

TweetPap: A Dataset to Study the Social Media Discourse of Scientific Papers

Naman Jain and Mayank Singh
 Department of Computer Science and Engineering
 Indian Institute of Technology Gandhinagar
 Gandhinagar, India
 singh.mayank@iitgn.ac.in

Abstract—Nowadays, researchers have moved to platforms like Twitter to spread information about their ideas and empirical evidence. Recent studies have shown that social media affects the scientific impact of a paper. However, these studies only utilize the tweet counts to represent Twitter activity. In this paper, we propose *TweetPap*, a large-scale dataset that introduces temporal information of citation/tweets and the metadata of the tweets to quantify and understand the discourse of scientific papers on social media. The dataset is publicly available at <https://github.com/lingo-iitgn/TweetPap>.

I. INTRODUCTION

As digitization led to the availability of a vast volume of publications, measuring research impact became important in order to analyze a publication’s performance and optimizing search engines to retrieve relevant articles [1]. Research institutes also use research impact to assess the performance of faculty members, make grant allocations, set policies, and recruit new members [3]. The eagerness to improve research impact led researchers to “push” their work on social media platforms like Twitter for greater reach [4], [5]. Existing studies have attempted to capture the discourse of scientific knowledge through social media platforms by analyzing the number of mentions of a publication on Twitter in small-scale datasets created manually or using surveys and analytics platforms [2]–[4], [6] whereas large-scale datasets lack analysis with citation data [5].

We propose a novel data collection scheme to broaden the study of quantifying the effect of social media activity on the research impact of scholarly documents and present *TweetPap* a large-scale dataset consisting of Twitter activity of arXiv papers created using this scheme. Other than the tweet count and citation count of papers, as introduced in previous studies, *TweetPap* also provides access to yearly citations of a paper, yearly retweets of a particular tweet mentioning a paper, likes on each tweet, and a list of embedded URLs for understanding the discourse of scientific knowledge.

TABLE I
STATISTICS OF THE *TweetPap* DATASET

Year range	2010–2019
Tweets	367,124
Papers	125,521
Links	418,494
Users	16,098

II. DATA COLLECTION SCHEME

To assess the diffusion of scientific papers on social media platforms, we need a comprehensive dataset with bibliographic information and the corresponding social media activity. Due to open-access permissions and public availability of citation data, *TweetPap* is compiled from the papers in arXiv. We collected the relevant data by searching the keyword “*arXiv*” in Twitter’s historical database. The tool¹ retrieves all the tweets mentioning the keyword “*arXiv*” in its metadata (text, username and comments). We collected tweets from a time range of 10 years (2010–2019) and mapped the papers to the tweets using arXiv identifiers (ID). We chose arXiv ID because it can be easily retrieved using regular expressions, whereas paper title matching may result in high computational cost and false matches due to variation in punctuation, grammar, and presence of incomplete titles². Apart from the tweet’s text, the arXiv IDs can also be retrieved from the links mentioned in the tweet in many cases. A major challenge in extracting IDs from the links in the tweets is that most of the links are shortened (discussed more in III-C). To extract arXiv IDs from such links, we un-shortened³ all the links from the tweets which do not have an arXiv ID in the non-hyperlink text of the tweet. Finally, using the regular expression “[0-9]{4}.[0-9]{4}[0-9]?” (based on the arXiv guidelines⁴), we extracted the arXiv ID from unshortened links. We do not consider the tweets without a retrievable arXiv ID from links or text. Further, we leverage the Semantic Scholar Corpus [7] to extract the yearly citation information of arXiv papers. From the arXiv IDs in the tweets, we extracted yearly citation information of 125,521 papers occurring in 367,124 tweets. Since the same paper can be mentioned in multiple tweets, we present a cumulative count of retweet/like from all the tweets corresponding to that paper. In the end, we obtain a dataset with statistics and attributes defined in Table I and Figure 1 respectively, along with yearly citations and tweet information such as text and yearly retweets data (pulled from Twitter API⁵).

¹github.com/Jefferson-Henrique/GetOldTweets-python

²Only 16,435 tweets were retrieved using title matching

³github.com/SMAPPNYU/urlExpander

⁴https://arxiv.org/help/arxiv_identifier

⁵<https://python-twitter.readthedocs.io/en/latest/>

“arXiv ID: <id>” “Citation: <map[year→int]>”
 “Publication Year: <int>” “Retweet: <map[year→int]>”
 “Tweet IDs: <list[int]>” “Likes: <int>”
 “Links: <list[string]>” “Users: <list[string]>”

Fig. 1. Attributes of *TweetPap* alongwith data types

III. DATASET ANALYSIS

This section describes the analysis of *TweetPap*, focusing on features showcasing the significance of this dataset.

A. Distribution of attributes

Figure 2 (a) shows the frequency distribution of tweets and papers based on the year of release. Contrary to our expectation, we found a significantly low number of tweets between 2015–2018. *TweetPap* confirms positive correlation between Twitter activity and citation count. For example, the average citation count increases from 7.88 to 21.2 as the number of tweets per paper increases from one to five.

B. Temporal peaks of citations and retweets

We extract the peak years for both of these attributes by utilizing the yearly citations and retweet data. A peak year refers to the year in which the corresponding attribute was the highest for a paper. In the case of multiple peak years (years with similar maximum frequency), we consider the first peak. Figure 2 (b) shows the frequency distribution of the difference of the years between citation peak and retweet peak of papers with citations count ≥ 20 and retweet count ≥ 10 (1881 papers). Higher thresholds are taken to showcase results for dense distributions. Interestingly, the positive difference suggests that the citation peak occurs after the retweet peak for the majority of the papers. Table II shows a high positive correlation between the attributes at the peak difference of two years. This experiment opens avenues for analysis of temporal peaks in causal relations of social media discourse.

C. Links and Users of tweets

The shortened links present in the tweets posed a challenge while mapping papers with tweets. Except for “fb”, “lnkd” and “github”, all other domains in the Top-10 most occurring domains are link shortening services supported by Twitter or a third party. These services change the text corresponding to the actual web address and provide a short version of the link which forwards the user to the actual address. For e.g. a paper with “2008.01342” as arXiv ID is embedded as “https://ift.tt/2Dkr4rw” in one of the tweets.

All users except “@animesh1977” amongst the top-10 users posting unique arXiv papers on Twitter are *automated bots* which regularly crawl arXiv and update new papers on Twitter. This shows that automated bots are very prominent in spreading scholarly documents.




Fig. 2. (a) Frequency distribution of tweets and papers over the years (b) Peak difference distribution

TABLE II
PEARSON’S CORRELATION VALUES OF PAPERS WITH HIGH CITATION (≥ 20) AND HIGH RETWEET (≥ 10). R=RETWEET, C=CITATION, L=LIKES

Peak Difference	Paper Count	RC Corr	LC Corr	RL Corr
2	902	0.236	0.421	0.301
1	431	0.079	0.072	0.056
0	678	-0.058	-0.037	0.042

IV. CONCLUSION AND FUTURE WORK

In this work, we introduce a novel method to collect better data for boosting the research in analyzing the relations between social media and scientific literature by exploiting temporal features, metadata of tweets, and citation information. This work provides directions that can act as an entry point to accurately analyze the chaotic behavior of social media in the scholarly ecosystem. In the future, *TweetPap* can be made more robust by incorporating semantics of tweets, keywords from popular scientific publishers with dense citation data, and tweets from the network of top researchers. Also, works in citation prediction have utilized features related to the metadata of a paper [1]. However, these prediction models still lack features concentrating on the medium of the spread of information. One can explore the possibility of utilizing *TweetPap* in predicting the impact.

REFERENCES

- [1] M. Singh, A. Jaiswal, P. Shree, A. Pal, A. Mukherjee, and P. Goyal, “Understanding the impact of early citers on long-term scientific impact,” in 2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL), 2017.
- [2] S. Klar, Y. Krupnikov, J. B. Ryan, K. Searles, and Y. Shmargad, “Using social media to promote academic research: Identifying the benefits of twitter for sharing academic work,” PLoS One, vol. 15, no. 4, p. e0229446, 2020.
- [3] J.-C. Na, “User motivations for tweeting research articles: A content analysis approach,” in Digital Libraries: Providing Quality Information, Cham: Springer International Publishing, 2015, pp. 197–208.
- [4] E. Mohammadi, M. Thelwall, M. Kwasny, and K. L. Holmes, “Academic information on Twitter: A user survey,” PLoS One, vol. 13, no. 5, p. e0197265, 2018.
- [5] A. Sesagiri Raamkumar, M. Erdt, H. Vijayakumar, A. Nagarajan, and Y.-L. Theng, “Understanding the twitter usage of science citation index (SCI) journals,” in Digital Libraries at the Crossroads of Digital Information for the Future, Cham: Springer International Publishing, 2019, pp. 54–60.
- [6] J. L. Ortega, “The presence of academic journals on Twitter and its relationship with dissemination (tweets) and research impact (citations),” Aslib j. inf. manag., vol. 69, no. 6, pp. 674–687, 2017.
- [7] W. Ammar, et al. “Construction of the literature graph in semantic scholar.” arXiv preprint arXiv:1805.02262, 2018.