

A Hybrid Approach for Dynamic Topic Models with Fluctuating Number of Topics

Christin Katharina Kreutz
Trier University
54286 Trier, DE
kreutzch@uni-trier.de

ABSTRACT

Scientific communities are always changing and evolving. Topics of today might split or even disappear in the future, other topics might merge or appear at some time. Nowadays, the closest we come to picture these developments are dynamic topic models which come with a fixed number of topics k . It would be desirable to omit k . This work outlines a research agenda for approaching that task by using LDA as a base in combination with the observation of state transitions in topics at consecutive times.

Categories and Subject Descriptors

H.1 [Models and Principles]: Document Topic Models;
I.5 [Pattern Recognition]: Trend Mining

General Terms

Algorithms

Keywords

Trend Mining, Dynamic Topic Models, LDA

1. INTRODUCTION

With today's publication methods, the number of papers increases rapidly. Losing track of the evolution of the majority of themes is common. Simultaneously, identifying important publications is difficult but cardinal for scientists.

Automatic detection of trends and their indicators in a scientific community (*trend mining*) could benefit researchers, politicians or entrepreneurs who are not ahead of current developments but want to get quick insights into promising areas.

Our goal is to construct a system, which autonomously identifies trends and accompanying influential persons and papers from a variety of bibliographic data. The appurtenant research plan is partitioned into three succeeding sections: First, the transformation of topics generated from a bibliographic data set over time, their assigned papers, authors




Figure 1: Simplified visualisation of our research plan.

and keywords should be mapped in a dynamic topic model with variable number of topics. Second, potential upcoming trends in the topics across the years should automatically be detected, predicted and extracted from this model, so they can be evaluated. And third, influential authors, papers and venues should be determined in these found trends. The resulting new insights about what supports the development of a topic can be used to enhance the identification of trends.

The steps are relatively independent of another, step two would be applicable on another suitable topic model without requiring a solution of step one. Figure 1 gives a schematic overview of our projected line of action.

In this work, we focus on outlining a research direction for the first step, present current state of research on related models and mark the problems at hand. We touch on trend mining, before we close with an evaluation plan and an outlook on possible application for our future model.

2. DEVELOPMENT OF TOPICS

We assume the importance and set of topics is not static over time. Topics might sprout, expand, diminish, split, merge or vanish. Terms that represent the topics change as new words appear [5]. To better understand the dynamics

of topics, we wanted to observe real bibliographical data.

2.1 Notation

Before diving into details of our experiments or the proposed model, some basic terms need to be set in order to formally discuss our concepts.

A paper has a number of fundamental, possibly latent, ideas. They can be grouped by motive to more general *topics* denoted by s_i . By observing co-occurring topics and terms in papers, conclusions about the assignment of terms to topics can be drawn. Topics can be term-wise alike or (partially) overlap with other topics. Assertions on this can be derived from the term distributions for topics.

The total time observed t can be sliced in disjunct consecutive intervals which are called times t_0, \dots, t_n . Given two times t_x and t_y , if $x < y$, t_x indicates an interval (and real period) before t_y . Given two times t_x and t_{x+1} , t_x describes the interval immediately before t_{x+1} .

Publications can be uniquely attached to intervals if the time is sliced by year and their year of publication determines the assignment. Exact publication dates are mostly not available. This classification is an approximate observation raster as in theory there is a time continuum and in reality we only have rough year specifications. States of topics are regarded at times.

A topic s_i is said to be *trending* at time t_{x+y} , $y \geq 1$, if it is unpopular or not even existing at time t_x , but its significance soars. This could be indicated by an increasing number of publications targeting this subject or its appearance in important journals or conferences. Essential members of the scientific community might start to work in this direction or the subjects builds its own experts which become widely known.

A topic that has not (yet) assigned any publications is described by s_0 . This case occurs before a topic is born or if it is inactive. A topic is *inactive*, if the number of publications assigned to the topic does not surpass a threshold or papers assigned with this topic do only cite papers from the same topic and are only cited by papers from this area. The topic has hardly any influence on the rest of the corpus. The community which works on this is very tightly connected but relatively isolated from the rest of the scientific world. These enclaves can be described as *sects*.

Opposing inactive topics are *active* topics. The set of active topics at a time t_x can be identified by k_x . The set of inactive topics at a time t_x can be described by $k_{\bar{x}}$.

2.2 Data Set

The data set used in this research is an incompletely enriched form of the dblp computer science bibliography data with part of the data from open academic graph. The dblp data contains bibliographic information related to publications, authors, conferences and journals from the field of computer science and adjacent areas [15]. As of February 2018, it holds metadata of over 4 million publications and more than 2 million authors. The Microsoft Academic Graph within open academic graph is used. It contains over 166 million publications and amongst others citation information, abstracts and details on authors [22, 21].

In our set, data from dblp was used completely. In addition, where publications could be matched based on DOI or title and author matches where DOI information was not available, information from open academic graph was in-

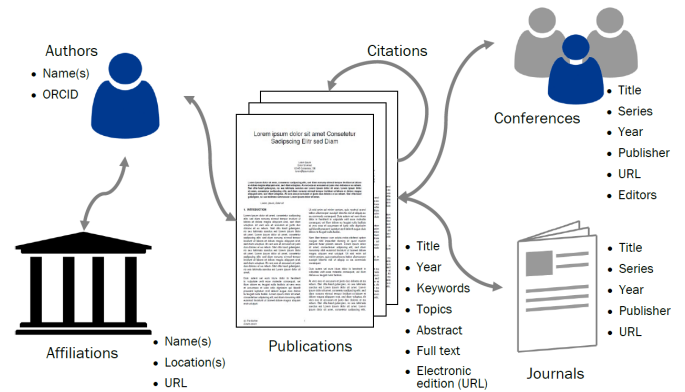


Figure 2: Simplified depiction of the composition of the extended dblp data set. Data is partial.

cluded. The extension contains author affiliations, citation data, abstracts, full texts, keywords and topics. The structure of the data set is depicted in Figure 2. Because we only focus on bibliographic information, further data sources like Twitter are not incorporated in our set.

For the experiments in this paper, only the data contained in dblp as well as abstracts were taken into consideration. At the moment, full texts are only available for a certain small area in computer science so the usage of them could have distorted the outcome of our initial trials drastically.

2.3 Methodology

Of the enriched dblp data, only English publications whose abstract was of considerable length (≥ 10 words, fewer words indicate flawed data) were taken into account. The titles and abstracts were purged and stemmed with a Porter stemmer. Afterwards, LDA [4] with $k = 100$ was run on all 2.5 million of them. We ignore terms occurring in over 50 percent of publications (collection dependent stop words) or in under 100 papers as they are often system names.


A visualisation of the data enabled us to draw conclusions about the characteristics of topics.

2.4 Initial Observations

In Figure 3, the popularity of a topic in relation to all topics in the corpus per year is visualised for the years 1990 to 2015 for four selected topics. We assume the number of topics is appropriate. Different settings can be observed:

- There are subjects, which are inactive and whose popularity rises, so they become active like topic 12, which is about mobile devices.
- There are subjects, which were always active and whose popularity increases as seen in topic 13, which covers terms like management, knowledge and business.
- There are subjects, whose popularity declines such as seen with topic 27, which includes papers concerning logic programming and reasoning.
- There are subjects, whose popularity does not really seem to change over the course of years such as topic 76, which deals with image processing.

In our data set, we found the case of a topic being active at a point in time but unrepresented by publications



(a) Overview of popularity of selected topics, topic distributions of papers are sliced by year. Size of bubble indicates relative importance of topic in all papers from this year.

Topic	10 most important stems
12	mobil, devic, network, commun, peer, music, ad, hoc, messag, wireless
13	manag, knowledg, studi, inform, research, technolog, organ, busi, factor, effect
27	program, logic, fuzzzi, oper, reason, gener, comput, base, languag, execut
76	imag, color, reconstruct, map, method, algorithm, base, render, resolut, pixel

(b) Topic number with corresponding assigned most important stems.

Figure 3: Exemplary illustration of the development of selected topics over time and their associated stems by running LDA with $k = 100$ on the whole extended dblp data set.

for a few following years. Later, it re-emerged. The topic's top keywords contained *cloud*, so early publications with a portion of this topic might have a background in weather, whereas the late publications which were (partly) assigned to the topic probably pick up on cloud computing.

The importance and number of active topics is highly varying throughout the years.

3. PROBLEM

Topics can be generated from a corpus by several probabilistic topic models. The most popular ones all have the significant weakness of an unchangeable number of topics. Before we dive into the problem, we present some existing methods.

3.1 Topic Models

The assignment of topics to papers can be performed by a number of approaches. The simplest one would be Latent Dirichlet Allocation *LDA*. Here, it is assumed that every document is a mixture of topics and every word in the documents comes from a specific drawn topic. There are no words that are partially assigned to no or even a residue topic. Hidden random variables contain information on the structure of topics in the documents. First, topic proportions for a document are drawn. After this step, for every position of a word in the document, a topic is drawn from this distribution. In the last part, actual words are drawn from the topic word distribution. LDA and constitutive models assume that documents are interchangeable in time. The number of topics k is fixed for a corpus and has to be chosen beforehand. The vocabulary of the corpus is also fixed. [4]

A lot of approaches build upon LDA, such as the Author-Topic Model *ATM*. Here, an additional dimension, the authors, is taken into account. The individual author codetermines the topic from which a word is drawn. [18]

The correlation of topics was presented with Correlating Topic Models *CTM*. Here, LDA was modified so instead of drawing topic distributions for documents from a dirichlet distribution, they were now taken from a logistic normal distribution. [2]

The temporal aspect of a collection and the development of topics has been widely disregarded until the introduction of Dynamic Topic Models *DTM*. This method extends CTM

by dividing a corpus by year so the topic distribution can change over time. Topics in slice t_{x+1} are derived from the topics in slice t_x . Words assigned to a subject are variable but k is still fixed. Information relating to authors is not used but papers are no longer interchangeable. [3]

3.2 Problem Description

The described methods cannot fully map the dynamics in a corpus, as the number of topics k is unchangeable. If data up until a point in time t_x is used to generate a DTM, at time t_{x+1} new publications can only be assigned to these already existing k topics. If DTM would be run with new publications and $k + n$ topics, the resulting topics would not necessarily represent the former k and additional n new ones even closely. Changing k slightly results in a different document topic distribution.

An easy way to capture the dynamics of topics would be to find a suitable k , perform LDA on the whole corpus, slice the corpus by year and look at topics changing over time like we did in our experiment. Trends could be found retrospectively. If new data is integrated, LDA could be used another time on all the publications. Again, trends could be located in retrospect. Big disadvantages are the determination of k and the inability to map the topics of the first run to the topics of the subsequent runs, especially if k is incremented. Terms which get mapped to subjects shift and it is impossible to regain old patterns. It would be unfeasible to measure if the identification of future trends was successful.

Emergence, disappearance, splitting and merging of topics over the course of time cannot be modelled with existing probabilistic topic models. Changes in subjects are indicators for trends and should thereby be observed.

There are other approaches to find trends which make use of a number of other features: Asooja et al. utilise keyword distributions on textural information [1], Glänzel et al. work on citations and textual information [9], Salatino et al. observe a topic network deployed from connections between keywords, publications, authors, venues and organisations [19].

Current methods usually only use a small portion on the spectrum of available data. A model which incorporates authors, affiliations as well as scientometric measures [20, 13, 10], publication information such as citations [17] and ve-

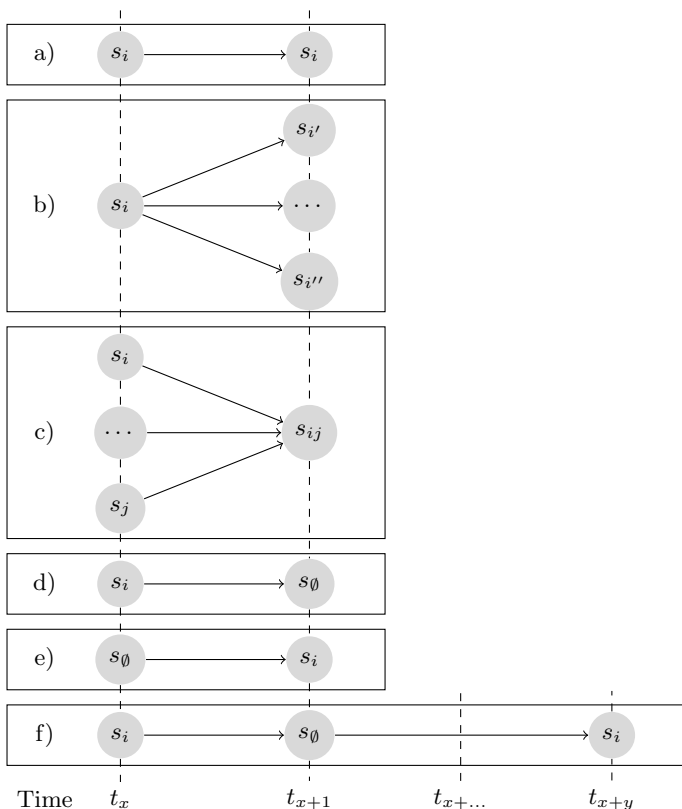


Figure 4: Possible state transitions of topics s_i over time t .

nues in addition to titles, abstracts, full texts, keywords and topics has the potential to detect trends reliably.

4. HYBRID APPROACH

Our theoretic approach is based on the assumption that there are different topic state transitions. They need to be represented by our model.

4.1 Evolution of Topics over Time

We identified possible state transitions with which the evolution of topics can be described, they are shown in Figure 4. There are six distinguishable forms: Case *a*) shows a topic which does not significantly change, *b*) shows the split of a topic s_i into possibly numerous topics $s_{i'}$, \dots , $s_{i''}$ that are somewhat coherent or the emergence of a topic $s_{i''}$ from an already existing (and persisting) topic s_i , *c*) shows the merging of possibly numerous disconnected topics s_i , \dots , s_j into one, *d*) shows a vanishing topic, *e*) shows the birth of a new topic and *f*) shows a combination of cases *d*) and *e*) with the anomaly of the topic s_i being inactive and re-emerging over a span of time being the same. The different transitions can be joined ad libitum.

An example for *a*) could be the image topic we already encountered in Figure 3. The distribution of words in the topic surely changes over time, because the fundamental terms vary, though the overall motive in them stays the same. As instance of case *b*), algorithms concerning depth first search could be the base, from which other algorithms, such as ones for the computation of strongly connected components, derived. The original topic persisted while new ones

were emerging from it. A topic describing machine learning might be a good example of case *c*). Many areas treating algorithms are collapsing into this big one, as machine learning has the potential to outperform even the most refined hand-knitted approaches. If a topic describes RSA, it could fall into category *d*), as it is no longer considered safe, therefore publications concerning this subject are most likely going to decrease over the next years until the topic is inactive. This is a good candidate for the forming of a sect. The development of a topic for quantum computers could be mapped to case *e*). It somewhat was the birth of this topic in computer science. There certainly were influences from different communities on the subject but in a corpus restricted to information technology, the representation might be fitting. As neural networks are currently experiencing a renaissance, they are an example of *f*).

4.2 Hybrid Topic Model

Our future model needs to be able to find and represent all described transitions of topics. In the following, we explain the core components of a hybrid model.

The rough plan would be to split t in years and use LDA to generate a baseline of topics for t_0 . For every new year, the topics of the prior year need to be considered when calculating the current developments. Citations are a key part in this as they indicate how information is being spread. At time t_{x+1} , we examine k_x as well as $k_{\bar{x}}$ and observe co-authorships, used words and how new publications cite already classified papers. By looking at the topic distributions and summing the percentages for each topic, it can be calculated, which topics are cited with corresponding weights by a new paper. With for example the Wasserstein metric [8], the distance between term distributions of topics $dist_{td}$ is calculated as their difference. A threshold th_{td} describes the distance value over which topic term distributions are considered dissimilar.

For every topic, the following strategies decide which state transition has occurred from t_x to t_{x+1} :

- With the first case, there is no major change in underlying motives from t_x to t_{x+1} . Publications in this topic reference about the same topics that were cited at t_x and $th_{td} > dist_{td}$. The content in cited publications is typically pretty similar to the content of the new ones.
- In this situation, we have the same phenomena as in case *a*) but a clustering on publications of this topic produces multiple distinguishable groups which are regarded as new topics split from the old one, $th_{td} < dist_{td}$ amongst the new topics. New words are likely to occur in the publications. If they solely appear in the papers from this area and not throughout the whole corpus, they strongly hint at a change or split in the topic.
- If a merging of topics occurs, the witnessed effects will resemble those of case *a*), although publications which would be ordered to prior topics harmonise their term distributions and citation behaviour. A clustering would group the topics together.
- A dying topic gets none or few new publications assigned to. The number of papers in this topic might already be declining for a few years. A topic getting inactive all of a sudden is highly unlikely.

- e) If a new topic emerges, publications do not really match term distributions of existing ones. They usually cite a lot of different topics as they have no clear predecessor. The overlap of content from cited papers (not topics) by a new publication and the citing paper should be calculated, as it is deemed to be rather small.
- f) With the sudden re-emergence of a topic, the term distribution of publications match a topic in $k_{\bar{x}}$.

After the topic distributions for the new publications are computed, the then active and inactive topics are assigned to k_{x+1} and $k_{\bar{x}+1}$ respectively. A run concludes with the processing of the next year of papers in the same manner.

4.3 Topic Development Prediction and Trend Mining

Predicting the development of a topic is directly linked to trend mining. Topics which are about to blow up are future trends. The upcoming number of publications in a field, the estimation of citations a new paper is going to gain [17] and possible collaborations between researchers can only be computed if the underlying author-publication-graph of the past is thoroughly analysed and influences on its evolution are discovered.

The computation of trends in currently active topics is a step which follows directly from the hybrid topic model. Topics which changed a lot from t_x to t_{x+1} are candidates for trends. Not only the development of topics from the last to the current time frame is going to be observed, the overall behaviour of the term distributions and cited topics are relevant. The appearance of new and popular words in the assigned terms of a topic could signal the beginning of a trend and is worth further investigation.

Often, popular papers are written by well-known and highly linked authors, they appear in journals with a lot of impact or are presented at seminal conferences. Here, the enriched data is going to be used. A co-author-graph with researchers' affiliations linked to a paper-citation-graph complete with venues and relationships between journals and conferences could help discover core persons [7], venues and publications in topics and trends. Sometimes, trends also develop from sects, so they have to be steadily looked at. Topics which were active in t_{x+1} are judged on whether they are likely going to be trending in the future. The evolution can be predicted based on the progress of the topic and the found influences.

5. FUTURE PROSPECTS

After completing the construction of our hybrid approach, an evaluation of the proposed system needs to prove and quantify its validity. Furthermore, several practical uses for the model are presented.

5.1 Evaluation Plan

The evaluation of our planned system, which includes the trend mining part, contains multiple steps. The results need to be cross-validated.

Our hybrid model is going to be run on a base of data up until 1995, then topic developments are computed by the iterative part with data for the next 10 years. For the following 5 years, trends are predicted. Afterwards, a manual evaluation of our model and the found trends involves expert

researchers from different domains within computer science. A list which contains our results is presented to them. They should rate it against the real trends with corresponding years.

Additionally, the trends, important researchers and venues identified by our system will be presented to those experts. They then should rank the correctness of the findings.

An automatic method to quantify the accuracy of the model would involve the observation of data up until a time t_x . Potential trends at this time will be detected, their evolution and future importance is going to be predicted for the succeeding five years and the predictions will be compared to the real development of significance of these topics. Numbers of papers from topics and citation behaviour could be prognosticated. If there are discrepancies in predicted and real data, a manual step could be put in, to question experts to explain the actual development.

The hybrid approach also needs to be tested against the purely incremental model which does not use LDA with a predetermined k as first step.

5.2 Applications

Possible applications of the dynamic topic model with varying number of topics complete with the identification of trends are manifold. A reviewer recommendation system for given publications, a citation recommendation system, a keynote speaker recommendation system or a visualisation tool for exploring bibliographic data with special focus on trends could be constructed.

Some reviewer recommendation systems work on word topic and topic citation distributions [11] or are only usable for already established conferences as they use former program committees [23]. Others are more refined and want to integrate the research interest and direction of scientists into the recommendations [16, 12]. Our model is independent of past conferences. It could make use of the enriched author-publication-graph to find scientists capable and willing to review new publications from the field of their current research interest. As the available data for this task is extensive, the results could be excellent.

Citation recommendation systems suggest fitting publications based on their content, but they do not focus on returning fundamental papers which lead the way of a topic or those written by influential authors for an area [11]. The relative importance of a paper for an area and its development is not considered. With our hybrid model, the identification of influential papers and persons is a by-product and could be easily incorporated in such a system.

Keynote speakers for a conference from topic s_i should be influential scientists from a different topic s_j , which is related to s_i . A linkage of the topics could be predicted, the term distributions of the topics harmonise or one topic adapts words from the other area. The findings in one topic could highly benefit the other. Our model contains this information so it could be used for this application.

A visualisation tool for the exploration of found topics, relationships and trends in the data would be beneficial for researchers, politicians and entrepreneurs [5]. Past work on the exploration of topics or trends in bibliographic data sometimes lacks the support for growing and big data sets [14] or base on a topic model with fixed number of topics [6]. A tool using our model and data would inherently dodge these weaknesses.

6. CONCLUSION

This work proposed a hybrid approach which aims at modelling the agile evolution of topics and trends in a growing corpus of bibliographic data without a fixed and predefined number of topics with help of an LDA base. Different state transitions were used to describe the development of topics over time in detail. A link to trend mining was drawn. The work concludes with the presentation of an evaluation concept to confirm the utility of the approach and numerous examples of use to underline the potential of our future model.

Acknowledgements

Special thanks goes to my supervisor Ralf Schenkel for his invaluable support.

7. REFERENCES

- [1] K. Asooja, G. Bordea, G. Vulcu, and P. Buitelaar. Forecasting emerging trends from scientific literature. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation LREC 2016, Portorož, Slovenia, May 23-28, 2016.*, 2016.
- [2] D. M. Blei and J. D. Lafferty. Correlated topic models. In *Advances in Neural Information Processing Systems 18 [Neural Information Processing Systems, NIPS 2005, December 5-8, 2005, Vancouver, British Columbia, Canada]*, pages 147–154, 2005.
- [3] D. M. Blei and J. D. Lafferty. Dynamic topic models. In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, pages 113–120, 2006.
- [4] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.
- [5] J. Boyd-Graber, Y. Hu, and D. Mimno. Applications of topic models. 11:143–296, 01 2017.
- [6] A. J. Chaney and D. M. Blei. Visualizing topic models. In *Proceedings of the Sixth International Conference on Weblogs and Social Media, Dublin, Ireland, June 4-7, 2012*, 2012.
- [7] A. Fiallos Ordoñez, K. Jimenes, C. Vaca, and X. Ochoa. Scientific communities detection and analysis in the bibliographic database: Scopus, 04 2017.
- [8] A. L. Gibbs and F. E. Su. On choosing and bounding probability metrics. *INTERNAT. STATIST. REV.*, pages 419–435, 2002.
- [9] W. Glänzel and B. Thijs. Using 'core documents' for detecting and labelling new emerging topics. *Scientometrics*, 91(2):399–416, 2012.
- [10] D. Herrmannova and P. Knoth. Semantometrics: Towards fulltext-based research evaluation. In *Proceedings of the 16th ACM/IEEE-CS on Joint Conference on Digital Libraries, JCDL 2016, Newark, NJ, USA, June 19 - 23, 2016*, pages 235–236, 2016.
- [11] W. Huang, Z. Wu, P. Mitra, and C. L. Giles. Refseer: A citation recommendation system. In *IEEE/ACM Joint Conference on Digital Libraries, JCDL 2014, London, United Kingdom, September 8-12, 2014*, pages 371–374, 2014.
- [12] J. Jin, Q. Geng, Q. Zhao, and L. Zhang. Integrating the trend of research interest for reviewer assignment. In *Proceedings of the 26th International Conference on World Wide Web Companion, Perth, Australia, April 3-7, 2017*, pages 1233–1241, 2017.
- [13] P. Knoth and D. Herrmannova. Towards semantometrics: A new semantic similarity based measure for assessing a research publication's contribution. *D-Lib Magazine*, 20(11/12), 2014.
- [14] B. Lee, G. Smith, G. G. Robertson, M. Czerwinski, and D. S. Tan. Facetlens: exposing trends and relationships to support sensemaking within faceted datasets. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI 2009, Boston, MA, USA, April 4-9, 2009*, pages 1293–1302, 2009.
- [15] M. Ley. DBLP - some lessons learned. *PVLDB*, 2(2):1493–1500, 2009.
- [16] X. Liu, T. Suel, and N. D. Memon. A robust model for paper reviewer assignment. In *Eighth ACM Conference on Recommender Systems, RecSys '14, Foster City, Silicon Valley, CA, USA - October 06 - 10, 2014*, pages 25–32, 2014.
- [17] A. Livne, E. Adar, J. Teevan, and S. Dumais. Predicting citation counts using text and graph mining. February 2013.
- [18] M. Rosen-Zvi, T. L. Griffiths, M. Steyvers, and P. Smyth. The author-topic model for authors and documents. In *UAI '04, Proceedings of the 20th Conference in Uncertainty in Artificial Intelligence, Banff, Canada, July 7-11, 2004*, pages 487–494, 2004.
- [19] A. A. Salatino and E. Motta. Detection of embryonic research topics by analysing semantic topic networks. In A. González-Beltrán, F. Osborne, and S. Peroni, editors, *Semantics, Analytics, Visualization. Enhancing Scholarly Data*, pages 131–146, Cham, 2016. Springer International Publishing.
- [20] S. Siebert, S. Dinesh, and S. Feyer. Extending a research-paper recommendation system with bibliometric measures. In *Proceedings of the Fifth Workshop on Bibliometric-enhanced Information Retrieval (BIR) co-located with the 39th European Conference on Information Retrieval (ECIR 2017), Aberdeen, UK, April 9th, 2017.*, pages 112–121, 2017.
- [21] A. Sinha, Z. Shen, Y. Song, H. Ma, D. Eide, B.-J. P. Hsu, and K. Wang. An overview of microsoft academic service (mas) and applications. In *Proceedings of the 24th International Conference on World Wide Web, WWW '15 Companion*, pages 243–246, New York, NY, USA, 2015. ACM.
- [22] J. Tang, J. Zhang, L. Yao, J. Li, L. Zhang, and Z. Su. Arnetminer: Extraction and mining of academic social networks. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '08*, pages 990–998, New York, NY, USA, 2008. ACM.
- [23] H. D. Tran, G. Cabanac, and G. Hubert. Expert suggestion for conference program committees. In *11th International Conference on Research Challenges in Information Science, RCIS 2017, Brighton, United Kingdom, May 10-12, 2017*, pages 221–232, 2017.