



# NLPExplorer: Exploring the Universe of NLP Papers

Monarch Parmar, Naman Jain, Pranjali Jain, P. Jayakrishna Sahit, Soham Pachpande, Shruti Singh, and Mayank Singh<sup>(✉)</sup>

Indian Institute of Technology Gandhinagar, Gandhinagar, Gujarat, India  
[singh.mayank@iitgn.ac.in](mailto:singh.mayank@iitgn.ac.in)

**Abstract.** Understanding the current research trends, problems, and their innovative solutions remains a bottleneck due to the ever-increasing volume of scientific articles. In this paper, we propose NLPExplorer, a completely automatic portal for indexing, searching, and visualizing Natural Language Processing (NLP) research volume. NLPExplorer presents interesting insights from papers, authors, venues, and topics. In contrast to previous topic modelling based approaches, we manually curate five course-grained non-exclusive topical categories namely *Linguistic Target* (Syntax, Discourse, etc.), *Tasks* (Tagging, Summarization, etc.), *Approaches* (unsupervised, supervised, etc.), *Languages* (English, Chinese, etc.) and *Dataset types* (news, clinical notes, etc.). Some of the novel features include a list of young popular authors, popular URLs and datasets, list of topically diverse papers and recent popular papers. Also, it provides temporal statistics such as yearwise popularity of topics, datasets, and seminal papers. To facilitate future research and system development, we make all the processed dataset accessible through API calls. The current system is available at <http://nlpexplorer.org>.

**Keywords:** Natural language processing · Search · Research papers

## 1 Introduction

Effective scientific literature understanding plays a critical role towards the research community's common goal of "*March for Science*". However, the yearly generated research volume shows an upward trend with an estimated increase of 8–9% per year. This results in information overload, impacting the literature review process and often, leads to '*reinventing the wheel*' syndrome. This negatively affects the efficiency of scientific progress on the knowledge frontier.

In the past, significant efforts have been made to curate peer-reviewed open access scientific information. In particular, the field of Natural Language Processing (NLP) witnessed the development of *ACL Anthology* since the year 2001, which curates papers (in PDF format) from more than 70 NLP venues

---

M. Parmar, N. Jain, P. Jain, P. Jayakrishna Sahit and S. Pachpande—Alphabetically ordered with equal contribution

including popular conferences like ACL, NAACL, EMNLP, etc. In the year 2008, Bird et al. [2] released the *ACL Anthology Reference Corpus (ACL ARC)*, consisting of OCR'd extracted text and metadata of PDF articles. In the year 2009, Radev et al. [8] developed *ACL Anthology Network (AAN)* by manually constructing paper citation, author citation and author collaboration network, along with other interesting statistics and citation summaries of the articles. The recently updated AAN system [9] indexes articles published till 2014. Another initiative, ACL Anthology SearchBench [10], provides bibliographic metadata filtering and full-text structured semantic search in ACL Anthology. CL Scholar [12] periodically crawls ACL Anthology and constructs a computational linguistic knowledge graph. It supports natural language queries and entity specific queries about the authors, venue, and paper. Almost all systems described above are either in dormant condition [2, 8–10] or not available [12].

In this paper, we discuss the development of NLPEXPLORER. NLPEXPLORER provides paper, author, venue and topic-specific temporal as well as aggregated statistics. For the first time, we also showcase, the usage of URLs over the years, popular top-level and sub-domains, survey papers, new authors, etc. The system also presents a visualization of the timeline for the first occurrences of sub-topics.

Our main contributions are as follows: (i) We periodically download, preprocess, and index the ACL Anthology dataset consisting of more than 55 thousand full-text PDF articles. (ii) We invest extensive effort in automatic curation, structured information extraction, cleaning, indexing, and other related preprocessing steps. (iii) We classify papers into a first-of-its-kind detailed list of NLP topics and subtopics. (iv) The proposed system presents content as well as bibliographic statistics along with basic keyword-based search facilities. (v) We deploy the current system in Google Cloud with an API-based retrieval facility.

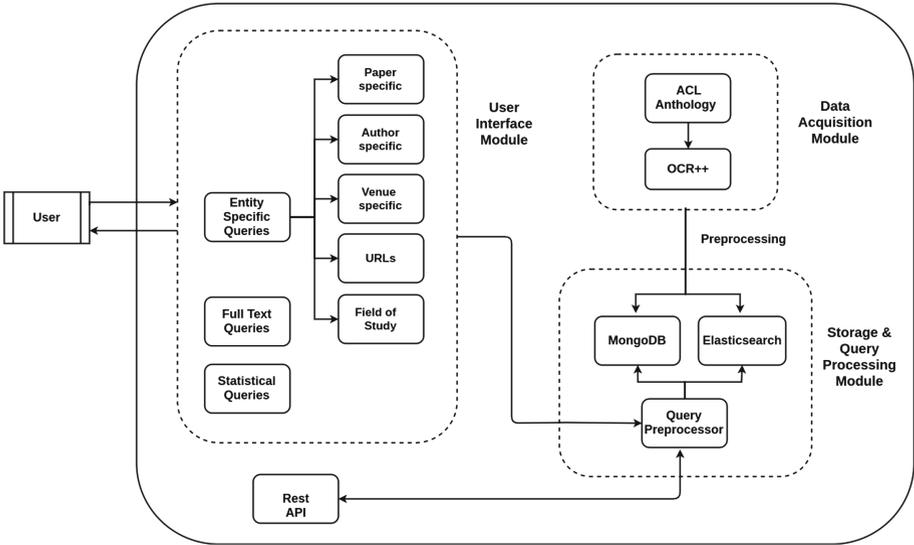
## 2 Dataset

We leverage the ACL Anthology dataset [1] that hosts articles dedicated to Computational Linguistics and Natural Language Processing. Each article is present as a PDF file along with the metadata namely author list, venue name, year of publication and a unique eight-character identifier. Overall, the dataset has 55,565 papers, 39,555 unique authors, 78 venues and 723,976 citations.

## 3 Architecture

The architecture of NLPEXPLORER is composed of three interdependent modules, (i) data acquisition, (ii) storage and query processing, and (iii) user interface module. Figure 1 presents the detailed description of the architecture. The *data acquisition* module periodically curates newly added semi-structured information such as metadata and the corresponding PDF articles from ACL Anthology. It leverages OCR++ [11] tool to extract structural and bibliographical information from each of the PDF articles, and then passes the metadata, structural and bibliographical information to the storage and query processing module.

The *storage and query processing* module indexes the periodical updates into MongoDB [5] database and Elasticsearch [3] engine. MongoDB database handles the basic storage and retrieval, whereas Elasticsearch supports full-text based query search and ranking. The third module, *the user interface*, fetches processed search results from the storage and query processing module. The user interface is designed using Python’s Flask library, and JavaScript libraries Plotly [7] and Timeline [13] are used to render statistics and graphical components of the interface.



**Fig. 1.** Architecture of NLPEXPLORER. Arrows represent data flow.

The current system is deployed at Google cloud [4]. The infrastructure consists of a 4 CPU - 7.5GB RAM, Linux VM instance that can be extended on demand. The system supports REST API requests served using MongoDB database.

## 4 Current Features

We categorize the wide range of functionalities of NLPEXPLORER as follows:

- **Entity-specific & Full-Text search:** NLPEXPLORER supports basic keyword based search leveraging the metadata and the full-text of articles. The search can be chosen to output results to any of the following domains - authors, papers, venues, URLs, and field of study. We also extend full-text search of articles in order to visualise n-gram trends over the period of time.

- **Paper statistics:** We provide standard paper related statistics such as the publication year and venue, author information, citation distribution over the years and the link to the corresponding PDF article. Additionally, we provide interesting insights like similar papers, topical distribution and mentioned URLs. We also provide statistics such as the list of popular recent papers, popular survey papers, seminal papers, papers with diverse topics, and publication count of last five years.
- **Author statistics:** We provide author statistics such as publication and citation distribution over the years, topical distribution of papers and venue preference. Additionally, we provide the list of popular recent authors, popular authors in the lifetime of ACL Anthology, authors with top publication counts, recent authors with high publication counts, and highly diverse authors.
- **Field of study statistics:** We curate a comprehensive list of topics categorized into five broad categories, (i) *Linguistic Targets* (Syntax, Discourse, Pragmatics, etc.), (ii) *Tasks* (Tagging, Summarization, Chunking, etc.), (iii) *Approaches* (unsupervised, supervised, etc.), (iv) *Languages* (English, Chinese, Hindi, etc.), and (v) *Dataset Types* (news, clinical notes, etc.). NLPEXPLORER provides interesting insights such as the temporal distribution of papers and authors in each topic and subtopics, distribution of papers in each conference, and a timeline to visualize the introduction of new topics. We make the list of topics and the processed data of the system available at the systems’ webpage [6].
- **Venue statistics:** Venue-related statistics include temporal distribution of publications and citations, topical distribution, and the list of papers in a year. We also include insights such as the top NLP venues citing and cited by the candidate venue, popular authors publishing in the candidate venue, and the shift in topical distribution over the years.
- **URL statistics:** We analyse URLs reported in the research papers. The URL-related statistics include top URLs in different categories such as universities, digital libraries, datasets and research groups, alongwith the analysis of top-level domains (TLDs) and corresponding sub-domains. Additionally, NLPEXPLORER provides year-wise usage distribution, total usage, and the list of top-most subdomains and associated papers.

## 5 Future Extensions and Conclusion

The current system provides a basic functionality for knowledge exploration in NLP domain. In future, we plan to incorporate advanced set of functionalities leveraging natural language understanding of research papers. Some of the main proposals include natural language query retrieval, intelligent ranking by leveraging citation sentiments and discourse-level citation information, visualization of topical flow from cited papers to the main text of citing papers, automatic generation of leaderboard for NLP tasks, visualization of author collaboration networks, paper citation networks, venue interaction networks, etc.

In this paper, we present an end-to-end automated system that periodically mines the ACL Anthology and serves as a tool to aid researchers in knowledge exploration and discovery. The goal of NLPexplorer is to serve as a retrieval engine for research papers, as well as a tool to assist researchers in knowledge discovery by helping them to better understand the problem domain, the top researchers in the field of study, and the latest research in the domain. Even though, the current system supports NLP research domain, we claim that similar systems can be built for any domain given the availability of full text articles and basic metadata.

## References

1. ACL Anthology Homepage. <https://www.aclweb.org/anthology/>. Accessed 13 Oct 2019
2. Bird, S., et al.: The ACL anthology reference corpus: a reference dataset for bibliographic research in computational linguistics. In: Language Resources and Evaluation Conference (2008)
3. Elasticsearch Homepage. <https://www.elastic.co>. Accessed 13 Oct 2019
4. Google Cloud Homepage. <https://cloud.google.com/>. Accessed 13 Oct 2019
5. MongoDB Homepage. <https://www.mongodb.com>. Accessed 13 Oct 2019
6. NLPexplorer Homepage. <http://nlpexplorer.org/>. Accessed 13 Oct 2019
7. Plotly JavaScript Homepage. <https://plot.ly/javascript>. Accessed 13 Oct 2019
8. Radev, D.R., Muthukrishnan, P., Qazvinian, V.: The ACL anthology network corpus. In: ACL-IJCNLP p. 54 (2009)
9. Radev, D.R., Muthukrishnan, P., Qazvinian, V., Abu-Jbara, A.: The ACL anthology network corpus. *Lang. Resour. Eval.* **47**(4), 919–944 (2013). <https://doi.org/10.1007/s10579-012-9211-2>
10. Schäfer, U., Kiefer, B., Spurk, C., Steffen, J., Wang, R.: The ACL anthology search-bench. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Systems Demonstrations, pp. 7–13. Association for Computational Linguistics (2011)
11. Singh, M., et al.: OCR++: a robust framework for information extraction from scholarly articles. In: Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers, pp. 3390–3400. The COLING 2016 Organizing Committee (2016)
12. Singh, M., et al.: CL scholar: the ACL anthology knowledge graph miner. In: Proceedings of the NAACL (2018)
13. TimelineJS Homepage. <https://cdnjs.com/libraries/timelinejs>. Accessed 13 Oct 2019