topics as well as to characterize the meaning of each topic. This type of analysis is a built-in function of a trained LDA model—in fact, it can also be used to evaluate the topic mixtures of new, unseen documents (for example, future PERC proceedings papers), which we see as a promising area of future research. After examining the most central words and papers for each topic, we created topical names and descriptions, which are shown in Sec. IV A. This analysis allowed us to address our first research question, “Based on an LDA analysis of a selection of the physics education research literature, what approaches and theories have been prevalent over the last two decades?”

Our second analysis involved graphing the “prevalence” of each topic over time, where prevalence was defined as the aggregated percentage of each topic across all documents in a given year. This leveraged the mixed-membership aspect of LDA by adding up the fractional contributions of each document to the different topics over time. Using this analysis, we can see large-scale shifts in the community over time, which may or may not have been visible in individual papers. In other words, this analysis allows us to see if certain topics have become more or less commonly addressed over time, even if those topics are

uniformly distributed across many papers. By aggregating the prevalence of such topics, we can identify broad trends and areas of focus over the last 18 years of PERC proceedings. This analysis helped answer our second research question, “How has the prevalence of these approaches and theories evolved over time?”

IV. RESULTS

A. Topic characterizations and most representative papers

We begin by presenting the topics produced by our model, in order of their appearance within the model, shown in Table I. For each topic, we provide the name we have assigned and the 10 most highly weighted words in that topic, both alone and with their relative weights. As a reminder, each of these topics is actually a probabilistic distribution over the 4645 unique words and bigrams in our corpus (our “bag of words”), with different weights assigned to each word. The weights on words are the normalized probability of choosing a word, given that specific topic—so, for example, the weight of 0.014 on the word “representation” in the first topic means that, on average, there would be a 1.4% chance of randomly choosing this word from the bag of words in our model each time one writes about the representations topic.

To clarify the meaning of each topic, in Table II we also present the 5 most representative papers according to the model, including the year of publication, title of the paper, authors, and the model’s rating of that topic’s prevalence within the paper. This rating of prevalence within the paper can be thought of as a kind of “percent match” between the paper and topic—a value of 0.9 means that the model rated the paper in question as being 90% composed of that topic, with the remaining 10% being a mixture across all other topics.

We then explicitly discuss each of the topics in greater detail and clarify our chosen names for each topic. Thereafter, in Sec. V, we discuss the validity of these groups and potential limitations of this analysis technique.

1. Topic 1: Use and understanding of representations.

This topic focuses on research on student understanding of and difficulties with representations, including mathematical representations (words: “equation”, “mathematical”) and graphical representations (“diagram”, “vector”) as well as some specific physics topics frequently used in this strand of research (“charge”, “force”). The most representative papers for this topic are all about student use and understanding of particular representational forms like graphs, vectors, and integrals, consistency across these representational forms, or subjects like electricity and magnetism that require students to become fluent in different representations.

⁵A selection of code used for this analysis can be found at https://github.com/uio-ccse/PERC_TopicModel.

TABLE I. Topic numbers, names, top 10 words, and top words with weights.

Topic number	Topic name	Top 10 words	Top 10 words with weights
1	Representations	Representation, diagram, force, equation, vector, charge, solve, difficulty, mathematical, energy	0.014 representation 0.013 diagram 0.011 force 0.009 equation 0.009 vector 0.008 charge 0.008 solve 0.007 difficulty 0.007 mathematical 0.007 energy
2	Problem solving	Problem_solve, solution, task, solve, write, strategy, expert, rubric, principle, feedback	0.021 problem_solve 0.019 solution 0.014 task 0.013 solve 0.011 write 0.008 strategy 0.008 expert 0.007 rubric 0.007 principle 0.007 feedback
3	Labs	Lab, activity, experiment, laboratory, simulation, experimental, circuit, report, section, code	0.042 lab 0.028 activity 0.022 experiment 0.012 laboratory 0.010 simulation 0.010 experimental 0.009 circuit 0.007 report 0.006 section 0.006 code
4	Quantitative assessment of concepts	Item, correct, score, assessment, post_test, version, choice, force, pre_test, instrument	0.031 item 0.011 correct 0.011 score 0.010 assessment 0.010 post_test 0.009 version 0.009 choice 0.007 force 0.007 pre_test 0.007 instrument
5	K-12	Teacher, classroom, activity, content, school, curriculum, scientific, practice, high_school, program	0.052 teacher 0.014 classroom 0.011 activity 0.010 content 0.009 school 0.009 curriculum 0.009 scientific 0.008 practice 0.008 high_school 0.007 program

(Table continued)

TABLE I. (Continued)

Topic number	Topic name	Top 10 words	Top 10 words with weights
6	Difficulties with quantum mechanics	State, vector, difficulty, quantum_mechanic, quantum, interview, measurement, particle, write, qm	0.026 state 0.017 vector 0.012 difficulty 0.012 quantum_mechanic 0.011 quantum 0.010 interview 0.009 measurement 0.009 particle 0.008 write 0.007 qm
7	Community, identity	Community, program, participant, identity, practice, project, interview, faculty, stem, graduate	0.016 community 0.013 program 0.011 participant 0.010 identity 0.009 practice 0.008 project 0.007 interview 0.006 faculty 0.006 stem 0.006 graduate
8	Qualitative methodologies and constructivist theory building	Energy, resource, interview, object, explanation, say, go, move, talk, frame	0.020 energy 0.011 resource 0.011 interview 0.007 object 0.006 explanation 0.006 say 0.006 go 0.005 move 0.005 talk 0.005 frame
9	Research-based instruction	Tutorial, faculty, la, ta, lecture, exam, semester, section, classroom, practice	0.022 tutorial 0.021 faculty 0.021 la 0.018 ta 0.018 lecture 0.011 exam 0.011 semester 0.008 section 0.007 classroom 0.007 practice
10	Quantitative surveys of demographic gaps	Score, performance, survey, average, grade, gain, measure, gender, factor, semester	0.040 score 0.017 performance 0.014 survey 0.013 average 0.011 grade 0.011 gain 0.011 measure 0.008 gender 0.008 factor 0.008 semester

TABLE II. Topics, along with the 5 most representative papers, including title, authors, and percent match with that topic according to the LDA model.

Topic	Year	Title	Authors	Prevalence in paper
Representations	2017	Student determination of differential area elements in upper-division physics	Benjamin P. Schermerhorn and John R. Thompson	0.90
	2012	Arrows as anchors: Conceptual blending and student use of electric field vector arrows	Elizabeth Gire and Edward Price	0.88
	2016	The impact of students' epistemological framing on a task requiring representational consistency	Alexandru Maries, Shih-Yin Lin, and Chandralekha Singh	0.87
	2011	Student interpretation of the signs of definite integrals using graphical representations	Rabindra R. Bajracharya, Thomas M. Wemyss, and John R. Thompson	0.87
	2012	Contrasting students' understanding of electric field and electric force	Alejandro Garza and Genaro Zavala	0.83
Problem solving	2012	Transfer of argumentation skills in conceptual physics problem solving	Carina M. Rebello and N. Sanjay Rebello	0.91
	2015	How do multimedia hints affect students' eye movements in conceptual physics problems?	Xian Wu, John Hutson, Lester C. Loschky, and N. Sanjay Rebello	0.87
	2004	Computer problem-solving coaches	Leonardo Hsu and Kenneth Heller	0.82
	2015	Effects of visual cues and video solutions on conceptual tasks	Tianlong Zu, Elise Agra, John Hutson, Lester C. Loschky, and N. Sanjay Rebello	0.80
	2003	Expert-novice comparisons to illuminate differences in perceptions of problem solutions	Kathleen A. Harper	0.79
Lab instruction	2007	Using students' design tasks to develop scientific abilities	Xueli Zou	0.81
	2004	Can computer simulations replace real equipment in undergraduate laboratories?	Noah D. Finkelstein, Katherine K. Perkins, Wendy K. Adams, Patrick B. Kohl, and Noah S. Podolefsky	0.79
	2006	When and how do students engage in sense-making in a physics lab?	Anna Karelina and Eugenia Etkina	0.76
	2017	Using lesson design to change student approaches to dorm-room design prelabs	Katherine Ansell and Mats Selen	0.76
	2016	Using real time physics with different instructional technologies in a circuits lab	Mónica Quezada-Espinoza, Angeles Dominguez, and Genaro Zavala	0.74
Quantitative assessment of concepts	2016	Assessing student learning and improving instruction with transition matrices	Paul J. Walter and Gary A. Morris	0.95
	2017	Showing the dynamics of student thinking as measured by the FMCE	Trevor I. Smith, Kerry A. Gray, Kyle J. Louis, Bartholomew J. Ricci, and Nicholas J. Wright	0.93
	2018	Determining a hierarchy of correctness through student transitions on the FMCE	Kyle J. Louis, Bartholomew J. Ricci, and Trevor I. Smith	0.93
	2012	Exploring student difficulties with pressure in a fluid	Matthew Goszewski, Adam Moyer, Zachary Bazan, and Doris J. Wagner	0.93
	2016	A multi-faceted approach to measuring student understanding	Ian T. Griffin, Kyle J. Louis, Ryan Moyer, Nicholas J. Wright, and Trevor I. Smith	0.89

(Table continued)

TABLE II. (Continued)

Topic	Year	Title	Authors	Prevalence in paper
K-12	2007	How elementary teachers use what we teach: The impact of PER at the K-5 level	Danielle B. Harlow	0.94
	2011	Teacher-driven professional development and the pursuit of a sophisticated understanding of inquiry	Michael J. Ross, Ben Van Dusen, Samson Sherman, and Valerie K. Otero	0.89
	2005	Different views on inquiry: A survey of science and mathematics methods instructors	Thomas Withee and Rebecca S. Lindell	0.87
Difficulties with quantum mechanics	2010	Uncovering the hidden decisions that shape curricula	Danielle B. Harlow	0.87
	2010	Learning pedagogy in physics	Danielle B. Harlow, Lauren Swanson, Hilary A. Dwyer, and Julie A. Bianchini	0.86
	2011	Students' difficulties with quantum measurement	Guangtian Zhu and Chandralekha Singh	0.97
	2015	Student difficulties with quantum states while translating state vectors in Dirac notation to wave functions in position and momentum representations	Emily Marshman and Chandralekha Singh	0.96
	2018	Student difficulties with the corrections to the energy spectrum of the hydrogen atom for the intermediate field Zeeman effect	Emily Marshman, Christof Keebaugh, and Chandralekha Singh	0.96
Community and identity	2006	Student difficulties with quantum mechanics formalism	Chandralekha Singh	0.95
	2014	Developing an interactive tutorial on a quantum eraser	Emily Marshman and Chandralekha Singh	0.95
	2016	Improving representation in physical sciences using a departmental action team	Katherine Rainey, Joel C. Corbo, Daniel L. Reinholz, and Meredith Betterton	0.96
	2018	Characterizing models of informal physics programs	Claudia Fracchiolla, Noah D. Finkelstein, and Kathleen A. Hinko	0.93
	2018	The intersection of identity and performing arts of black physicists	Tamia Williams, Simone Hyater-Adams, Kathleen A. Hinko, Claudia Fracchiolla, Kerstin Nordstrom, and Noah D. Finkelstein	0.92
	2017	An analysis of community formation in faculty online learning communities	Alexandra Lau, Melissa H. Dancy, Joel C. Corbo, Charles R. Henderson, and Andy Rundquist	0.92
	2018	Intense outreach: Experiences shifting university students' identities	Brean Prefontaine, Claudia Fracchiolla, Manuel Vasquez, and Kathleen A. Hinko	0.91
	2018	Examining the productiveness of student resources in a problem-solving interview	Lisa M. Goodhew, Amy D. Robertson, Paula R. L. Heron, and Rachel E. Scherr	0.94
	2010	Generating explanations for an emergent process: The Movement of sand dunes	Lauren Barth-Cohen	0.93
	2012	Examining the positioning of ideas in the disciplines	Vashiti Sawtelle, Tiffany-Rose Sikorski, Chandra Turpen, and Edward F. Redish	0.91
Qualitative methodologies and constructivist theory building	2007	Conceptual dynamics in clinical interviews	Bruce L. Sherin, Victor R. Lee, and Moshe Krakowski	0.89
	2011	Intuitive ontologies for energy in physics	Rachel E. Scherr, Hunter G. Close, and Sarah B. McKagan	0.87

(Table continued)

TABLE II. (Continued)

Topic	Year	Title	Authors	Prevalence in paper
Research-based instructional methods	2009	Towards understanding classroom culture: Students' perceptions of tutorials	Chandra Turpen, Noah D. Finkelstein, and Steven J. Pollock	0.94
	2009	Student perspectives on using clickers in upper-division physics courses	Katherine K. Perkins and Chandra Turpen	0.89
	2011	But does it last? Sustaining a research-based curriculum in upper-division electricity & magnetism	Stephanie V. Chasteen, Rachel E. Pepper, Steven J. Pollock, and Katherine K. Perkins	0.86
	2007	Research-based practices for effective clicker use	C. J. Keller, Noah D. Finkelstein, Katherine K. Perkins, Steven J. Pollock, Chandra Turpen, and Michael Dubson	0.86
Quantitative assessment of demographic gaps	2009	The impact of physics education research on the teaching of introductory quantitative physics	Charles R. Henderson and Melissa H. Dancy	0.83
	2017	Racial and ethnic bias in the Force Concept Inventory	Rachel Henderson and John Stewart	0.95
	2008	The persistence of the gender gap in introductory physics	Lauren E. Kost, Steven J. Pollock, and Noah D. Finkelstein	0.94
	2014	The impacts of instructor and student gender on student performance in introductory modeling instruction courses	Daryl McPadden and Eric Brewster	0.92
	2010	Improved student performance in electricity and magnetism following prior MAPS instruction in mechanics	Saif Rayyan, Andrew Pawl, Analia Barrantes, Raluca E. Teodorescu, and David E. Pritchard	0.92
	2010	Gender differences in physics 1: The impact of a self-affirmation intervention	Lauren E. Kost-Smith, Steven J. Pollock, Noah D. Finkelstein, Geoffrey L. Cohen, Tiffany A. Ito, and Akira Miyake	0.92

2. Topic 2: Problem-solving strategies and instructor support.

This topic focuses on research on problem solving in physics, including words related to problem-solving expertise (“expert,” “strategy”) and feedback (“rubric,” “feedback”). Representative papers for this topic focus on different facets of problem solving such as argumentation skills, computer aids, video solutions, and differences between novices and experts.

3. Topic 3: Laboratory instruction theory and practice.

This topic focuses on research on physics laboratories, including the written components (“report”) and specific activities (“simulation,” “circuit”). The word “code” here might refer to computer code, or the process of coding data for analysis. The most representative papers for this topic focus on different approaches to laboratory instruction such as design tasks, computer simulations, and active learning transformations.

4. Topic 4: Quantitative assessment of conceptual understanding.

This topic focuses on use of quantitative methods to measure student conceptual understanding, with a specific focus on concept inventories. The topic includes a mix of quantitative assessment-related words (such as “item,” “correct,” “score,” “pre_test,” “post_test,” “choice”) and concept-inventory words (such as “force,” “instrument,” and “version”). Representative papers for this topic focus on assessment of student understanding, application of existing concept inventories, methods for analyzing concept inventory results, and development of new concept inventories.

5. Topic 5: K–12 teacher training, professional development, and curriculum design.

This topic focuses on research physics teaching below the college level, including teacher professional development and pedagogy. It includes a number of words that are usually used to refer to teaching in K–12 settings (teacher, “high_school,” “school,” teacher) as well as several more general teaching-related words that could encompass multiple contexts (“content,” “curriculum,” “scientific,” “practice,” “program”). The most representative papers for this topic focus primarily on K–12 teacher professional development, with some papers also focusing on K–12 curriculum design.

6. Topic 6: Student understanding of and difficulties with quantum mechanics.

This topic focuses on student understanding of and difficulties with quantum mechanics. It includes numerous quantum mechanics-related words, as well as two words that are likely more methodological in nature (“interview,”

“write”). Representative papers for this topic focus exclusively on student difficulties with various aspects of quantum mechanics such as the concepts of measurement and energy levels, notation, as well as development of tutorials to help with these difficulties.

7. Topic 7: Communities of practice, identities, and institutional change.

This topic focuses on research on student identities and physics education communities, with a specific focus on institutional change. It includes words such as “community” and “practice,” which we suspect might refer to the communities of practice framework, and the word “identity.” It also includes words that likely refer to institutional change (“program,” “faculty,” “stem,” “graduate”). Representative papers for this topic focus on the development of different types of communities, both for physics learning and for departmental change, and intersectional analyses of physics student identities.

8. Topic 8: Qualitative methodologies and constructivist theory building.

This topic focuses on more qualitative, theoretical case studies of physics student cognition with a specific focus on student discourse and explanation. The topic includes several words commonly associated with studies of student cognition and understanding (“resource,” “frame”) and references to qualitative methodologies such as interviews and observations (interview, “explanation,” “say,” and “talk”). We suspect the primary word, “energy,” refers to the fact that many of the papers aligned with this topic used the concept of energy as the foci for their case studies (e.g., Ref. [33]). Representative papers for this topic focus on interview methodologies, student explanations, and cognitive theories of student understanding such as conceptual resources or ontological categories.

9. Topic 9: Use of research-based instructional techniques.

This topic focuses on use of research-based instructional techniques, such as tutorials and clickers, in lectures and recitation sections. It includes words related to specific techniques (“tutorial”), settings (“lecture,” “classroom,” “section”), assessment (“exam”), and implementation (“faculty,” “la,” “ta”). The most representative papers for this topic focus on use of research-based instructional tools and the effects of these instructional tools on student learning.

10. Topic 10: Quantitative assessment of demographics-based learning gaps.

This topic focuses on quantitative PER research about topics other than conceptual understanding, with an especial focus on gender and ethnicity. It includes one

FIG. 7. Cumulative topic prevalence over time, with shaded areas based on a jackknife estimation of variance. A wider shaded area corresponds to a greater variance in the distribution of prevalence values, while a narrower shaded area indicates a more uniform distribution. The Y axis corresponds to the effective number of papers each year on that topic.

overlapping word with Topic 4 (“score”) and several other words related to quantitative measurement and analysis (“performance,” “survey,” “average,” “grade,” “gain,” “gender,” “factor”). Representative papers for this topic mostly focus on measuring the effects of gender and race or ethnicity on student concept inventory performance, with a smaller subset focusing on assessing the impact of modeling instruction interventions.

B. Topic prevalence over time

In this section, we examine the changing prevalence of topics over time. We define “prevalence” as the sum of particular topic percentages across documents in a particular year, which can be measured both cumulatively and averaged by year. For example, if in a certain year all documents were evaluated by the model to contain 10% of topic 1, the average prevalence of Topic 1 for that year would be 10%, while the cumulative prevalence would be $0.1 \times N$, where N is the number of documents produced that year. Summing the cumulative prevalence across all topics for a particular year would return the total number of documents in that year. Because LDA is a mixed-membership model, evaluating topic prevalence in this way allows us to see large-scale trends across the PER community that might be thinly distributed across many documents.

Plots of cumulative prevalence over time for each of the 10 topics are shown in Fig. 7. Based on our definition of cumulative prevalence, the y axis of the figure can be thought of as a count of the “effective number of papers” written about a topic each year. For example, a cumulative prevalence of 5.0 for the representations topic would indicate that the equivalent of 5 complete papers about student understanding of representations was produced

that year. We use the term “effective” because in practice every topic is distributed (to varying degrees) across the entire corpus.

The shaded areas in Fig. 7 provide an indication of the distribution of topic prevalence for each year. The heights of these shaded areas were calculated using the jackknife resampling technique⁶ [34], which works as follows: for each topic and year, the sample of prevalence values was divided into 5 equal subsamples, each of which contained 20% of the values. We then calculated the cumulative prevalence 5 times, each time holding out one of the 5 subsamples and normalizing for the missing 20% of the prevalence values. We iterated this process of random splitting-and-calculating 20 times. This procedure yielded, for each topic and year, a new sample of $5 \times 20 = 100$ cumulative prevalence values. We then calculated the standard deviation σ of this sample, and used 3σ as the height of the shaded area.

Based on this approach, the shaded areas can be interpreted as a measure of the level of variation in the topical distribution during a single year. If, during one year, all papers have identical topic prevalence, the jackknife resampling method described above will produce a spread of 0. If, on the other hand, only a small fraction of papers in a certain year focus heavily on that topic, the spread will be much larger.

The results shown in Fig. 7 suggest that all of the topics have seen some sustained growth over time. This, however, is not surprising as there are significantly more papers being published in the PERC proceedings in recent years than in the early years. Certain topics have small spikes at

⁶Necessary because prevalence counts for each topic per year are not Gaussian distributed.

FIG. 8. Cumulative prevalence of topics over time, smoothed using a rolling three-year window.

different points in the history of PERC: specifically, the representations topic has a spike around 2013; problem solving has a dramatic uptick after 2005; and the qualitative theory-building topic has spikes around 2003 and 2012. Both the qualitative theory-building topic and the K–12 topic also see significant downward shifts around 2012–2013. By far the most dramatic visible trend is in the “Communities of practice, identities, and institutional change” topic (Topic 7) which has seen an especially sharp upward trend since 2010, to nearly 30 effective papers per year in recent years. The width of the shaded area in this trend (proportional to the level of variance of the topic distribution within the papers in a certain year) indicates that this topic seems to be more concentrated within a small number papers, suggesting that this trend may be due to the concerted efforts of a dedicated group of researchers rather than an overall shift within the community.

In order to reveal trends over time, we have used a data-smoothing technique which dampens out some of the minor year-to-year variations. Specifically, we have used a “rolling window” function of width 3, which averages the prevalence values for each year with those of the preceding and following year. Although this technique neglects the values of the first and last years, it serves to dampen out small shifts from year to year while keeping larger-scale trends intact allowing us to more clearly see how these trends have developed over time. These smoothed graphs of cumulative and average prevalence over time are shown in Figs. 8 and 9.

Because the number of documents published in PERC proceedings has increased substantially over its years of publication, we present both cumulative and average prevalence over time in Figs. 8 and 9. The graph of cumulative prevalence in Fig. 8 allows us to see the magnitude of

changes throughout the years, while the average prevalence graph in Fig. 9 allows us to see shifts over time that are normalized for the growing PER membership.

Looking at these graphs, we first notice that many of the topics have remained relatively constant over time, even though they are not necessarily very widespread overall. For example, the quantum difficulties topic remains fairly consistent on Fig. 9, and roughly tracks with the other topics based on cumulative prevalence. We suspect that this consistency is due to the fact that, unlike other topics that incorporate multiple research methods and subjects, this topic has focused on a specific subject, based on a distinct vocabulary, over a long period of time. Similarly, both assessment topics remain fairly stable over time on average, which might indicate that this method has seen continued use in recent years through the regular development of targeted concept inventories (e.g., Refs. [35,36]).

Beyond this overall stability, there are three trends immediately visible on both graphs: First, there is a spike in topic 8, “qualitative methodologies and constructivist theory building” during the early years of PER. We suspect this spike may be related to the early push to construct and adapt cognitivist learning theories to physics education (e.g., Refs. [37,38]). A smaller spike in this topic, along with a parallel spike in topic 1 (“representations”) occurs around 2012. Second, there is a wave of research on problem solving starting around 2006, peaking around 2008, and decreasing from there (relative to the rest of the topics). Although we cannot say for certain where this trend comes from, based on the representative papers for this topic we suspect this trend corresponds to a simultaneous focus on problem solving by several PER groups in the Midwestern US, especially Kansas State University (e.g., Refs. [39,40]). Third, we again see a major increase in the

FIG. 9. Average prevalence of topics over time, smoothed using a rolling three-year window.

“communities of practice, identities, and institutional change” topic prevalence since about 2011. This is by far the largest trend on the graph, and based on its low prevalence prior to that point, this seems to be a new trend in the field. Although the most recent year is absent in the smoothed graphs of prevalence over time, one can see from Fig. 7 that the upward trend has continued through 2018. Based on the combined focus on identities, communities of practice, and institutional change visible in the topic’s words and representative papers, we would argue that this trend indicates an increasing shift in the way we, as a community, are conceptualizing teaching and learning physics, towards a more sociocultural and community-based perspective.

V. DISCUSSION: IMPLICATIONS AND LIMITATIONS OF LDA MODEL RESULTS

To answer our first research question, using LDA we have identified 10 enduring themes in PER that have been explored over the last two decades: (i) research on representations; (ii) research on problem solving; (iii) research on laboratory instruction, (iv) assessment of student understanding; (v) research on physics teacher preparation and physics instruction at the K–12 level; (vi) research on student difficulties with quantum mechanics; (vii) Research on communities of practice, identities, and institutional change; (viii) qualitative, constructivist theory building; (ix) research-based instructional methods; (x) assessment of demographics-based learning gaps.

To answer our second research question, this analysis provides suggestions for where the field of PER has come from and where it seems to be going. Based on the trends discussed above, it appears that much of the research in PER has remained fairly steady over the years and that there

remains a strong, enduring interest in all of the research topics mentioned above. However, the field has also had waves of increased interest in several of the topics, beginning with initial interest in qualitative, theory-building studies of student cognition. Over time, that gave rise to an increased focus on problem solving. In more recent years, the field has increasingly moved towards more sociocultural views of teaching and learning, with a special focus on communities of practice, student identities, and institutional change. We note that certain researchers (ourselves included) have anecdotally noticed an increase in studies leveraging these types of sociocultural frameworks in recent years, and this model provides support for that observation.

We realize that many of these topics and trends may not be particularly surprising, especially to long-standing members of the PER community. However, we would argue that this is actually a strength of the analysis since, in addition to the results reported above, this study serves a methodological purpose: it provides a proof-of-concept that natural language processing techniques can successfully be used to analyze large amounts of educational research literature. If readers find these topics and trends to be recognizable and familiar, that should lend credence to our hypothesis that machine learning techniques can contribute to our understanding of large bodies of literature.

One way to check these trends it to compare them to those described in more traditional large-scale literature reviews, such as the 2014 PER literature review compiled by Docktor and Mestre [3]. In that review, they break the field apart into six primary topical areas of study:

- Conceptual understanding
- Problem solving
- Curriculum and instruction
- Assessment

- Cognitive psychology
- Attitudes and beliefs about teaching and learning

Comparing our model to this review, there are several immediate points of overlap: for example, both our model and Docktor and Mestre's review include physics problem solving by name. Both also contain themes explicitly focused on assessment, though the research on assessment described by Docktor and Mestre seems to encompass both of our assessment-related topics, Topics 4 (Assessment: Concepts) and 10 (Assessment: Demographics). Docktor and Mestre's Curriculum and Instruction topical area is defined to cover research on lectures, recitations, and laboratories, which solidly encompasses our topic 3 (Lab Instruction) and topic 9 (Research-Based Instructional Strategies). We also suspect that their Conceptual Understanding topic would likely cover our topic 6 (Quantum difficulties) and topic 8 (Qualitative Theory-Building). All in all, our model seems to have significant overlap with the topical areas they define.

Of course, there are also some areas that do not overlap. For starters, our model has grouped together research on representations (and to some degree electricity and magnetism concepts) into its own topic, while Docktor and Mestre place research on representations under the heading of problem solving. Docktor and Mestre deliberately excluded research on precollege physics education and physics teacher preparation, while our model explicitly covers this research under the heading of topic 5 (K–12). We also do not see any reference to research on student identities, communities of practice, or institutional change in their review, although we note that their article was published in 2014 which (if the trends discussed above are to be believed) was before research on this topic had really begun to take off. On the other hand, our model does not explicitly recognize research on either cognitive psychology or attitudes and beliefs about teaching and learning as their own distinct topics.

In making these comparisons, we do not mean to question the usefulness or thoroughness of large-scale reviews like Docktor and Mestre's. Rather, we aim to demonstrate that our model has produced results which seem to be well in line with many long-standing and long-recognized trends in the field. Thus, as previously mentioned, we argue that this model provides a quantitative, independent, and replicable (but not objective) approach to analyzing the field, as well as a complimentary perspective to Docktor and Mestre's review.

We wish to be clear, however, about the limitations of this type of analysis. One limitation with this method is that it is particularly sensitive to research on distinctive topics that is sustained over time. This type of sensitivity is inherent to the way LDA classification works—if there are distinctive sets of words that consistently co-occur, they are likely to be grouped into a single topic. We suspect this is part of the reason why research on student difficulties with quantum mechanics was identified as its own topic, despite the fact that it is not very

widespread throughout the community—it is a topic of study that has been sustained over time, which uses distinct terminology uncommon to the remainder of the literature. LDA also tends to group together texts that use a similar set of vocabulary, even if they are referring to different phenomena; for example, the community and identity topic includes research on both informal learning communities and institutional change communities. Although a shared vocabulary might indicate a shared set of theoretical commitments (for example, sociocultural views of teaching and learning) the set of work encompassed by these kinds of topics is significantly more diverse than, for example, the work grouped into the laboratory instruction topic.

Additionally, returning to our original description of the LDA method, we wish to revisit the many simplifying assumptions inherent to this type of topic modeling. In order to prepare our dataset for analysis, we removed a significant number of words (including some that are frequently used by certain PER communities), stripped out much of the semantic meaning of sentences, and removed all numbers and symbols, many of which might be of great importance to a paper's argument. Doing so allowed us to analyze large amounts of text in a short amount of time, but this means that the analysis may miss fine-grained distinctions between individual papers. To some degree, this is assumed to be offset by the aggregate trends in word usage (that is, if two papers are actually substantially similar, they will use a similar subset of words), however there is a limit to the amount of information that can be inferred by such trends. As previously stated, this analysis should be considered a kind of toy model of the PER literature, and any results it produces should be evaluated with this in mind.

In the course of this project, we also had to grapple with several problems that are not well addressed in the LDA literature, especially the instability of topic models with regard to random initializations. We addressed this limitation using our ensemble clustering method. As a check on this method, in addition to the model presented above we performed two additional 500-model replication runs with different random seed values. Our first replication run yielded an extremely similar set of 10 topics (with minor reshuffling of words); the second differed from our presented model by one topic (it had two topics focused on representations and only one on quantitative methods). However, both replication runs showed similar large-scale trends over time to those presented above. This, we feel, provides evidence that our ensemble clustering method adequately addressed the instability limitation of LDA.

We note that our choice of dataset likely had much to say for the trends we observed. For example, it is relatively common for researchers to publish multiple PERC proceedings papers on different phases of a single project; thus, trends extracted from PERC proceedings are likely to be “smeared out” across time when compared to those we might see from a journal like *PRPER*. PERC proceedings are

also commonly viewed as a way for junior researchers to gain experience in the publication process, meaning that the papers we analyzed may be skewed towards more conservative research projects (i.e., projects drawing heavily from well-established theories and methods). However, as previously argued, PERC proceedings capture work that may never be published in venues like PRPER, either because the projects never mature into a full paper or because they were published outside of PER. Additionally, we have found that the stability of LDA models depends on the number of documents analyzed, and in this regard it is more advantageous to have a large number of short documents than a smaller number of long documents. Rather than considering this to be a limitation, we see interesting opportunities in using techniques like LDA to make comparisons between journals like the PERC proceedings and PRPER.

Finally, we note that many of the model-building decisions made along the way, from the pre-processing steps [25], to the number of models included in the clustering analysis, to the technique chosen for clustering, likely affected our final results. As previously stated, this is in line with standard practice in the field of natural language processing. As Grimmer and Stewart [24] argue, text analysis can be used to gain insights into data, but cannot replace careful analysis by humans. This means that this model should not be taken as a completely “true” result—rather, it is a simplified model of the system under study, with all of the limitations that entails.

Despite all of these limitations, we would like to reiterate the fact that all of these topics and trends, many of which we suspect will be familiar to members of the PER community, arose from an unsupervised machine learning model. Although we did make numerous decisions throughout the modeling process, we did not specify any of the topics *a priori*—they all arose from the algorithm utilized by the model.

VI. CONCLUSIONS

In this article, we have presented a method for applying natural language processing techniques to extract the topics of a large educational research literature base. Using this technique, we have been able to define and analyze themes present in the PER literature over the last 18 years, as seen through the PERC proceedings. We hope this analysis will shed additional light on where the field has come from and where it seems to be going.

We see great promise for machine learning techniques like LDA and believe that this kind of work could provide a rich avenue for future research. For example, by analyzing analogous publications to PERC from other fields, one might be able to compare the key topics of physics education research with those from other fields of discipline-based educational research like chemistry, biology, or math education. We also see great promise for analyzing larger datasets, over longer timescales. For example, some journals

in the field of science education have over 100 years’ worth of publications. These would be very interesting to analyze, especially with an eye to how the practice of physics education has developed over the last century [41].

We also see room for methodological development, for example, in exploring the effects of different pre-processing decisions or clustering approaches on model results. These could, for example, reveal other trends in the PER literature not explored here. As previously mentioned, there are extensions to LDA that have been developed in order to analyze additional covariate factors, such as word sentiment [24]. Such methods could potentially be used to add in factors like citation rate to see how they relate to the different topics and provide another metric to evaluate topical prevalence. However, we feel these covariate factors would likely be more useful with full journal articles, such as *Physical Review Physics Education Research* or *Science Education*.

In summary, we look forward to using this method for future explorations of the educational research literature.

ACKNOWLEDGMENTS

This work was funded by Norges Forskningsråd Project No. 288125 “International Partnership for Computing in Science Education” and by the Norwegian Agency for International Cooperation and Quality Enhancement in Higher Education (DIKU) which supports the Center for Computing in Science Education. We also thank Anders Malthe-Sørenssen, Christine Lindstrøm, Chandra Turpen, Lin Ding, Elise Lockwood, Marieke Kuijjer, and John M. Aiken for their support and feedback on this project.

APPENDIX A: MATHEMATICAL EXPLANATION OF LDA’S PROBABILISTIC TOPIC MODELING TECHNIQUE

LDA is based on a particular probabilistic model for how text documents are generated, which helps to explain the parameters it uses. In this idealized writing process, a corpus, or complete set of documents D are “written” through the following generative process [1,20]. First, a writer determines the “topics” they will write about, which are conceptualized as distributions across all possible words that could be included in the document:

- (1) For each topic $k \in \{1, \dots, K\}$, draw a distribution β_k over the set of words, V , where $\beta_k \sim \text{Dir}(\eta)$.

This determines the how different topics weight each word from the full set of possible words (or vocabulary) V , where the word-topic distribution β_k is a multinomial distribution based on the Dirichlet distribution parametrized by η , $\text{Dir}(\eta)$. The Dirichlet distribution is a multi-variable generalization of the Beta distribution, which essentially allows one to specify a set of probable outcomes across $(k - 1)$ -dimensional probability simplex. This simplex can be visualized as a kind of “triangle” generalized

to $k - 1$ dimensions with k vertices, where each vertex represents 100% of a particular outcome, and locations in between represent mixtures of several outcomes. For example, in the 3-sided distribution (a triangle) shown in Fig. 1, the “outcomes” correspond to mixtures across the three different topics. Thus, headlines determined to be entirely aligned with one single topic would be located at the vertices, while those weighted as 50/50 would be on located at the midpoint between two vertices. The Dirichlet parameter η specifies how likely points are to fall into a specific corner, with $\eta < 1$ leading to more clumping around the corners, and $\eta > 1$ leading to a smoother distribution across the entire shape.

Once one has determined the topics, the next step is to determine which topics each document will hold:

- 2) For every document $d \in \{1, \dots, D\}$ draw a distribution over topics, $\theta_d \sim \text{Dir}(\alpha)$

where the document-topic distribution θ_d is also distributed Dirichlet, parametrized by α . This means that the particular distribution of topics in the documents is imagined to depend on the Dirichlet parameter α , with small α corresponding to a few topics per document and a large α corresponding to a more even mixture.

Finally, the writer uses these two distributions to generate the words in each document:

- 3) To generate each of the N words in the document d
 - a. Choose a topic, $z_{d,n} \sim \text{Multinomial}(\theta_d)$, where $z_{d,n} \in \{1, \dots, K\}$
 - b. Choose a word $w_{d,n} \sim \text{Multinomial}(\beta_{z_{d,n}})$, where $w_{d,n} \in \{1, \dots, V\}$

Several of these distributions are latent variables (not known *a priori*): the set of topics β_K , the proportion of each topic across documents θ_D , and the assignments of each word to a topic z_N . Only the words in each document w_N can be observed directly. α and η are input parameters (commonly known as “hyperparameters”) which we specify beforehand. Therefore, the joint distribution of all variables (that is, the probability of generating a particular corpus) becomes

$$p(\beta_{1:K}, \theta_{1:D}, z_{1:N}, w_{1:N} | \alpha, \eta) = \prod_{k=1}^K p(\beta_k | \eta) \prod_{d=1}^D p(\theta_d | \alpha) \prod_{n=1}^N p(z_{d,n} | \theta_d) p(w_{d,n} | z_{d,n}, \beta_{d,k}). \quad (\text{A1})$$

Here, $p(\beta_k | \eta)$ is the probability of drawing a particular topic distribution β_k , given a value of the hyperparameter η ; $p(\theta_d | \alpha)$ is the probability of drawing a certain mixture of topics in a document d , given the hyperparameter α ; $p(z_{d,n} | \theta_d)$ is the probability of choosing a specific topic k , given the specific document-topic mixture θ_d ; and $p(w_{d,n} | z_{d,n}, \beta_{d,k})$ is the probability of choosing a certain word w , given the particular topic k and its corresponding topic-word distribution, β_k .

In practice, the latent variables must be inferred, since neither β_K , θ_D , or z_D are known *a priori*. So, our goal becomes to infer these hidden distributions using statistical inference. That is, we wish to compute the conditional probability of a certain set of topic-word distributions, document-topic distributions, and word-topic assignments, given the observed set of words and Dirichlet parameters:

$$p(\beta_{1:K}, \theta_{1:D}, z_{1:N} | w_{1:N}, \alpha, \eta) = \frac{p(\beta_{1:K}, \theta_{1:D}, z_{1:D}, w_{1:D} | \alpha, \eta)}{p(w_{1:N} | \alpha, \eta)}. \quad (\text{A2})$$

However, the marginal probability $p(w_N | \alpha, \eta)$ in the denominator, which is the probability of observing a particular set of documents given a particular α and η , is intractable to calculate exactly because it requires marginalizing over all possible instantiations of the latent distributions and there is a coupling between β and θ [1,20]. In other words, it would require us to calculate

$$p(w_N | \alpha, \eta) = \iint p(\beta' | \eta) p(\theta' | \alpha) \prod_{n=1}^N p(z_n | \theta') p(w_n | z_n, \beta') d\beta' d\theta', \quad (\text{A3})$$

where the integration variables are all possible sets of topic-word distributions and document-topic distributions. Although this is intractable to calculate exactly, it can be approximated using sampling [1,8] or variational inference-based techniques [2]. Our chosen LDA library, Gensim, uses the latter approach.

- [1] D. M. Blei, A. Y. Ng, and M. I. Jordan, Latent Dirichlet allocation, *J. Mach. Learn. Res.* **3**, 993 (2003); <https://dl.acm.org/doi/10.5555/944919.944937>.
- [2] M. D. Hoffman, D. M. Blei, and F. Bach, *Advances in Neural Information Processing Systems*, edited by J. D. A. Lafferty, C. K. I. Williams, J. Shawe-Taylor, and R. S. Zemel (2010), Vol. 23, pp. 856–864, <http://papers.nips.cc/paper/3902-online-learning-for-latent-dirichlet-allocation.pdf>.
- [3] J. L. Docktor and J. P. Mestre, Synthesis of discipline-based education research in physics, *Phys. Rev. ST Phys. Educ. Res.* **10**, 020119 (2014).
- [4] J. Grimmer and B. M. Stewart, Text as data: The promise and pitfalls of automatic content analysis methods for political texts, *Polit. Anal.* **21**, 267 (2013).
- [5] S. J. Blair, Y. Bi, and M. D. Mulvenna, Aggregated topic models for increasing social media topic coherence, *Appl. Intell.* **50**, 138 (2020).
- [6] S. Tirunillai and G. J. Tellis, Mining marketing meaning from online chatter: Strategic brand analysis of big data using latent Dirichlet allocation, *J. Mark. Res.* **51**, 463 (2014).
- [7] G. C. Nunez-Mir, B. V. Iannone, B. C. Pijanowski, N. Kong, and S. Fei, Automated content analysis: addressing the big literature challenge in ecology and evolution, *Methods Ecol. Evol.* **7**, 1262 (2016).
- [8] T. L. Griffiths and M. Steyvers, Finding scientific topics, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 5228 (2004).
- [9] A. Spangher, *Building the Next New York Times Recommendation Engine* (The New York Times, 2015), <https://open.blogs.nytimes.com/2015/08/11/building-the-next-new-york-times-recommendation-engine/>.
- [10] E. F. Redish, Oersted Lecture 2013: How should we think about how our students think?, *Am. J. Phys.* **82**, 537 (2014).
- [11] R. Misra, News Category Dataset, ResearchGate (2018).
- [12] J. Amatulli, *Mark Zuckerberg Met with Pope Francis and Gave Him a Drone* (Huffington Post, 2016), https://www.huffpost.com/entry/mark-zuckerberg-pope-francis-drone_n_57c4337ee4b085c1ff2a2ed7?ncid=engmodushpmg00000006.
- [13] I. Bıró, J. Szabó, and A. A. Benczúr, Latent Dirichlet allocation in web spam filtering, in *AIRWeb 2008—Proceedings of the 4th international workshop on Adversarial information retrieval on the web* (Association for Computing Machinery, New York, 2008), p. 29, <https://doi.org/10.1145/1451983.1451991>.
- [14] G. Maskeri, S. Sarkar, and K. Heafield, Mining business topics in source code using latent Dirichlet allocation, in *Proceedings of the 1st India software engineering conference ISEC'08* (Association for Computing Machinery, New York, 2008), p. 113, <https://doi.org/10.1145/1342211.1342234>.
- [15] C. Sugimoto, D. Li, T. G. Russel, S. C. Finlay, and Y. Ding, The shifting sands of disciplinary development: Analyzing North American library and information science dissertations using latent Dirichlet allocation, *J. Am. Soc. Inf. Sci. Technol.* **62**, 80 (2011).
- [16] C. Jacobi, W. Van Atteveldt, and K. Welbers, Quantitative analysis of large amounts of journalistic texts using topic modelling, *Digit. J.* **4**, 89 (2016).
- [17] M. Paul and R. Girju, Topic modeling of research fields: An interdisciplinary perspective, in *Proceedings of the International Conference Recent Advances in Natural Language Processing RANLP-2009* (Association for Computational Linguistics, Borovets, 2009), p. 337, <https://www.aclweb.org/anthology/R09-1061>.
- [18] D. M. Blei and J. D. McAuliffe, *Advances in Neural Information Processing Systems*, edited by J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis (2007), Vol. 20, pp. 121–128, <https://papers.nips.cc/paper/3328-supervised-topic-models.pdf>.
- [19] Z. S. Harris, *Distributional structure*, *Word* **10**, 146 (1954).
- [20] S. Syed and M. Spruit, Full-Text or abstract? Examining topic coherence scores using latent dirichlet allocation, in *2017 IEEE International Conference on Data Science and Advanced Analytics (DSAA)* (Institute of Electrical and Electronics Engineers, New York, 2018), pp. 165–174, <https://doi.org/10.1109/DSAA.2017.61>.
- [21] S. Debortoli, O. Müller, I. Junglas, and J. vom Brocke, Text mining for information systems researchers: An annotated topic modeling tutorial, *Commun. Assoc. Inf. Syst.* **39**, 110 (2016).
- [22] M. Röder, A. Both, and A. Hinneburg, Exploring the space of topic coherence measures, in *WSDM 2015—Proceedings of the Eighth ACM International Conference on Web Search and Data Mining* (Association for Computing Machinery, New York, 2015), p. 399, <https://doi.org/10.1145/2684822.2685324>.
- [23] J. Chang, J. Boyd-Graber, S. Gerrish, C. Wang, and D. M. Blei, in *Advances in Neural Information Processing Systems*, edited by Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta (2009), Vol. 22, pp. 288–296, <https://papers.nips.cc/paper/3700-reading-tea-leaves-how-humans-interpret-topic-models.pdf>.
- [24] M. E. Roberts, B. M. Stewart, and D. Tingley, Navigating the local modes of big data: The case of topic models, in *Computational Social Science: Discovery and Prediction* (Cambridge University Press, New York, 2016), pp. 51–97, <https://doi.org/10.1017/CBO9781316257340.004>.
- [25] M. J. Denny and A. Spirling, Text preprocessing for unsupervised learning: Why it matters, when it misleads, and what to do about it, *Polit. Anal.* **26**, 168 (2018).
- [26] T. I. Smith, D. B. Mountcastle, and J. R. Thompson, Student understanding of the Boltzmann factor, *Phys. Rev. ST Phys. Educ. Res.* **11**, 1 (2015).
- [27] M. D. Caballero, J. B. Burk, J. M. Aiken, B. D. Thoms, S. S. Douglas, E. M. Scanlon, and M. F. Schatz, Integrating numerical computation into the modeling instruction curriculum, *Phys. Teach.* **52**, 38 (2014).
- [28] C. Fazio, O. R. Battaglia, and B. Di Paola, Investigating the quality of mental models deployed by undergraduate engineering students in creating explanations: The case of thermally activated phenomena, *Phys. Rev. ST Phys. Educ. Res.* **9**, 020101 (2013).
- [29] R. S. Russ and T. O. B. Odden, Intertwining evidence-and model-based reasoning in physics sensemaking: An example from electrostatics, *Phys. Rev. Phys. Educ. Res.* **13**, 020105 (2017).
- [30] M. Belford, B. Mac Namee, and D. Greene, Stability of topic modeling via matrix factorization, *Expert Systems with Applications* **91**, 159 (2018).

- [31] M. V. Mantyla, M. Claes, and U. Farooq, Measuring LDA topic stability from clusters of replicated runs, in *Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement* (Association for Computing Machinery, New York, 2018), pp. 1–5, <https://doi.org/10.1109/CIDM.2014.7008665>.
- [32] V. Mehta, R. S. Caceres, and K. M. Carter, Evaluating topic quality using model clustering, in *2014 IEEE Symposium on Computational Intelligence and Data Mining (CIDM)* (2015), p. 178.
- [33] R. E. Scherr, H. G. Close, and S. B. Mckagan, Intuitive ontologies for energy in physics, *AIP Conf. Proc.* **1413**, 343 (2012).
- [34] B. Efron and C. Stein, The Jackknife estimate of variance, *Ann. Stat.* **9**, 586 (1981).
- [35] P. Barniol and G. Zavala, Mechanical waves conceptual survey: Its modification and conversion to a standard multiple-choice test, *Phys. Rev. Phys. Educ. Res.* **12**, 010107 (2016).
- [36] E. Marshman and C. Singh, Validation and administration of a conceptual survey on the formalism and postulates of quantum mechanics, *Phys. Rev. Phys. Educ. Res.* **15**, 020128 (2019).
- [37] J. Tuminaro, A cognitive framework for analyzing and describing introductory students' use and understanding of mathematics in physics, PhD thesis, University of Maryland, Department of Physics, 2004, p. 154.
- [38] E. F. Redish, A theoretical framework for physics education research: Modeling student thinking, *Proceedings of the International School of Physics "Enrico Fermi"* (IOS Press, Amsterdam, 2003), Vol. 156, pp. 1-64.
- [39] N. S. Rebello, Can we assess efficiency and innovation in transfer?, *AIP Conf. Proc.* **1179**, 241 (2009).
- [40] F. Mateycik, Z. Hrepic, D. Jonassen, and N. S. Rebello, Students' perceptions of case-reuse based problem solving in algebra-based physics, *AIP Conf. Proc.* **951**, 144 (2007).
- [41] D. E. Meltzer and V. K. Otero, A brief history of physics education in the United States, *Am. J. Phys.* **83**, 447 (2015).