

Predicting Research Trends From Arxiv

Steffen Eger¹, Chao Li², Florian Netzer¹, and Iryna Gurevych¹

¹ Technical University Darmstadt, Darmstadt, Germany

² Frankfurt School of Finance and Management, Frankfurt am Main, Germany.

Abstract. We perform trend detection on two datasets of Arxiv papers, derived from its machine learning (cs.LG) and natural language processing (cs.CL) categories. Our approach is bottom-up: we first rank papers by their normalized citation counts, then group top-ranked papers into different categories based on the tasks that they pursue and the methods they use. We then analyze these resulting topics. We find that the dominating paradigm in cs.CL revolves around *natural language generation* problems and those in cs.LG revolve around *reinforcement learning* and *adversarial principles*. By extrapolation, we predict that these topics will remain lead problems/approaches in their fields in the short- and mid-term.

Keywords: Trend prediction · Citation Counts · Deep Learning · Arxiv.

1 Introduction

Predicting trends in research has been a long-standing dream of scientists. Projects on popular research topics often lead to higher acceptance rates at conferences and journals, as well as funding application approvals. Further, knowing future research trends also has implications for society as a whole, because these trends will most likely directly affect the labor market, technological orientation and biases, consumer end products, as well as cultural metaphors and definitions of the human identity—this is even more true for fields such as artificial intelligence, as we focus on here. However, with the accelerating number of papers made available each year, it becomes ever more difficult to digest the incoming information and thereby manually identify topics that will have long-term scientific impact. We have developed an automatic system whose goal is to uncover important research trends, and, therefore, aims at helping researchers better plan their academic endeavors.

Our system crawls papers published in the Machine Learning (cs.LG) and natural language processing (cs.CL) categories of Arxiv,³ with information about how often they were cited. In this dataset, we identify promising papers using normalized citation counts and then categorize them by hand and automatically. Using Arxiv papers for our exploration appears promising, because Arxiv is a

³ <https://arxiv.org/>

very popular pre-print (and post-print) server for scientific publications, whose impact has, moreover, considerably increased over the last few years.⁴

2 Data & Annotation

2.1 Data

We created two datasets: One with papers from the machine learning (cs.LG) category of Arxiv and one with papers from computation and language (cs.CL). We focus on these two prominent subfields of artificial intelligence because they appear concurrently particularly dynamic, with drastic changes and performance improvements witnessed each year, mainly due to the impact of artificial neural network (aka deep learning) models. The data includes papers with their title, abstract and authors. We also harvested citation information for the papers from semanticscholar.⁵ Our crawled papers date between June, 2017 and December, 2018. We crawled roughly 4.8k papers from cs.CL and 12.4k papers from cs.LG.

2.2 Annotation

One co-author of this study manually annotated the abstracts of the top-100 (according to their normalized citation count; see Section 3) papers in cs.LG and cs.CL, respectively, for three aspects: **task**, **method**, and **goal/findings**. These aspects answer the questions *what* a paper researches, and *how* and *why* it does so.

Title	task	method	goal
IEST: WASSA-2018 (18, 13.0) Implicit Emotions Shared Task	Emotion Det.	Data	Difficult task
BERT: Pre-training of Deep (20, 12.0) Bidirectional Transformers for Language Understanding	Text repr.	Transformer	Better accuracy
Deep contextualized (261, 11.2) word representations	Text repr.	Language Model	Better accuracy

Table 1. Top-3 papers from cs.CL according to their normalized citation counts, as well as **task**, **method** and **goal** classification. In brackets: absolute and normalized citations by end of December 2018.

Table 1 lists the top-3 papers from cs.CL, together with their **task**, **method**, and **goal** annotations. For cs.CL, we decided upon 15 fine-grained **task** label

⁴ See submission statistics to Arxiv at https://arxiv.org/help/stats/2018_by_area/index.

⁵ <https://www.semanticscholar.org/>

categories, given in Table 2. Similarly, we decided upon 28 **method** categories, and 7 **goal** categories. For cs.LG, we used 13,15, and 13 **task**, **method**, and **goal** labels, respectively. We note that it was sometimes difficult to annotate the abstracts for any of the three categories because the information in question may not have been available from the abstract or an abstract could not be clearly assigned to one of the agreed upon categories. We also note that we used different labeling schemes for cs.LG and cs.CL, respectively. This is because papers in cs.CL often deal with a fine-grained natural language processing (NLP) problem whereas wording in cs.LG abstracts ignored such fine-grained differences or claimed to address *various* problems from, e.g., image or text fields. We also mention the general vagueness and arbitrariness (to a certain degree) of such classifications. Overall, it was easiest and least ambiguous to classify cs.CL papers to **tasks** and cs.LG papers to **methods**. Hence, we mostly focus on these two combinations.

task -cs.CL	Generation, Machine Translation (MT), Text representations, Speech, Language Modeling, Sentence Classification, Style Transfer, Reasoning, Relation extraction, Sequence Tagging, Emotion Detection, Argument Mining, Human-Computer Interaction, Parsing, Rest
method -cs.LG	Reinforcement Learning (RL), Other Deep Learning architect., Representation Learning, GAN, Generation, Architecture Search, Distillation, Analysis, Interpretability, Learning Aspects, Various, Data, Rest, Adversarial

Table 2. **task** and **method** labels for cs.CL and cs.LG, respectively.

3 Normalized citation counts

The simplest measure for a paper’s influence is the number of citations it has received. But plain citation counts may be misleading. They may vary depending on the research field and the date of the publication. Instead, citation counts can be normalized by comparing only papers in the same research fields and adjusting citation count scores by the paper’s age. This is the idea of the *z-score* approach suggested by Newman [8, 9]. This is calculated by subtracting the mean citation count of papers within a time window from the citation count of a paper and dividing by the standard deviation. With this method, Newman [9] does not merely identify promising papers as those that have a high number of citations because they were published earlier than other papers from the same field, but also papers with only a few citations. To illustrate the effectiveness of the method, after 5 years, the 50 papers with the initially highest z-scores in Newman’s sample received 15 times as many citations as the average paper in a randomly drawn control group *that started with the same number of citations*.

Thus, Newman argues that the z-score can indeed identify short- and mid-term research trends.⁶

To find promising papers, we calculated the z-score for the papers in our datasets, using a time-window of ± 10 days. We ignored papers with less than 4 citations, because we deemed such small numbers unreliable.

4 Results

Figure 1 shows the distribution of **tasks** in the top-100 list of cs.CL, with topic labels given in Table 2. There are few very interesting patterns: for example, traditional NLP tasks like sequence tagging and parsing are marginalized and make up only 5% of the top papers, while the most prominent tasks dealt with in the top-100 papers are almost all about natural language generation: machine translation, language modeling (where the goal is to predict the next word given the previous text sequence) and generation proper (consisting of subcategories dialogue, question answering, text2SQL generation and summarization). The **task** speech mostly also deals with generation, namely, generation of a speech signal from written text (text2speech). The only major category not dealing with text generation is ‘text representations’ which deals with either word embeddings [6] or sentence embeddings [1], that is, vector representations of either words or sentences that summarize their semantic and syntactic properties. One speculation is that this latter field, text representations, will now be gradually declining in importance, after having dominated most of the 2010s. The more high-level topic text generation appears to replace it in its leadership role.

Figure 2 shows the **method** distribution in the top-100 list of cs.LG. The most important methods here are Reinforcement Learning (RL) and Adversarial techniques. The latter either deal with adversarial attacks on neural networks—i.e., fooling these networks into misclassification or shielding from such fooling—or with adversarial training, which aims to make machine learning systems more robust. Generation is likewise prominent, through Generative Adversarial Networks (GANs), which have been used to generate, e.g., artificial images with impressive quality. Indeed, the leading paper in the top-100 list uses GANs to generate ‘fake celebrities’ [5]. Another interesting observation is that Reinforcement Learning is much less prominent in cs.CL (only 5% of all papers indicate to use this method in their abstracts). It remains open whether NLP papers will catch up with the trends in general machine learning or whether reinforcement learning is generally less suited for natural language processing problems.

Finally, Figure 3 shows the distribution of **goals** in the top-100 list in cs.CL. About half of the papers present a better system that outperforms the previous state-of-the-art. Another major motivation in our cs.CL list is to create new and better datasets, which cover more realistic conditions, for example. Finally, exposing weaknesses either of NLP systems—such as their failure to adversarial attacks—or of researchers—such as drawing insufficiently supported conclusions

⁶ We assume here that high future citation counts are a proxy for the research trend prediction task, and ignore cases of, e.g., negative citations.




Fig. 1. task distribution in top-100 list in cs.CL.




Fig. 2. method distribution in top-100 list in cs.LG.

regarding superiority of systems—are further main goals stated in the paper abstracts.




Fig. 3. goal distribution in top-100 list in cs.CL.

5 Related work

Analyzing and predicting research trends is of interest to a variety of stakeholders. For example, Clarivate Analytics publishes a yearly “Research fronts” report discussing leading scientific topics.⁷ Their analysis also uses citation count statistics, even though these are not normalized as in the z-scores that we use. A key difference to such studies is the database we use, namely, Arxiv. Arxiv is arguably a lead market place for the dissemination of machine learning research connected to deep learning (as a main driver of artificial intelligence research).

There are other methods discussed in the literature for trend detection besides those based on citations. The authors of [7] use the age distribution of the references of a paper to identify prospectively highly cited papers. The authors of [10] use the rhetorical status of topics in paper abstracts to identify topics that will increase or decline over time. Further studies combine citation network analysis with temporal topic distributions obtained from LDA [4], model the success of a paper by the centrality of its authors in co-authorship networks [12], or track trends by analyzing keyword networks [2].

⁷ <https://clarivate.com/essays/research-fronts/>

6 Concluding remarks

We have developed a system to rank Arxiv papers according to their z-score, in order to detect short- and mid-term research trends. Our manual evaluation and clustering of ranked papers indicated some potentially interesting paradigm shifts in machine learning and NLP in the upcoming years. It appears that the fields will be more heavily concerned with fooling and attacking deep learning systems, as well as defending against such attacks, with data generation (instead of classification, as in traditional machine learning and NLP research), and with (deep) reinforcement learning, which can be seen as another endeavor towards proper artificial intelligence where artificial agents learn in a way similar to humans, performing actions and receiving rewards for them.⁸

Further automatization of our methodology is necessary: we want to assign topics to ranked Arxiv papers automatically rather than manually. We will use neural text embeddings [1], keyword analysis [3] and ontologies [11] for this. In the same vein or alternatively, clustering of Arxiv papers must be done automatically. As a secondary signal next to the z-score ranks, we want to predict citation counts from a model that incorporates the text and the meta-data (such as authors and publication venue) of a paper. Finally, we want to conduct a diachronic study: how does the ranking of high z-score papers/topics change over time? This may be considered as an additional layer of estimating and quantifying trends and their shifts.

A final note is on our data. As with each dataset, Arxiv has its particular biases. One such bias is its connection to deep learning (due to deep learning’s fast pace, which makes it an ideal candidate for a pre-print server). Indeed, almost all papers in our top-100 lists are about deep learning. As such, the topic distributions obtained from Arxiv may look considerably different from those derived from standard machine learning or NLP conferences—such conferences have their own biases, namely, to achieve a certain balancedness and democracy in choosing topics, which may in turn not truly reflect where the action is.

Our tool plus the annotated data and the ranked Arxiv lists are available at <https://github.com/UKPLab/refresh2018-predicting-trends-from-arxiv>

References

1. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. CoRR **abs/1810.04805** (2018), <http://arxiv.org/abs/1810.04805>
2. Duvvuru, A., Kamarthi, S., Sultornsane, S.: Undercovering research trends: Network analysis of keywords in scholarly articles. In: 2012 Ninth International Conference on Computer Science and Software Engineering (JCSSE)
3. Florescu, C., Caragea, C.: Positionrank: An unsupervised approach to keyphrase extraction from scholarly documents. In: ACL (1). pp. 1105–1115. Association for Computational Linguistics (2017)

⁸ This analysis assumes that the normalized citation counts we use indeed help us to identify future citation behavior.

4. He, Q., Chen, B., Pei, J., Qiu, B., Mitra, P., Giles, L.: Detecting topic evolution in scientific literature: How can citations help? In: Proceedings of the 18th ACM Conference on Information and Knowledge Management. pp. 957–966. CIKM '09, ACM, New York, NY, USA (2009). <https://doi.org/10.1145/1645953.1646076>, <http://doi.acm.org/10.1145/1645953.1646076>
5. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=Hk99zCeAb>
6. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. CoRR **abs/1301.3781** (2013), <http://arxiv.org/abs/1301.3781>
7. Mukherjee, S., Romero, D.M., Jones, B., Uzzi, B.: The nearly universal link between the age of past knowledge and tomorrows breakthroughs in science and technology: The hotspot. *Science advances* **3**(4), e1601315 (2017)
8. Newman, M.E.: The first-mover advantage in scientific publication. *EPL (Europhysics Letters)* **86**(6), 68001 (2009)
9. Newman, M.: Prediction of highly cited papers. *EPL (Europhysics Letters)* **105**(2), 28002 (2014)
10. Prabhakaran, V., Hamilton, W.L., McFarland, D., Jurafsky, D.: Predicting the rise and fall of scientific topics from trends in their rhetorical framing. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). vol. 1, pp. 1170–1180 (2016)
11. Salatino, A.A., Thanapalasingam, T., Mannocci, A., Osborne, F., Motta, E.: Classifying research papers with the computer science ontology. In: International Semantic Web Conference (P&D/Industry/BlueSky). CEUR Workshop Proceedings, vol. 2180. CEUR-WS.org (2018)
12. Sarigöl, E., Pfitzner, R., Scholtes, I., Garas, A., Schweitzer, F.: Predicting scientific success based on coauthorship networks. *EPJ Data Science* **3**(1), 9 (Sep 2014). <https://doi.org/10.1140/epjds/s13688-014-0009-x>, <https://doi.org/10.1140/epjds/s13688-014-0009-x>