

Report on the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL 2018)

Philipp Mayr

GESIS – Leibniz Institute for the Social Sciences, Germany

philipp.mayr@gesis.org

Muthu Kumar Chandrasekaran

NUS School of Computing, Singapore

cmkumar087@gmail.com

Kokil Jaidka

University of Pennsylvania, USA

jaidka@sas.upenn.edu

Abstract

The 3rd joint BIRNDL workshop was held at the 41st ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2018) in Ann Arbor, USA. BIRNDL 2018 intended to stimulate IR researchers and digital library professionals to elaborate on new approaches in natural language processing, information retrieval, scientometrics, and recommendation techniques that can advance the state-of-the-art in scholarly document understanding, analysis, and retrieval at scale. The workshop incorporated three paper sessions and the 4th edition of the CL-SciSumm Shared Task.

1 Introduction

The goal of the BIRNDL workshop at SIGIR is to engage the information retrieval (IR), natural language processing (NLP), digital libraries, bibliometrics and scientometrics communities to advance the state-of-the-art in scholarly document understanding, analysis and search and retrieval at scale [10]. Scholarly documents are indexed by large, cross-domain digital repositories such as the ACL Anthology, ArXiv, ACM Digital Library, IEEE database, Web of Science, Google Scholar and Semantic Scholar. Currently, digital libraries collect and allow access to papers and their metadata — including citations — but mostly do not analyze the items they index. The scale of scholarly publications poses a challenge for scholars in their search for relevant literature. Information seeking and sensemaking from the large body of scholarly literature is the key theme of BIRNDL and sets the agenda for tools and approaches to be discussed and evaluated at the workshop.

Papers at the 3rd BIRNDL workshop incorporate insights from IR, bibliometrics and NLP to develop new techniques to address the open problems such as evidence-based searching, measurement of research quality, relevance and impact, the emergence and decline of research

problems, identification of scholarly relationships and influences and applied problems such as language translation, question-answering and summarization. We also address the need for established, standardized baselines, evaluation metrics and test collections. Towards the purpose of evaluating tools and technologies developed for digital libraries, we are organizing the 4th CL-SciSumm Shared Task based on the CL-SciSumm corpus, which comprises over 500 computational linguistics (CL) research papers, interlinked through a citation network.

2 Overview of the papers

This year six full papers were submitted to the workshop, three of which were finally accepted as full papers for presentation and inclusion in the proceedings. In addition three poster papers were accepted for inclusion in the proceedings. The workshop featured one keynote talk, one full paper session, one session with presentations of systems participating in the CL-SciSumm Shared Task (see the CL-SciSumm overview paper [6]) and a poster session. The following section briefly describes the keynote and sessions.

2.1 Keynote

Byron Wallace gave an inspiring keynote on “Automating Biomedical Evidence Synthesis: Recent Work and Directions Forward” [15]. He talked about recent progress in biomedical evidence synthesis and called on the NLP, IR and machine learning communities to take up the challenges that remain unaddressed in this critical field. He said that the field puts forth technically challenging problems of interest such as, building models with low-supervision, joint inference and extraction over long documents and hybrid crowd-expert annotations, to the aforementioned technical communities.

2.2 Research papers

Alzogbi [2] presented a time-aware recommender system (Time-aware Collaborative Topic Regression - T-CTR) that accounts for the concept-drift in user interest by computing user-specific concept drift scores. The paper considered the use case of scientific papers recommendation and conducted experiments on data from citeulike. Results showed the superiority of the time-aware recommendation system T-CTR over the state-of-the-art systems.

Shinoda and Aizawa [14] proposed an unsupervised query-based summarization of scientific papers. Importance scores for words calculated from word embeddings trained on an auxiliary corpus are used to compute sentence vectors. Finally, a random walk is performed on sentences which leverages distributional similarities between query terms and words in the sentence, as well as the similarities between pairs of sentences.

Brochier et al. [4] applied a new document-query methodology to evaluate experts retrieval from a set of queries sampled directly from the experts documents. They provided a formal definition of the expert finding task and worked on a topic-query and a document-query evaluation protocol. They performed a detailed evaluation with three baseline expert recommender algorithms on two AMiner expert data sets.

2.3 Poster papers

Scharpf et al. [13] proposed a Wikidata markup to link semantic elements of a mathematical formula in MathML to Wikidata items. They suggested Formula Concept Discovery as a concept to develop automatic retrieval functions (e.g. formula clustering) on labeled full text corpora. They argued in favor of a larger MathML benchmark for evaluation purposes.

Jia and Saule [7] proposed “Keyphraser” to alleviate the “over-generation error” when extracting key phrases from scientific documents. KeyPhraser is an unsupervised method for identifying document phrases using features such as concordance, popularity, informativeness and position of the first occurrence of the phrase.

Luan et al. [9] presented their SemEval-2018 Task 7 system Semantic Relation Extraction and Classification in Scientific Papers as an invited poster. They were invited due to the work’s close relevance to the workshop and the authors’ interest.¹

2.4 CL-SciSumm

We hosted the 4th Computational Linguistics Scientific Summarization Shared Task, sponsored by Microsoft Research Asia as part of the BIRNDL workshop. The Shared Task is aimed at creation of an open corpus for citation based faceted summarization of scientific documents and evaluation of systems over three sub-tasks to output a summary. This is the first medium-scale shared task on scientific document summarization in the computational linguistics (CL) domain. The task and its corpus have the potential to spur further interest in related problems in scientific discourse mining, such as citation analysis, query-focused question answering and text reuse.

We have been incrementally building our annotated corpora over the four editions of CL-SciSumm. CL-SciSumm’18 systems were provided with 50 document sets for training and evaluation was run over 10 test document sets. Eleven teams registered and ten teams participated in this year’s shared task, on a corpus that was 33% larger than the 2017 corpus. A total of 60 runs were evaluated for Task 1, and 33 runs were evaluated for Task 2. For Task 1A, on using sentence overlap (F1 score) as the metric, the best performance was achieved by using voting methods on the results of a random forest classifier trained on lexical features and word embeddings, by NUDT [16]. When ROUGE-based F1 is used as a metric, the best performance was also based on a BM25 voting mechanism by Klick Labs [3]. The best performance in Task 1B was using kernel-based methods by CIST [8].

For Task 2, the team from LaSTUS/TALN+INCO had the best performance using CNN to learn the relation between a sentence and a scoring value indicating its relevance, and choosing the most relevant sentences for the summary [1].

Table 1 summarizes the best results, together with information about the corpus size in the different years. It also mentions the best performing approaches for Task 1. We observe that although the corpus size has doubled over the years, there is no perceptible change in the best scores. It should be noted that because of the difference in corpus sizes, the numbers are perhaps not comparable – model overfitting has surely decreased since the efforts in 2016. We believe that to observe a real spurt in progress, there is a need to try to automatically scale dataset development and achieve $O(N) = 1000$ for any significant nudge in results, and we invite other researchers to collaborate with us on this long term goal.

¹BIRNDL organising committee and [9] thank the BIRNDL reviewers for reviewing this accepted SemEval paper without prejudice.

	Year		
	2016	2017	2018
Corpus size	30	40	60
Task 1A results			
F1-score	0.13	0.12	0.14
Best performing approach	Tf-idf	Weighted ensemble	CNN
Task 1B results			
F-1 score	0.65	0.40	0.38
Task 2 (Rouge-2 scores)			
vs. Abstracts	0.68	0.35	0.33
vs. Humans	0.22	0.27	0.25
vs. Community	0.25	0.20	0.21

Table 1: Task 1A, 1B and Task 2 scores and best performing approaches over the years.

Based on the trends in system performance, we anticipate that lexical methods complemented with domain-specific word embeddings in a deep learning framework would be an ideal approach for a scientific summarization task such as ours. For a detailed breakdown of the results, please see the CL-SciSumm’18 overview paper [6]. It is inspiring that each year, we have had new research teams participating, including participants from the industry in 2018. Many other teams have consistently participated over the years and patiently worked with us to improve the quality of our corpus, and we are glad to have their support as well as to support their research with our efforts.

3 Outlook and further reading

With this continuing workshop series we have built up a sequence of explorations, visions, results documented in scholarly discourse, and created a sustainable bridge between bibliometrics, IR and NLP. We see the community still growing.

We will continue to organize similar workshops at IR, DL, Scientometric, NLP and CL premier venues. The combination of research paper presentations, and a shared task like CL-SciSumm with system evaluation has proven to be a successful agile format. We propose to continue with this format. CL-SciSumm ’19 will see a major expansion in the training data size.

In 2015 we published a first special issue on “Combining Bibliometrics and Information Retrieval” in the *Scientometrics* journal [12]. Recently, a special issue on “Bibliometrics, Information Retrieval and Natural Language Processing in Digital Libraries” appeared in September 2018 in the *International Journal on Digital Libraries*, see an overview in [11]. Another special issue on “Bibliometric-enhanced Information Retrieval and Scientometrics” appeared in *Scientometrics*, see editorial [5]. Since 2016 we maintain the “Bibliometric-

enhanced-IR Bibliography”² that collects scientific papers which appear in collaboration with the BIR/BIRNDL organizers. We invite interested researchers to join this project and contribute related publications.

Acknowledgments

We thank Microsoft Research Asia for their generous support in funding the development, dissemination and organization of the CL-SciSumm dataset and the Shared Task³. We are also grateful to the co-organizers of the 1st BIRNDL workshop - Guillaume Cabanac, Ingo Frommholz, Min-Yen Kan and Dietmar Wolfram, for their continued support and involvement. Finally we thank our programme committee members who did an excellent reviewing job. All PC members are documented on the BIRNDL website⁴.

References

- [1] A. Abura’ed, A. Bravo, L. Chiruzzo, and H. Saggion1. LaSTUS/TALN+INCO @ CL-SciSumm 2018 - Using Regression and Convolutions for Cross-document Semantic Linking and Summarization of Scholarly Literature. In *Proceedings of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL2018)*, Ann Arbor, Michigan, July 2018.
- [2] A. Alzogbi. Time-aware collaborative topic regression: Towards higher relevance in textual item recommendation. In *Proc. of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL)*, 2018.
- [3] G. Baruah and M. Kolla. Klick Labs at CL-SciSumm 2018. In *Proceedings of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL2018)*, Ann Arbor, Michigan, July 2018.
- [4] R. Brochier, A. Guille, B. Rothan, and J. Velcin. Impact of the query set on the evaluation of expert finding systems. In *Proc. of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL)*, 2018.
- [5] G. Cabanac, I. Frommholz, and P. Mayr. Bibliometric-enhanced information retrieval: preface. *Scientometrics*, 116(2):1225–1227, 2018.
- [6] K. Jaidka, M. Yasunaga, M. K. Chandrasekaran, D. Radev, and M.-Y. Kan. The cl-scisumm shared task 2018: Results and key insights. In *Proc. of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL)*, 2018.
- [7] H. Jia and E. Saule. Addressing overgeneration error: An effective and efficient approach to keyphrase extraction from scientific papers. In *Proc. of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL)*, 2018.

²https://github.com/PhilippMayr/Bibliometric-enhanced-IR_Bibliography/

³<http://wing.comp.nus.edu.sg/cl-scisumm2018/>

⁴<http://wing.comp.nus.edu.sg/birndl-sigir2018/>

-
- [8] L. Li, J. Chi, M. Chen, Z. Huang, Y. Zhu, and X. Fu. CIST@CLSciSumm-18: Methods for Computational Linguistics Scientific Citation Linkage, Facet Classification and Summarization. In *Proceedings of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL2018)*, Ann Arbor, Michigan, July 2018.
- [9] Y. Luan, M. Ostendorf, and H. Hajishirzi. The UWNLP system at SemEval-2018 Task 7: Neural Relation Extraction Model with Selectively Incorporated Concept Embeddings. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 788–792. ACL, 2018.
- [10] P. Mayr, M. K. Chandrasekaran, and K. Jaidka. Report on the 2nd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL 2017). *SIGIR Forum*, 51(3):107–113, 2017.
- [11] P. Mayr, I. Frommholz, G. Cabanac, M. K. Chandrasekaran, K. Jaidka, M.-Y. Kan, and D. Wolfram. 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, 2018.
- [12] P. Mayr and A. Scharnhorst. Scientometrics and information retrieval: weak-links revitalized. *Scientometrics*, 102(3):2193–2199, 2015.
- [13] P. Scharpf, M. Schubotz, and B. Gipp. Representing mathematical formulae in content mathml using wikidata. In *Proc. of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL)*, 2018.
- [14] K. Shinoda and A. Aizawa. Query-focused scientific paper summarization with localized sentence representation. In *Proc. of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL)*, 2018.
- [15] B. Wallace. Automating biomedical evidence synthesis: Recent work and directions forward. In *Proc. of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL)*, 2018.
- [16] P. Wang, S. Li, T. Wang, H. Zhou, and J. Tang. NUDT @ CLSciSumm-18. In *Proceedings of the 3rd Joint Workshop on Bibliometric-enhanced Information Retrieval and Natural Language Processing for Digital Libraries (BIRNDL2018)*, Ann Arbor, Michigan, July 2018.