



Scholarly literature mining with information retrieval and natural language processing: Preface

Guillaume Cabanac¹ · Ingo Frommholz² · Philipp Mayr³

Received: 9 October 2020 / Published online: 17 November 2020
© Akadémiai Kiadó, Budapest, Hungary 2020

Introduction

This special issue features the work of authors originally coming from different communities: bibliometrics/scientometrics (SCIM), information retrieval (IR) and, as an emerging player gaining more relevance for both aforementioned fields, natural language processing (NLP). The work presented in their papers combine ideas from all these fields, having in common that they all are using the scholarly data well known in scientometrics and solving problems typical to scientometric research. They model and mine citations, as well as metadata of bibliographic records (authorships, titles, abstracts sometimes), which is common practice in SCIM. They also mine and process fulltexts (including in-text references and equations) which is common practice in IR and requires established NLP text mining techniques. IR collections are utilised to ensure reproducible evaluations; creating and sharing test collections in evaluation initiatives such as CLEF eHealth¹ is common IR tradition that is also prominent in NLP, eg., by the CL-SciSumm shared task.²

From an IR perspective, surprisingly, scholarly information retrieval and recommendation, though gaining momentum, have not always been the focus of research in the past. Besides operating on a rich set of data for researchers in all three disciplines to play with, scholarly search poses challenges in particular for IR due to the complex information needs that require different approaches than known from, e.g., Web search, where information needs are simpler in many cases. As an example, the current COVID-19 crisis shows that hybrid SCIM/IR/NLP approaches are increasingly required to ensure researchers get access

¹ <https://clefehealth.imag.fr>.

² <https://github.com/WING-NUS/scisumm-corpus>.

✉ Guillaume Cabanac
guillaume.cabanac@univ-tlse3.fr

Ingo Frommholz
ifrommholz@acm.org

Philipp Mayr
philipp.mayr@gesis.org

¹ Computer Science Department, IRIT UMR 5505 CNRS, University of Toulouse, 118 Route de Narbonne, 31062 Toulouse Cedex 9, France

² University of Bedfordshire, Luton LU1 3JU, UK

³ GESIS – Leibniz Institute for the Social Sciences, Cologne, Germany

to important relevant and high-quality information, often only available on preprint servers, in a short period of time (Brainard 2020; Fraser et al. 2020; Kwon 2020; Palayew et al. 2020). These kinds of complex information needs pose challenges which have been recognised by the Information Retrieval community that quickly launched the TREC-COVID initiative run by NIST (Roberts et al. 2020), demonstrating the timeliness of our endeavour and this special issue. Working on scholarly material thus has incentives for researchers in Information Retrieval but we believe the challenges can only be tackled effectively by all three communities as a whole. The NLP community has initiated a similar activity with a dedicated workshop series NLP COVID-19 Workshop³ which is running at major NLP conferences (ACL & EMNLP) in 2020.

With the surge of “scholarly big data” (Giles 2013), Bibliometrics and Information Retrieval in combination with NLP methods have seen a recent renaissance that resulted in a series of special issues:

- “Combining Bibliometrics and Information Retrieval” (Mayr and Scharnhorst 2015) in *Scientometrics* (2015).
- “Bibliometric-enhanced Information Retrieval” (Cabanac et al. 2018) in *Scientometrics* (2018).
- “Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries” (Mayr et al. 2018) in *International Journal on Digital Libraries* (2018).
- “Mining Scientific Papers: NLP-enhanced Bibliometrics” (Atanassova et al. 2019) in *Frontiers in Research Metrics and Analytics* (2019).

Special issue papers

This special issue on “Scholarly literature mining with Information Retrieval and Natural Language Processing” presents works intersecting Bibliometrics and Information Retrieval, utilising Natural Language Processing (NLP). The special issue was announced via an open call for papers⁴. In response to the CFP, we received 24 submissions which were reviewed by two to three reviewers (for overlapping papers, eg., IR and NLP, we selected reviewers from both domains). Eventually, the guest editors accepted 14 papers. Nine papers have been rejected and one paper was withdrawn by the authors during the reviewing rounds.

In the following we provide an overview of the 14 papers organised into 3 clusters. We introduce the paper ordering of the special issue in Table 1. To generate a lightweight overview of the variety of the papers we identified the research *Tasks* and *Area of Application*, the used *Corpus*, *Objects*, and *Methods* of each contribution.

The papers in this special issue appear in the following sequence. We decided to start with a set of more classical papers featuring scientometric methods like network analysis and bibliographic data from the Web of Science, Scopus or similar resources. The second set of papers is more IR oriented: papers mine fulltexts and they use techniques like

³ <https://www.nlpcovid19workshop.org/acl2020/>.

⁴ <https://sites.google.com/view/scientometrics-si2019-bir>.

Table 1 Overview of the articles in this special issue

Task	Area of application	Corpus	Objects	Methods
Lietz				
Field delineation	Social network science	Web of science	Metadata (title, abstract, keywords), references	Clustering, network analysis
Schneider, Ye, Hill, & Whitehorn				
Analysing citing papers of a retracted study	Clinical science	Google scholar, web of science	Seed paper, citations, retraction notices	Network analysis, citation context analysis, retraction status visibility analysis
Kreutz, Sahitaj, & Schenkel				
Spotting seminal work; classifying papers	Computer science	DBLP	Fulltext	Classification using words, semantics, topics and publication years
Haunschild & Marx				
Spotting seminal work	Physics	Microsoft academic, web of science	References, time	Reference publication year spectroscopy
Lamirel, Chen, Cuxac, Al Shehabi, Dugué & Liu				
Mapping the evolution of a country's scientific production	Science in China	China national knowledge infrastructure database	metadata (title, abstract, authors), dictionary of Chinese names	Clustering, topic modeling, network analysis
Nogueira, Jiang, Cho, & Lin				
Ranking citation recommendations	Computer science, biomedicine	DBLP, open research, PubMed	Fulltext	Document ranking model, embeddings
Greiner-Petter, Youssef, Ruas, Miller, Schubotz, Aizawa & Gipp				
Discovering mathematical term similarity and analogy and query expansions	Mathematics	arXiv	Fulltext	Embeddings
Carvalho, Parra, Lobel, & Soto				
Paper screening for evidence-based medicine	Medicine	CLEF eHealth, Epistemonekos	Fulltext	Document ranking model, query expansion, embeddings
Saier & Färber				
Dataset creation	Fields of arXiv preprints	arXiv, Microsoft academic graph	Fulltext, in-text citations, linked data	Data integration, descriptive statistics

Table 1 (continued)

Task	Area of application	Corpus	Objects	Methods
Zerva, Nghiem, Nguyen, & Ananiadou				
Paper summarization (from citations)	Natural language processing	CL-SciSumm	Fulltext, in-text citations	Neural networks
La Quatra, Cagliero, & Baralis				
Discourse facet summarization	Natural language processing	CL-SciSumm	Fulltext, in-text citations	Neural networks
AbuRa'ed, Saggion, Shvets, & Bravo				
Citation sentence production	Text summarization	ScisummNet, Open academic graph, microsoft academic graph, RWSData	Fulltext	Neural networks
Jimenez, Avila, Dueñas, & Gelbukh				
Citation forecasting	The scientific literature	Scopus	Metadata (title + abstract)	Statistics, stylometry
Portenoy & West				
Generation of a literature review of a field	Community detection in graphs, misinformation studies, science communication	Web of science	References, paper titles	Text similarity, supervised learning, embeddings

embeddings and neural networks. The third cluster of papers contains NLP-oriented papers which are, for instance, specialised in summarisation and utilise scholarly documents.

Cluster 1. *SCIM with IR and NLP*

- Lietz: *Drawing impossible boundaries: field delineation of Social Network Science.*
- Schneider et al.: *Continued post-retraction citation of a fraudulent clinical trial report, eleven years after it was retracted for falsifying data.*
- Kreutz et al.: *Evaluating semantometrics from computer science publications.*
- Haunschild & Marx: *Discovering seminal works with marker papers.*
- Lamirel et al.: *An overview of the history of Science of Science in China based on the use of bibliographic and citation data: a new method of analysis based on clustering with feature maximization and contrast graphs.*

Cluster 2. *IR and Text-mining of scholarly literature*

- Nogueira et al.: *Navigation-based candidate expansion and pretrained language models for citation recommendation.*
Greiner-Petter et al.: *Math-word embedding in math search and semantic extraction.*
- Carvallo et al.: *Automatic document screening of medical literature using word and text embeddings in an active learning setting.*
- Saier & Färber: *unarXive: a large scholarly data set with publications' full-text, annotated in-text citations, and links to metadata.*

Cluster 3. NLP-oriented papers on scholarly literature

- Zerva et al.: *Cited text span identification for scientific summarisation using pre-trained encoders.*
- La Quatra et al.: *Exploiting pivot words to classify and summarize discourse facets of scientific papers.*
- AbuRa'ed et al.: *Automatic related work section generation: experiments in scientific document abstracting.*
- Jimenez et al.: *Automatic prediction of citability of scientific articles by stylometry of their titles and abstracts.*
- Portenoy & West: *Constructing and evaluating automated literature review systems.*

We hope the selection of papers in this special issue will be interesting and enjoyable for researchers coming from all relevant fields and provides a starting point for future explorations in the field.⁵

Acknowledgements We wish to thank all contributors to this special issue: The researchers who submitted papers, the many reviewers who generously offered their time and expertise, and the participants of the BIR and BIRNDL workshops (Cabanac et al. 2020).

References

- Atanassova, I., Bertin, M., & Mayr, P. (2019). Editorial: mining scientific papers: NLP-enhanced bibliometrics. *Frontiers in Research Metrics and Analytics*, <https://doi.org/10.3389/frma.2019.00002>.
- Brainard, J. (2020). New tools aim to tame pandemic paper tsunami. *Science*, 368(6494), 924–925. <https://doi.org/10.1126/science.368.6494.924>.
- Cabanac, G., Frommholz, I., & Mayr, P. (2018). Bibliometric-enhanced information retrieval: Preface. *Scientometrics*, 116(2), 1225–1227. <https://doi.org/10.1007/s11192-018-2861-0>.
- Cabanac, G., Frommholz, I., & Mayr, P. (2020). Bibliometric-Enhanced Information Retrieval 10th Anniversary Workshop Edition. In J. M. Jose, E. Yilmaz, J. Magalhães, P. Castells, N. Ferro, M. J. Silva, & F. Martins (Eds.), *Advances in Information Retrieval, LNCS*, (Vol. 12036, pp. 641–647). Berlin: Springer International Publishing. https://doi.org/10.1007/978-3-030-45442-5_85.
- Fraser, N., Brierley, L., Dey, G., Polka, J.K., Pálffy, M., Nanni, F., & Coates, J.A. (2020). Preprinting the COVID-19 pandemic. *bioRxiv*. <https://doi.org/10.1101/2020.05.22.111294>
- Giles, C.L. (2013). Scholarly big data. In *CIKM'13: Proceedings of the 22nd ACM international conference on conference on information and knowledge management*, p. 1. ACM, New York, NY. <https://doi.org/10.1145/2505515.2527109>
- Kwon, D. (2020). How swamped preprint servers are blocking bad coronavirus research. *Nature*, 581(7807), 130–131. <https://doi.org/10.1038/d41586-020-01394-6>.
- Mayr, P., Frommholz, I., Cabanac, G., Chandrasekaran, M. K., Jaidka, K., Kan, M. Y., et al. (2018). Introduction to the special issue on bibliometric-enhanced information retrieval and natural language processing for digital libraries (BIRNDL). *International Journal on Digital Libraries*, 19(2–3), 107–111. <https://doi.org/10.1007/s00799-017-0230-x>.
- Mayr, P., & Scharnhorst, A. (2015). Combining bibliometrics and information retrieval: Preface. *Scientometrics*, 102(3), 2191–2192. <https://doi.org/10.1007/s11192-015-1529-2>.
- Palayew, A., Norgaard, O., Safreed-Harmon, K., Andersen, T. H., Rasmussen, L. N., & Lazarus, J. V. (2020). Pandemic publishing poses a new COVID-19 challenge [Comment]. *Nature Human Behaviour*, 4(7), 666–669. <https://doi.org/10.1038/s41562-020-0911-0>.

⁵ Since 2016 we maintain the “Bibliometric-enhanced-IR Bibliography” https://github.com/PhilippMayr/Bibliometric-enhanced-IR_Bibliography/ that collects scientific papers which appear in collaboration with the BIR/BIRNDL organizers.

Roberts, K., Alam, T., Bedrick, S., Demner-Fushman, D., Lo, K., Soboroff, I., et al. (2020). TREC-COVID: Rationale and structure of an information retrieval shared task for COVID-19. *Journal of the American Medical Informatics Association*, 27(9), 1431–1436. <https://doi.org/10.1093/jamia/ocaa091>.