

5335 Days of Implementation Science: Using Natural Language Processing to Examine Publication Trends and Topics

Jonathan P. Scaccia (✉ jon@dawnchorusgroup.com)

The Dawn Chorus Group <https://orcid.org/0000-0001-6800-1286>

Victoria C. Scott

UNC Charlotte: University of North Carolina at Charlotte

Methodology

Keywords: Implementation science, natural language processing, synthesis and translation, bibliometric study, systematic review

Posted Date: December 1st, 2020

DOI: <https://doi.org/10.21203/rs.3.rs-113353/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Abstract

Introduction: Moving evidence-based practices into the hands of practitioners requires the synthesis and translation of research literature. However, the growing pace of scientific publications across disciplines makes it increasingly difficult to stay abreast of research literature. Natural Language Processing (NLP) methods are emerging as a valuable strategy for conducting content analyses of academic literature. We sought to apply NLP to identify publication trends in the flagship journal field, *Implementation Science*, including key topic clusters and the distribution of topics over time. A parallel study objective was to demonstrate how NLP can be used in research synthesis.

Methods: We examined 1711 *Implementation Science* abstracts published from February 22, 2006 to October 1, 2020. We retrieved the study data using PubMed's Application Programming Interface (API) to assemble a database. Following standard preprocessing steps, we use topic modeling with latent Dirichlet Allocation (LDA) to cluster the abstracts following a minimization algorithm.

Results: We examined 30 topics and computed topic model statistics of quality. Analyses revealed that published articles largely reflect i) characteristics of research, or ii) domains of practice. Emergent topic clusters encompassed key terms both salient and common to implementation science. HIV and stroke(s) represent the most commonly published clinical areas. Systematic reviews have grown in topic prominence and coherence, whereas articles pertaining to knowledge translation (KT) have dropped in prominence since 2013. Articles on HIV and implementation effectiveness have increased in topic exclusivity over time.

Discussion. We demonstrated how NLP can be used as a synthesis and translation method to identify trends and topics across a large number of (over 1700) articles. With applicability to a variety of research domains, NLP is a promising approach to accelerate the dissemination and uptake of research literature. For future research in implementation science, we encourage the inclusion of more equity-focused studies to expand the impact of implementation science on disadvantaged communities.

Contribution To Implementation Science

- Natural Language Processing (NLP) is an emerging strategy for qualitative research in a variety of disciplines, but has not yet been widely adopted in implementation science.
- NLP can efficiently synthesize a large volume of scientific literature to provide information on topics and trends.
- It can usefully complement existing synthesis approaches (e.g., systematic and scoping reviews) to promote dissemination and uptake of implementation research literature.

Introduction

The classic implementation science axiom says that it takes an average of 17 years for evidence-based practices (EBPs) to be routinized in real-world settings.¹⁻³ This statistic has become so commonplace it is almost an implementation science shibboleth. However, the underlying truth is that good ideas don't always get into practice. Making the first leap from research to community-based trials and then the bigger jump into routine use is laborious and complex. Addressing this challenge is the central purpose of the field.⁴

The Interactive Systems Framework for Dissemination and Implementation was developed as a model to bridge the gap between research and practice.⁵ It delineates three critical systems that make this happen: i) the Delivery System, comprised of individuals and settings involved in the implementation of evidence-based practices (EBPs), (e.g., healthcare clinics, school, rehabilitation centers); ii) the Support System, which builds the readiness of the Delivery System for EBP implementation through training, technical assistance, and other capacity-building strategies;⁶ and iii) the Synthesis and Translation System, which organizes, summarizes, and translates research findings into practitioner-friendly formats. This article hones in on synthesis and translation activities through examining trends in the publication of implementation research.

The accumulation of academic literature has grown exponentially across disciplines in the last 20 years,⁷ including within the field of implementation science. Since the inception of *Implementation Science* - the flagship journal for peer-reviewed research on the scientific study of methods to promote the uptake of research findings into routine healthcare - the number of yearly submissions has increased eight-fold, from 100 submissions in 2006⁸ to over 800 in 2018.⁹ In response to the rise and diverse nature of submissions pertaining to implementation research, a companion journal titled *Implementation Science Communication* was launched in 2019, which accepts a broader range of methodological studies. Two other implementation-specific journals were established in 2019 (*Implementation Research and Practice*; *Global Implementation Research and Action*), signifying the expanding appetite for implementation research. This does not include the multitude of journals that publish implementation science-adjacent literature (*The American Journal of Community Psychology*, *Health Promotion and Practice*, *Frontiers in Public Health*, etc.)

With the rapid growth in publications, an emerging challenge to bridge the research-practice gap is the time required to synthesize research studies. The average time to complete and publish a systematic review is 67.3 weeks (1.3 years).¹⁰ By the time a review is published, the study data are approximately a year dated. The average time to complete a scoping review is 5.2 months, with some efforts requiring as long as 20 months.¹¹ The time-consuming and labor-intensive nature of systematic and scoping reviews partially account for persistent lags in moving research to practice. A fundamental challenge to synthesis and translation activities is how to extract information from research literature more effectively and efficiently.

To understand trends in academic publications and accelerate the synthesis of research literature, scientists have increasingly turned to the use of Natural Language Processing (NLP). NLP is a set of

methods and computer-aided algorithms designed to detect patterns in textual data. By treating words and clustering of words as meaningful, NLP extracts meaning from texts more efficiently than humans are capable of doing. NLP is emerging as a valuable strategy for conducting content analyses of academic literature through bibliometric and scientometric studies. These studies involve the examination of digital data objects (e.g., author, journal, publisher, citations, co-wording of abstracts) to quantify study characteristics within a publication dataset.^{7,12} They can inform the breadth of impact of a search tool (e.g., Web of Science⁷), or shed light on trends within a specific research field.¹²⁻¹³ NLP has been used across a variety of disciplines and consumer applications to understand scholarly publication trends,¹³⁻¹⁶ but has not yet filtered into the methodological mainstream in the social services.

The substantial growth of implementation science, and a personal affinity for this field, sparked our interest in how to leverage these advances to examine how the field has evolved over time. We utilize an NLP approach to identify publication trends in the flagship journal in the field, *Implementation Science*, including the distribution of content areas and association among concepts. Specifically, this study aimed to answer the following research questions:

1. What is the composition of content areas published in *Implementation Science*?
2. How have these content areas changed over time?

While answering these questions, we also have a methodological aim: to demonstrate how NLP can be used in research syntheses. At this time, we do not see NLP-aided methods as replacing the need to actually engage with articles. After all, there is still a large gap between NLP and Natural Language Understanding (NLU), though cutting-edge research moves closer and closer to those benchmarks (such as OpenAI's GPT-3¹⁷). We also note the human reluctance to engage with language models due to the fear of deeply embedded bias.¹⁸ However, we wish to show that NLP can be a powerful synthesis strategy for understanding trends in a field of study and a way that describes and potentially guides research.

Methods

Sample

This study includes the abstracts and associated meta-data for *Implementation Science* articles published between February 22, 2006 (*Welcome to Implementation Science*¹⁹) and October 1, 2020. We retrieved the study data using PubMed's Application Programming Interface (API) to assemble a database that included the following publication attributes: title, authors, affiliation, abstract, keywords, and date indexed. Of primary interest to us for this study were the publication abstracts and the data indexed. There were no applicable EQUATOR standards since i) we limited our database to articles published in one journal and ii) we used all data that was indexed.

As of October 1, 2020, there were 1711 unique article abstracts available in PubMed, which we used in this analysis. We note that there appears to be a day or two lag time between when articles are uploaded to the *Implementation Science* website and when they appear in PubMed. Additionally, articles published in 2008 and 2009 were not indexed until 2010 for unknown reasons. We chose to preserve the PubMed indexing date because that is when a consumer of research would have come across the studies in PubMed.

Data Analysis

We used several NLP methods to review the content in the abstracts. The core of NLP is tokenization, which treats the words as “tokens” that are meaningful information units in and of themselves. Generally, the more frequently that a token occurs, the more relevant it is in descriptive analysis. This is known as the *Bag of Words* approach. However, a fair amount of preprocessing is needed to prepare the text for analysis. We first removed *stop words*, or words that add little informative value to the overall text, such as “and,” “the,” and “with.” We then *lemmatized* the remaining words. Lemmatization involves converting a word down to its root form.²⁰ In some cases, this involves making plural words singular. It also involves normalizing the tense of the word (which is especially relevant when dealing with irregular verbs in English, such as “be”). A similar, less intensive method, *stemming*, is sometimes used because it is less computationally expensive than lemmatization. However, our dataset was not large enough for this to be an issue.

What is the composition of content areas published?

To examine study question 1, we used a topic modeling approach using Latent Dirichlet Allocation (LDA). LDA assumes that each topic is a cluster of words that co-occur together, and that documents (in this case: abstracts) are clusters of topics.²¹⁻²² In order to identify the ideal number of topics (k), researchers want to find the lowest perplexity score across a number of possible topics. *Perplexity* is a metric that looks at how well a probability distribution predicts a sample. To arrive at this minimum, we tested potential topic clusters from two up to 60. While our goal was to minimize the perplexity score, we also wanted to maintain some topic interpretability. This can be a tradeoff; while we may show mathematical improvements, several of the emergent topics may strain a clear interpretation. We decided on a parameter of $k = 30$ topics. While we were continuing to show improvements in the perplexity at higher k s, too many of the emergent topics did not appear to add conceptual value. Topic modeling is a Bayesian process,²² so to ensure replicability of our analysis we set a seed number.

An LDA yields a matrix where each document (in this case: an abstract) is assigned a likelihood of belonging to a specific topic; the γ (“gamma”) value.²³ Researchers generally take the highest γ value to assign a document to a topic, with the caveat that this statistic is a likelihood, rather than a classification. Because we end up with a k -dimensional matrix, we needed a dimensionality reduction algorithm to appropriately visualize the research space. We used Uniform Manifold Approximation and Projection (UMAP), a general-purpose manifold learning and dimension reduction algorithm²⁴ that has advantages

over principal components analysis (PCA) and more advanced techniques like t-distributed stochastic neighbors estimation (tSNE) because the global structure of the data is better preserved. We have posted an interactive UMAP visualization online at www.dawnchorusgroup.com, as the color scheme and embedded data does not translate to print, .pdf., or static .html pages.

Topic modeling is a primary data-driven process, though there are methods to seed the model with specific words.²⁵ Because of this, the LDA outputs may either not be interpretable or so broad that there is little meaning that a human could sense. Boyd-Graber and colleagues²¹ developed a number of metrics to judge the quality of topics generated by an LDA approach (listed below). Like most metrics, there is no rule of thumb that guides absolute cutoff scores.

- *Topic size*: Total number of tokens by topic. There is a strong relationship between topic size and topic quality because common topics are generally represented in many documents, and are not as susceptible to being “diluted” by smaller topics.²¹
- *Mean token length*: The average number of characters for the top tokens in a topic.
- *Prominence*: How many unique documents in which a topic appears. In this type of analysis, we generally find that the methodology-specific topics have higher prominence because they are more likely to feature in many articles.
- *Coherence*: How often the top tokens in each topic appears together in the same document.
- *Exclusivity*: How unique the top tokens in each topic are when compared to the other topics.

How have these content areas changed over time?

We assigned each article to a topic based on its maximum γ value, and then plotted the number of articles in that topic over time against the years indexed in PubMed. To decide which articles to visualize, we used the above metrics to filter the results so we did not end up with a time-series plot with 30 lines, all along a greyscale.

All statistics were computed in R 4.0.2 using a number of open-source packages (tidytext, topicmodels, quanteda, umap). We used what are known as “out-of-the-box” algorithms, meaning we did not need to make any substantial changes to the underlying mathematics. All data that we used in this analysis is available, either directly as a database from the authors, through PubMed, or on *Implementation Science's* website.

Results

We first present general publication trend descriptive data. Apart from the PubMed indexing error for 2008-2009, analyses indicate a growth in the early years of *Implementation Science*, along with a steady volume of articles since 2014 (see Figure 1).

We examined the 25 most frequently published authors in *Implementation Science* (Figure 2). The total number of publications per author in this category ranged from 27-110. We made no distinction in authorship order because different fields (and even different research groups) use different conventions about the meaning of different positions (lead vs. junior researchers).

We then examined the top 25 most frequently occurring terms across all abstracts for a very high-level visual of common themes (see Figure 3). Consistent with the focus of *Implementation Science*, we found that terms relating to health and healthcare frequently appeared, with “implementation” and “intervention” appearing most often.

We examined the composition of content areas published in article abstracts through the association of words within topics. Figure 4 depicts the top five words associated with each topic. The β (“beta”) coefficient indicates the likelihood that a word would appear in a particular topic across all the abstracts in our data.²³ First, we used an inductive process to examine the composition of content areas published based on word associations in each topic.

We observe that the topics reflect qualities of research (e.g. Topic 5: trial, outcome, control) and domains of practice (e.g., Topic 28: measure, assess, construct, scale, item; Topic 29: practice, primary, patient, care, management; Topic 26: patient, physician, prescribe, gp, test). Two clinical health topics emerge as highly published topics: HIV (Topic 24), stroke (Topic 28). Prevalent settings for research are community health (Topic 21) and healthcare (Topic 11, Topic 14, Topic 29.) Of note, we eliminated Topics 4 and 25 from this analysis because these topic clusters appeared to collect terms that had no significant meaning. For example, the key terms of Topic 4 included *Ãç, nÃç, 95, increase, It*. This type of clustering is not uncommon in high-k LDA models.

While we can intuitively identify some themes across topics (e.g., “quality improvement” in Topic 16, “clinical guidelines” in Topic 13, and “evidence” in Topic 17), there needs to be a better method to determine the overall quality of the topic apart from a prima facie glance. Table 1 lists each topic’s empirically derived keywords by β coefficient, along with five LDA core statistics. *Topic size* refers to the number of tokens, or meaningful units of information, per topic. The largest topic sizes pertain to patient safety (potentially around strokes, Topic #11). The topic with the *longest average token length*, that is the longest words, was #7 (about implementation strategy). *Topic prominence* is the number of abstracts in which a topic appears. The most prominent topics were #30 (pertaining to systematic reviews) and #12 (regarding behavioral domains). *Topic coherence* measures how the top tokens in each topic appear together in the same document. Values closer to zero indicated greater coherence. The most coherent topics were #16 (pertaining to quality improvement) and #29 (regarding primary care practice management). *Topic exclusivity* is how unique the tokens in the topic are compared to other topics. Even after we scaled the data, we found the top 5 exclusive topics had identical values. These were Topics 5, 7, 15, 17, 19, & 29. Figure 5 shows the top five topics for each LDA statistics. Because we picked the topic five, there is, by definition, less variation in the scores than if we displayed all topics by these metrics.

Research Question 2: How have these content areas changed over time?

Because a linear plot with 30 topic clusters is not visually pragmatic, we focused on the top five topics as ranked by *prominence, coherence, and exclusivity* (see Figure 6). The spikes in 2010 are a byproduct of the PubMed 2008-2009 indexing error and are excluded from interpretation. The most notable trend is in Topic 30 (systematic reviews). We see Topic 30 rising in prominence and coherence. This indicates that a growth in the appearance of systematic reviews over time and that the terms in this topic cluster tightly.

We observed a drop-off in coherence for topic #5 (RCTs) since 2016. This indicates there are fewer abstracts that have the set of tokens representing these particular topics. We also see a very large drop in the prominence of knowledge translation articles (#8) after 2013, meaning the KT article appears in fewer unique abstracts overall. For exclusivity, there is a drop in primary care practice management (#29) since 2013, and a rise in HIV-specific (#19) and implementation effectiveness (#7) articles.

Discussion

We intended this article to serve two purposes: i) to highlight publication trends in *Implementation Science* across the tenure of the journal, and ii) to demonstrate how NLP can be used as a method of synthesis and translation to support analyses of a large volume of publication content. Below we discuss the implications for both of these aims.

Implications for Implementation Science

Implementation science is a growing field of study. Consistent with publication trends reported by Sales et al,⁹ we found a positive trend in the number of articles published in *Implementation Science* between 2006-2020. Analyses of 28 key topics (per Research Question 1) revealed that published articles largely reflect characteristics of research or domains of practice. Topic clusters encompass key terms that are relevant and common to implementation science (e.g., research synthesis, decision-support, implementation effectiveness, quality improvement). HIV and stroke(s) represent the most commonly published clinical areas.

In examining trends over time (Research Question 2), we found that systematic reviews have grown in topic prominence and coherence in the journal. The prominence of systematic review is consistent with trends reported by Boyd-Graber and colleagues,²¹ indicating that systematic reviews continue to be a unique and popular approach to research synthesis. Articles pertaining to knowledge translation (KT) have dropped in prominence since 2013. We note that KT activities have had more of a political inroad and impact in non-U.S. countries, and especially Canada.²⁶ As such, the decreasing prominence of KT as a topic area may be unique to the study sample, and not reflective of global trends in implementation science. We also found a rise in HIV-related articles overtime, with a publication spike in 2019.

In reflecting on the observed trends, we asked: Could these trends relate to specific editorial changes? We believe it is possible. The editor-in-chief position turned over in 2012 (B. Mittman → A. Sales & M.

Wensig) and 2019 (A. Sales → P. Wilson, with M. Wensig continuing as co-editor-in-chief). We see changes in the LDA core metrics of prominence, exclusivity, and coherence in 2013 and 2020, years which follow changes in editorial leadership. We do not have information about how much influence the editors-in-chief have over types of articles published but recognize that it is customary for journal editors write large to encourage topic-focused publications akin to special issues.

One health-related term of emerging importance that did not appear within any topic clusters was *health equity*. Broadly, health equity refers to the opportunity for all people to experience optimal health.²⁷ It is probable that a health equity focus is embedded within the set of examined implementation science articles, without health equity serving as a primary research outcome (e.g. study by Mizen et al.²⁸ and Woodward et al.²⁹). However, as community psychologists, we strongly believe that quality implementation of evidence-based practices is a precursor to ameliorating underlying disparities in communities. More equity-specific research which includes the development of equitable methodologies and interventions would meaningfully expand the reach (more people have access) and bottom-line impact (it achieves better and more sustainable outcomes) of implementation science. After all, this is the primary goal: to ensure the successful uptake of evidence-based interventions in real-world settings.

Implications for Synthesis and Translation

Full-time, academic researchers can find it arduous to keep pace with the scientific literature given the precipitous rise in academic publications in the recent decades.⁷ The challenge of consuming research literature is greater among working practitioners involved in intervention delivery and/or the provision of care.³⁰⁻³² We have demonstrated how several NLP algorithms can be used to help identify trends and topics across a large number of (over 1700) articles. We observe that several of the topics that emerged in this analysis are evidence that NLP can pick up on these trends because they correspond to what is likely known to *Implementation Sciences'* readers and editors.

NLP is not a panacea for a deep understanding of qualitative data. There is no substitute for deeply engaging with ideas, questioning results, or contemplative analysis of issues to inform research and practice. What we do see is that NLP can be an efficient method to provide a more refined set of guideposts than one could normally obtain from querying PubMed, Google Scholar, or PLoS One. NLP provides a more advanced search and synthesis method, which may be particularly useful when approaching a new field or staying current with research. For example, when using "implementation science" as a search term, 941 articles are returned just for 2020 (as of October 22, 2020) The most dedicated researcher would be challenged to stay currently with that, let alone the practitioner pulled between client and organizational needs.

More efficient and effective synthesis and translation methods are needed in order to reduce lag in uptake of evidence-based practices. For this article we looked at the journal *Implementation Science*. This method can be ported to other search terms like, "decision-support," "barrier and facilitators," and "COVID prevention." A more robust and refined clustering approach can assist the dissemination of scientific

findings. Additionally, NLP methods can be applied to examine correspondence between research literature and societal goals,¹³ implications of particular research terms,¹⁶ along with a host of other issues bearing significance in our communities.³³⁻³⁴

Limitations

Similar to systematic and scoping reviews, the study sample is subjected to publication bias (which is especially true for a high-quality journal like *Implementation Science*). Also, we clustered topics, not findings. This is a critical distinction. We used topic modeling to cluster what the articles were about, not what articles reported as outcomes. Many researchers are working on developing better extractive methods to pull findings and sources of bias out of the article and into a summarization algorithm.³⁵

We limited our search to just *Implementation Science* the journal, not “implementation science” as a search term. This was to constrain our findings to a specific dissemination outlet. It would be possible to replicate this analysis with all 3353 articles returned by PubMed with implementation science as a search term or by pulling in the grey literature. Indeed, looking at this larger corpus would be an interesting next step for our research, which could include more deep semantic modeling with word vectorization models that depend on data-rich inputs.³⁶

But fundamentally, we recognize the inherent challenges in human interpretation of machine-learned models. NLP has not advanced to the point where it can replace human understanding. Because the algorithms are data-driven that can lead to poor interpretations and messy results. A pure summarization-and-synthesis solution is not available yet.

Any applications of NLP, to date, require human intervention. We acknowledge that more advanced language models are yielding incredible results. And yet, these results are only as good as the training data, and we have much work to do ensure our processes yield accurate, actionable, and ethical results.^{18,37-38}

Conclusion

By using NLP, we have demonstrated the ebbs and flows of *Implementation Science* (the journal). Our methods have captured the rise in prominence of systematic reviews, and the clustering of health and methodological-related content. We believe that similar methods can be used in a synthesis and translation process to further the bottom-line goals of *Implementation Science*. We can leverage advances in technology to accelerate the dissemination and uptake of good ideas.

Abbreviations

API

Application Programming Interface; LDA:Latent Dirichlet Allocation; NLP:Natural Language Processing; KT:Knowledge Translation; HIV:Human Immunodeficiency Virus; EBP:Evidence Based Practices; NLU:Natural Language Understanding; UMAP:Uniform Manifold Approximation and Projection; PCA:Principle Component Analysis; tSNE:t-Distributed Stochastic Neighbors

Declarations

Ethics Approval and Consent to Participate:

Human Rights: This study does not include human subjects and is exempt from IRB review.

Consent for Publication: All authors listed approve the paper for submission to Implementation Science.

Availability of Data and Materials: Study data is available at PubMed and Implementation Science. Study materials are available at: www.dawnchorusgroup.com

Competing Interests and Funding: The authors of this manuscript declare that they have no competing interests or funding to report.

Conflicts of Interest: The authors of this manuscript declare that they have no conflicts of interest.

Authors' Contributions: Both authors, JS and VS listed contributed to the design of the study. JS led the data collection and analysis for this study. Both JS and VS drafted the manuscript.

Acknowledgements: None

References

1. Balas EA, Boren SA. Managing clinical knowledge for health care improvement. In: Bemmell J, McCray AT, editors. Yearbook of Medical Informatics 2000: Patient-Centered Systems. Stuttgart, Germany: Schattauer Verlagsgesellschaft mbH; 2000:65-70
2. Grant J, Green L, Mason B. Basic research and health: a reassessment of the scientific basis for the support of biomedical science. *Res Eval.* 2003;12:217–24.
3. Morris ZS, Wooding S, Grant J. The answer is 17 years, what is the question: Understanding time lags in translational research. *J Roy Soc Med.* 2011;104:510–20.
4. Bauer MS, Damschroder L, Hagedorn H, Smith J, Kilbourne AM. An introduction to implementation science for the non-specialist. *BMC Psychol.* 2015;3(1):32.
5. Wandersman A, Duffy J, Flaspohler P, Noonan R, Lubell K, Stillman L, et al. Bridging the gap between prevention research and practice: The interactive systems framework for dissemination and implementation. *Am J of Community Psychol.* 2008;41(3-4):171-81.
6. Wandersman A, Chien VH, Katz J. Toward an evidence-based system for innovation support for implementing innovations with quality: Tools, training, technical assistance, quality

- assurance/quality improvement. *Am J Community Psychol.* 2012;50(3-4):445-459.
7. Li K, Rollins J, Yan E. Web of Science use in published research and review papers 1997–2017: A selective, dynamic, cross-domain, content-based analysis. *Scientometrics.* 2018;115(1):1-20.
 8. Eccles MP, Mittman BS. Welcome to Implementation Science. *Implement Sci.* 2006;1:1 <https://doi.org/10.1186/1748-5908-1-1>
 9. Sales AE, Wilson PM, Wensing M, Aarons GA, Armstrong R, Flottorp S, et al. Implementation Science and Implementation Science Communications: Our aims, scope, and reporting expectations. *Implement Sci.* 2019;14:77. <https://doi.org/10.1186/s13012-019-0922-2>
 10. Borah R, Brown AW, Capers PL, Kaiser KA. Analysis of the time and workers needed to conduct systematic reviews of medical interventions using data from the PROSPERO registry. *BMJ Open.* 2017;7:e012545. doi: 10.1136/bmjopen-2016-012545
 11. Pham MT, Rajic A, Greig JD, Sargeant JM, Papadopoulos A, McEwen SA. A scoping review of scoping reviews: Advancing the approach and enhancing the consistency. *Res Synth Methods.* 2014 Jul;5(4):371-85.
 12. Chen H, Luo X. An automatic literature knowledge graph and reasoning network modeling framework based on ontology and natural language processing. *Adv Eng Inform.* 2019 Jul;42:100959.
 13. Asatani K, Takeda H, Yamano H, Sakata I. Scientific attention to sustainability and SDGs: Meta-analysis of academic papers. *Energies.* 2020 Feb;13(4):975.
 14. Gal D, Thijs B, Glänzel W, Sipido KR. Hot topics and trends in cardiovascular research. *Eur Heart J.* 2019 Jul 21;40(28):2363-2374. doi: 10.1093/eurheartj/ehz282
 15. Shaikh A, Mahoto N, Unar M. Bringing shape to textual data – A feasible demonstration. *Mehran Univ Res J Eng Technol.* 2019;38(4):901-14. doi:10.22581/muet1982.1904.04
 16. Selles OA, Rissman AR. Content analysis of resilience in forest fire science and management. *Land Use Policy.* 2020 May;94:104483.
 17. Vincent, J. OPENAI's latest breakthrough is astonishingly powerful, but still fighting its flaws [Internet]. *The Verge*, 2020 [updated 2020 July 30; Cited 2020 November 8]. Available from: <https://www.theverge.com/21346343/gpt-3-explainer-openai-examples-errors-agi-potential>
 18. O'Neil C. *Weapons of math destruction: How big data increases inequality and threatens democracy.* New York, NY: Broadway Books; 2016.
 19. Eccles MP, Foy R, Sales A, Wensing M, Mittman B. Implementation Science six years on—our evolving scope and common reasons for rejection without reviews. *Implement Sci.* 2012;7:71.
 20. Jivani AG. A comparative study of stemming algorithms. *Int J Comp Tech Appl.* 2011;2(6):1930-38.
 21. Boyd-Graber J, Mimno D, Newman D. Care and feeding of topic models: Problems, diagnostics, and improvements. In: Blei D, Fienberg S, Airoldi E, Erosheva EA editors. *Handbook of mixed membership models and their applications.* Boca Raton, Florida :Chapman and Hall/CRC; 2014.
 22. Jelodar H, Wang Y, Yuan C, Feng X, Jiang X, Li Y, Zhao L. Latent Dirichlet Allocation (LDA) and Topic modeling: Models, applications, a survey. *Multimed Tools Appl.* 2019;78(11):15169-15211.

23. Silge J, Robinson D. Text mining with R: A tidy approach. Sebastopol, CA: O'Reilly Media, Inc. 2017.
24. McInnes L, Healy J, Melville J. Umap: Uniform manifold approximation and projection for dimension reduction. arXiv [preprint]. 2018;arXiv:1802.03426.
25. Grün B, Hornik K. Topicmodels: An R package for fitting topic models. J Stat Softw. 2011;40(13): 1-30.
26. Grimshaw JM, Eccles MP, Lavis JN, Hill SJ, Squires JE. Knowledge translation of research findings. Implement Sci. 2012;7(1):50.
27. Scott VC, Scaccia JP, Alia K. Using Evaluation to Promote Improvements in Health Service Settings. In: Kilmer R, Cook J, editors. The practice of evaluation: Partnership approaches for community change. Thousand Oaks, CA: Sage Publications; 2020.
28. Mizen LA, Macfie ML, Findlay L, Cooper SA, Melville CA. Clinical guidelines contribute to the health inequities experienced by individuals with intellectual disabilities. Implement Sci. 2012;7:42.
<https://doi.org/10.1186/1748-5908-7-42>
29. Woodward EN, Matthieu MM, Uchendu US. et al. The health equity implementation framework: Proposal and preliminary study of hepatitis C virus treatment. Implementation Sci. 2019;14:26.
<https://doi.org/10.1186/s13012-019-0861-y>
30. Atkinson M, Turkel M, Cashy J. Overcoming barriers to research in a magnet community hospital. J Nurs Care Qual. 2008 Dec;23(4):362-8.
31. Bahadori M, Raadabadi M, Ravangard R, Mahaki B. The barriers to application of research findings from the nurses' perspective: A cast study at a teaching hospital. J Educ Health Promot. 2016 Jun;5:14.
32. Lessick S, Perryman C, Billman BL, Alpi KM, De Groote SL, Babin TD Jr. Research engagement of health sciences librarians: A survey of research-related activities and attitudes. J Med Libr Assoc. 2016;104(2):166-173. doi:10.3163/1536-5050.104.2.015
33. Xu JM, Jun KS, Zhu X, Bellmore A. Learning from bullying traces in social media. Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. 2012;665-666.
34. Conway M, Hu M, Chapman WW. Recent advances in using Natural Language Processing to address public health research questions using social media and consumer generated data. Yearb of Med Inform. 2019 Aug;28(1):208-17.
35. Michie S, Thomas J, Johnston M, Mac Aonghusa P, Shawe-Taylor J, Kelly MP, et al. The Human Behaviour-Change Project: Harnessing the power of artificial intelligence and machine learning for evidence synthesis and interpretation. Implement Sci. 2017;12(1):121.
36. Wu L, Yen IE, Xu K, Xu F, Balakrishnan A, Chen PY, et al. Word mover's embedding: From word2vec to document embedding. arXiv [preprint]. 2018;arXiv:1811.01713.
37. Bergstrom CT, West JD. Calling bullshit: the art of skepticism in a data-driven world. New York, NY: Penguin Random House, LLC; 2020.

38. McGuffie K, Newhouse A. The radicalization risks of GPT-3 and advanced neural language models. arXiv [preprint]. 2020;arXiv:2009.06807.

Tables

Table 1. Topic model statistics						
Topic	Keywords	Topic Size	Mean Token Length	Prominence	Coherence	Exclusivity
2	research, policy, health, public, evidence	322.13	6.5	34	-88.49	9.94
3	innovation, context, factor, level, influence	305.88	7.8	18	-87.81	9.94
5	intervention, trial, control, outcome, month	298.51	7.6	33	-62.29	9.95
6	datum, service, method, evaluation, process	328.83	7.5	19	-85.92	9.9
7	implementation, strategy, effectiveness, tailor, implement	287.95	8.6	11	-132.68	9.95
8	kt, network, research, knowledge, translation	362.37	7.5	42	-99.38	9.87
9	decision, professional, patient, system, healthcare	374.13	8.4	25	-94.91	9.91
10	adaptation, tool, framework, guide, develop	326.06	6.3	22	-107.14	9.93
11	patient, hospital, stroke, include, safety	456.9	5.9	37	-101.35	9.82
12	theory, professional, behaviour, domain, theoretical	356.79	7.7	56	-91.61	9.92
13	recommendation, guideline, practice, evidence, clinical	381.98	8.3	39	-107.01	9.9
14	care, nurse, barrier, facilitator, staff	374.16	6	27	-99.46	9.91
15	research, quot, fund, researcher, science	372.39	6.6	32	-131.6	9.95
16	care, team, feedback, quality, improvement	317.84	6.8	31	-81.22	9.93
17	evidence, development, approach, method, develop	323.87	8	14	-91.72	9.95
18	report, information, survey, participant, complete	414.36	7.2	13	-103.17	9.9

19	care, hiv, site, provider, clinic	397.64	5.6	35	-120.15	9.95
20	resource, health, service, low, country	444.38	5.6	42	-120.73	9.9
21	health, sustainability, program, community, prevention	368.73	7.9	24	-107.15	9.91
22	implementation, organizational, organization, support, base	336.46	8.1	49	-103.89	9.9
23	train, fidelity, treatment, base, mental	383.56	7.2	35	-91.36	9.88
24	theory, model, approach, context, process	374.4	7.1	30	-89.96	9.89
26	prescribe, patient, gp, test, physician	452.39	6.9	46	-120.72	9.9
27	change, intervention, target, design, effective	307.89	8.4	14	-96.92	9.94
28	item, scale, measure, construct, assess	407.97	7.4	40	-110.9	9.94
29	care, practice, patient, management, primary	294.12	6.9	32	-76.22	9.95
30	review, report, study, systematic, include	323.62	7.2	62	-75.42	9.93
Notes:						
<i>Topic Size:</i> The total number of tokens by topic; higher is indicative of high quality.						
<i>Mean token length:</i> The average number of characters for the top tokens in a topic, with longer words potentially indicative of better quality of topics.						
<i>Prominence:</i> The number of unique abstracts in which a topic appears.						
<i>Coherence:</i> How often each topic's top tokens appear together in the same abstract.						
<i>Exclusivity:</i> How unique the top tokens in each topic are when compared to the token in other topics.						

Figures

Distribution of when articles were indexed in PubMed

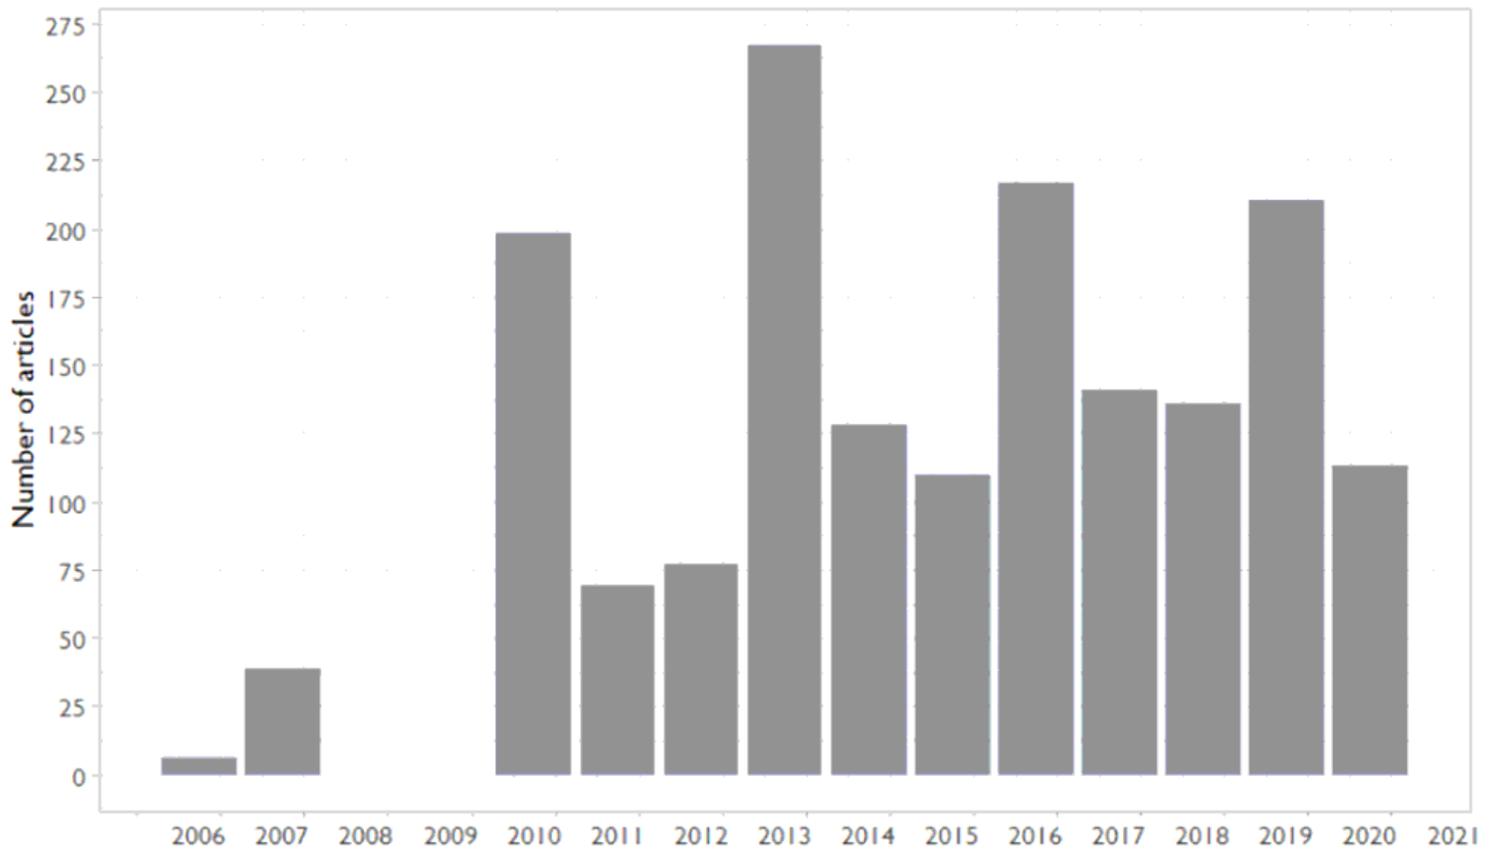


Figure 1

Number of Published Articles Indexed in PubMed between 2006-2020

Authors appearing in Implementation Science

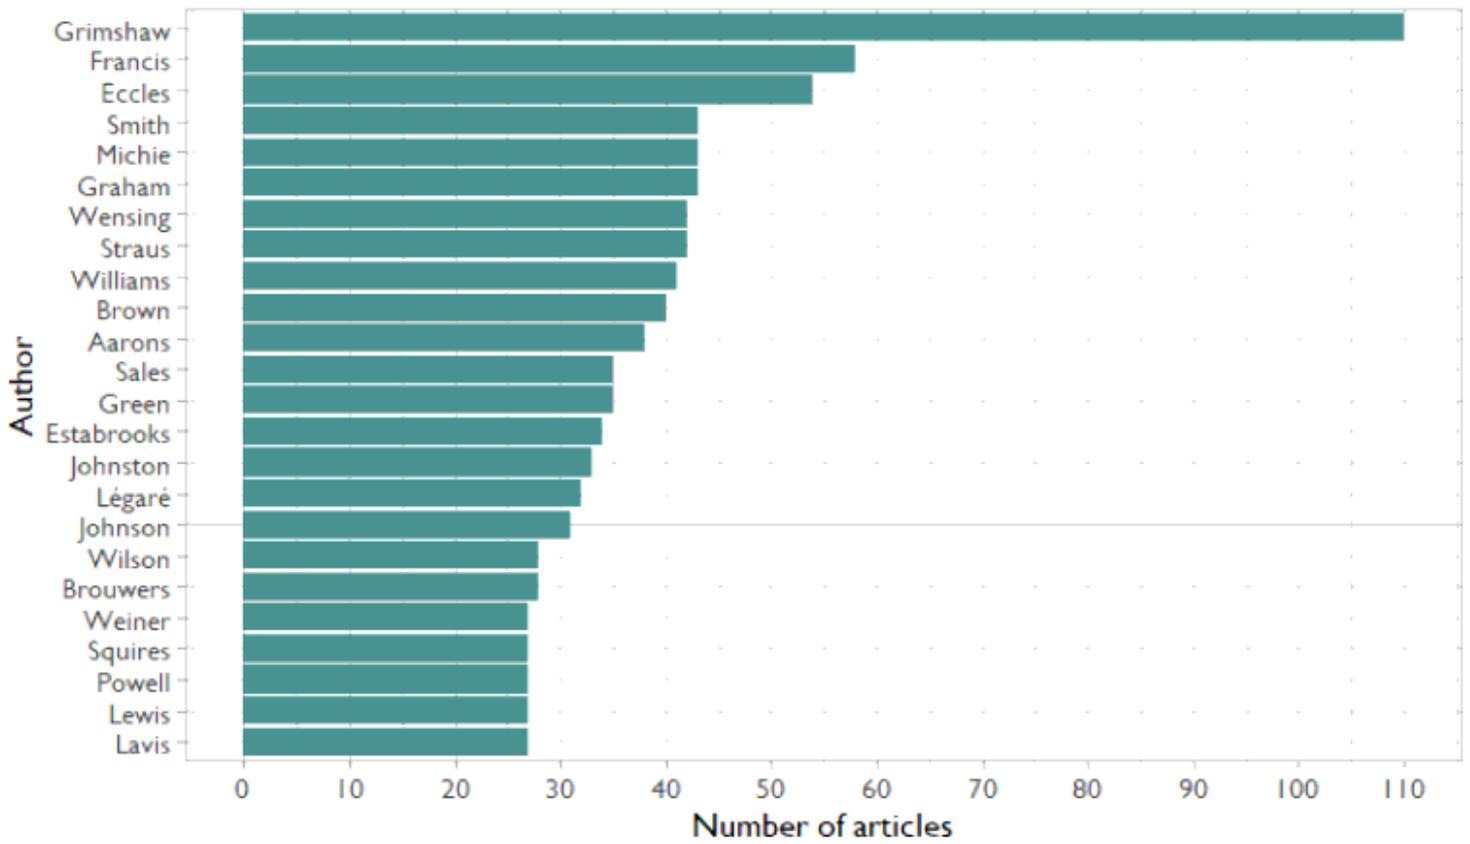


Figure 2

Top 25 Published Authors in Implementation Science between 2006-2020

Most frequently occurring terms

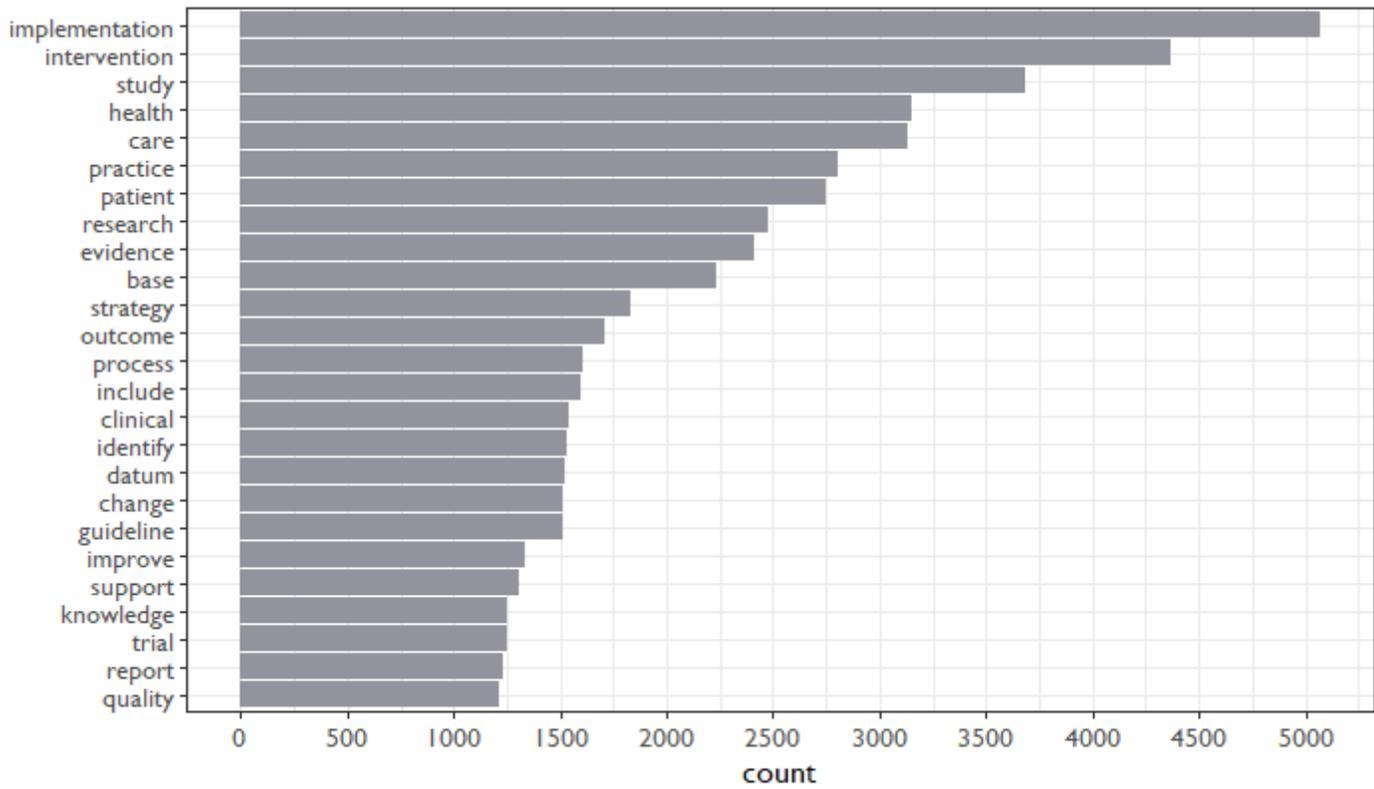


Figure 3

Most Frequently Occurring Terms in Implementation Science Abstracts.

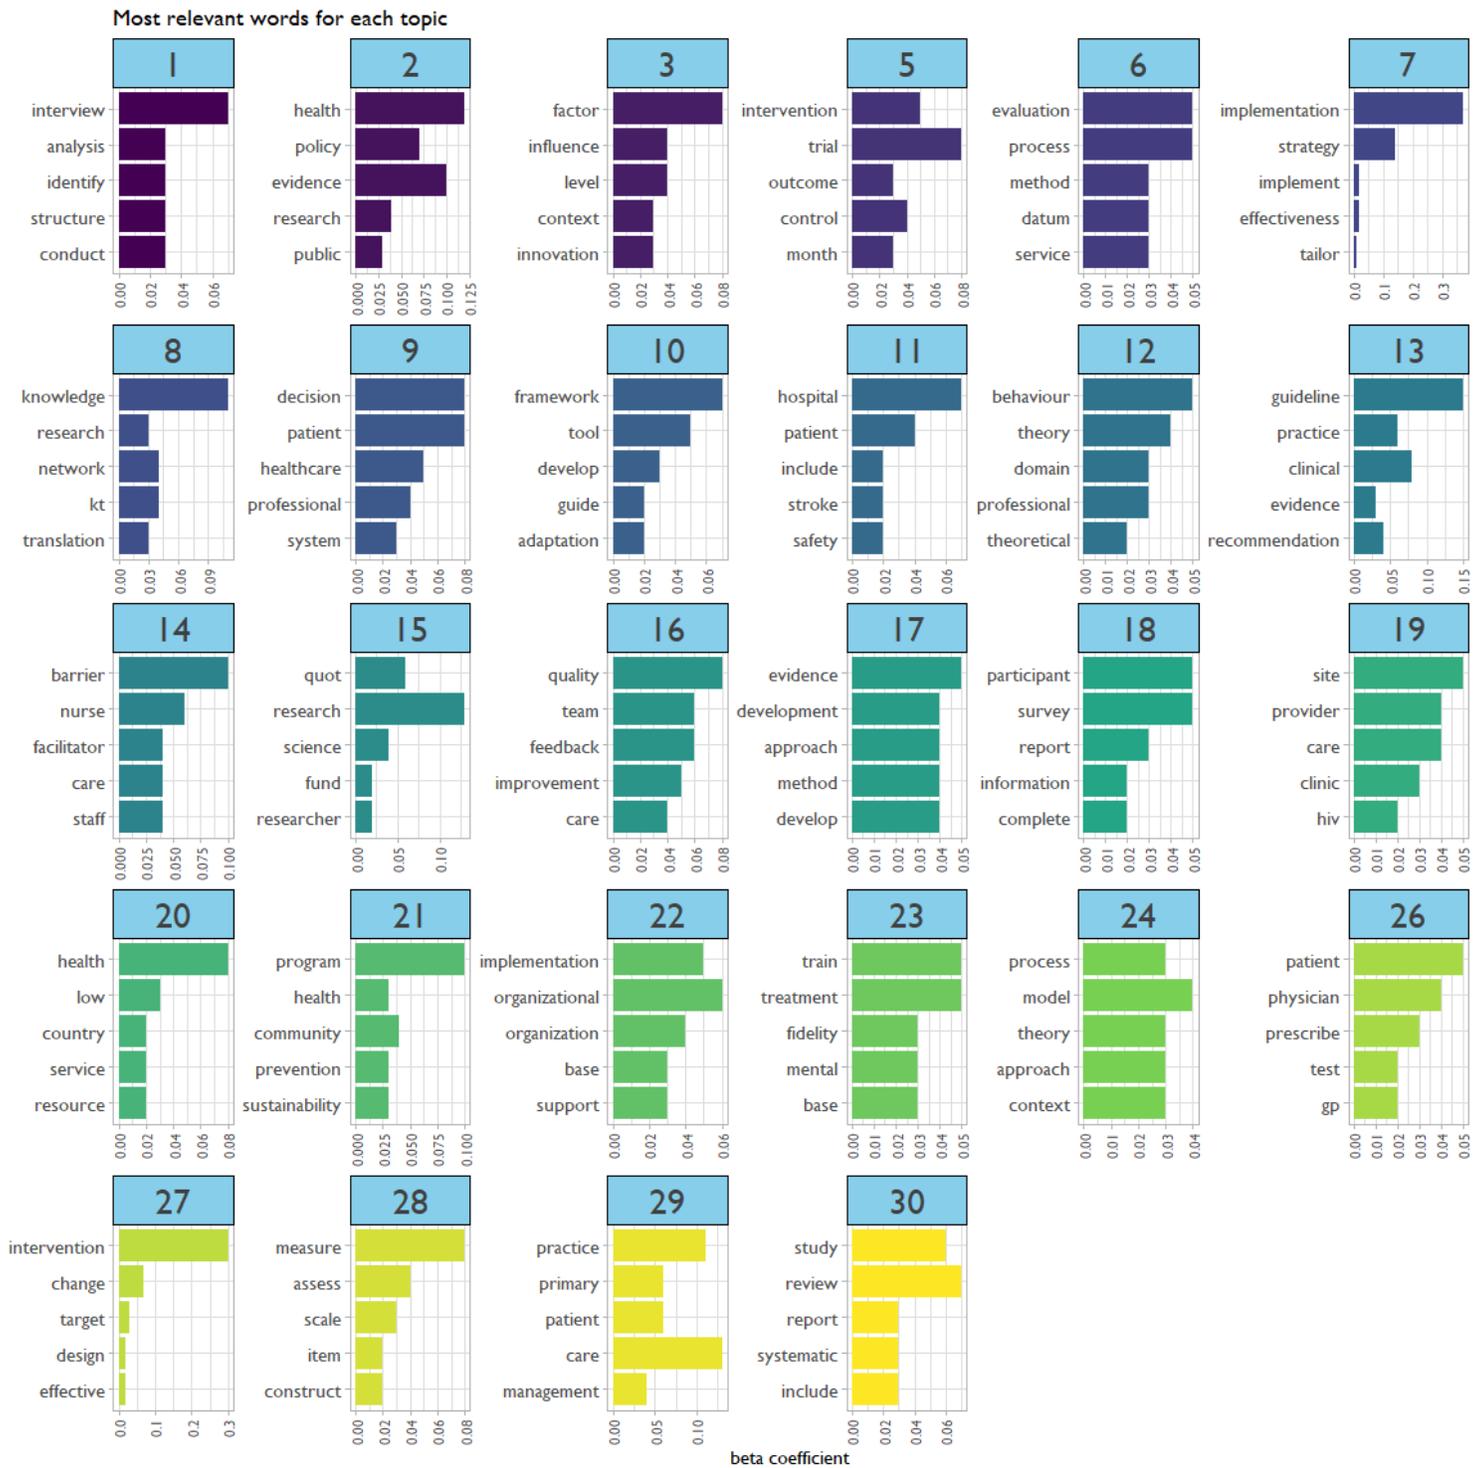


Figure 4

Top 5 Word Associations across Implementation Science Abstract Topic Clusters.

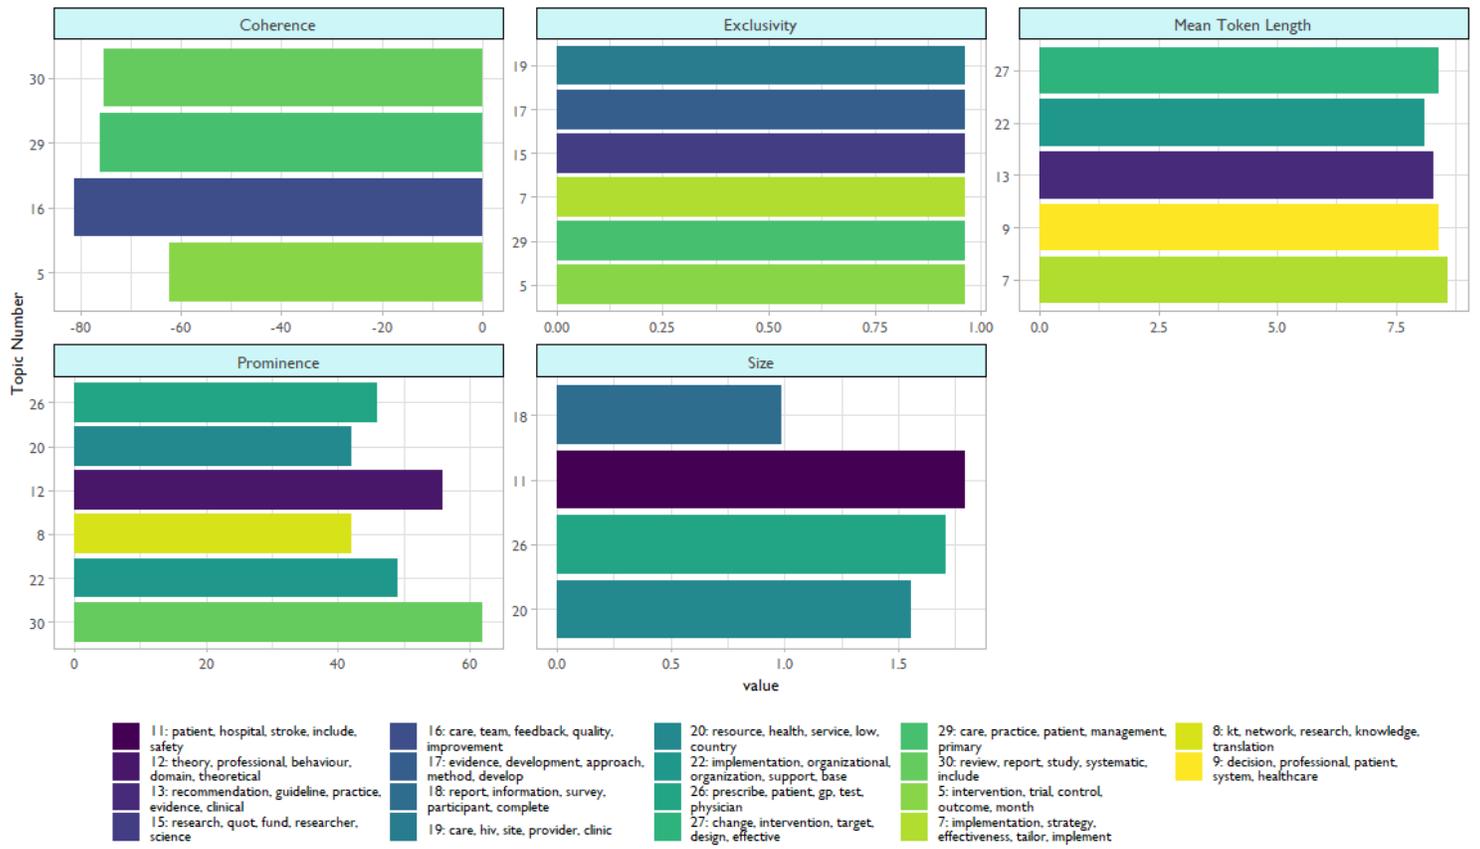


Figure 5

Topic Clusters Across the five latent Dirichlet allocation (LDA) Statistics

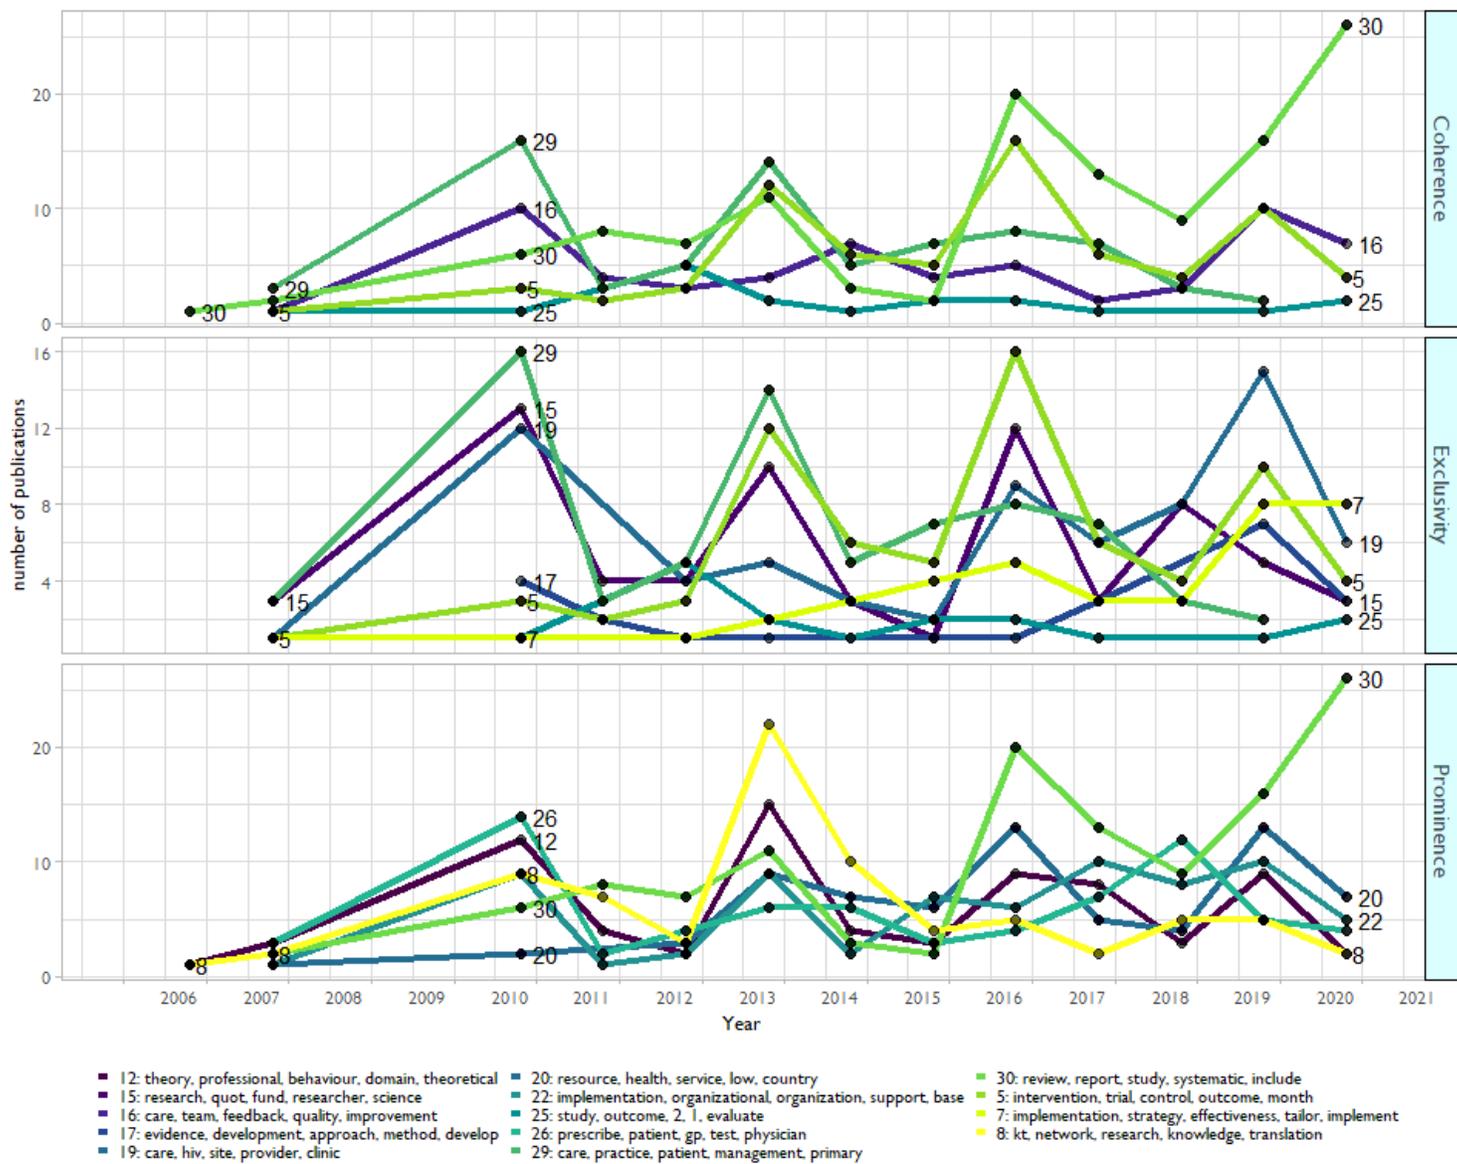


Figure 6

Changes in Topics over time.