

WikiM: Metapaths based Wikification of Scientific Abstracts

Abhik Jana
Indian Institute of Technology
Kharagpur, India
abhikjana1@gmail.com

Animesh Mukherjee
Indian Institute of Technology
Kharagpur, India
animeshm@gmail.com

Sruthi Mooriyath
SAP Labs India Pvt Ltd
India
sruthi.m01@sap.com

Pawan Goyal
Indian Institute of Technology
Kharagpur, India
pawang.iitk@gmail.com

ABSTRACT

In order to disseminate the exponential extent of knowledge being produced in the form of scientific publications, it would be best to design mechanisms that connect it with already existing rich repository of concepts – the Wikipedia. Not only does it make scientific reading simple and easy (by connecting the involved concepts used in the scientific articles to their Wikipedia explanations) but also improves the overall quality of the article. In this paper, we present a novel *metapath* based method, **WikiM**, to efficiently wikify scientific abstracts – a topic that has been rarely investigated in the literature. One of the prime motivations for this work comes from the observation that, wikified abstracts of scientific documents help a reader to decide better, in comparison to the plain abstracts, whether (s)he would be interested to read the full article. We perform mention extraction mostly through traditional tf-idf measures coupled with a set of smart filters. The entity linking heavily leverages on the rich citation and author publication networks. Our observation is that various *metapaths* defined over these networks can significantly enhance the overall performance of the system. For mention extraction and entity linking, we outperform most of the competing state-of-the-art techniques by a large margin arriving at precision values of **72.42%** and **73.8%** respectively over a dataset from the ACL Anthology Network. In order to establish the robustness of our scheme, we wikify three other datasets and get precision values of **63.41%-94.03%** and **67.67%-73.29%** respectively for the mention extraction and the entity linking phase.

CCS CONCEPTS

•Information systems →Wikis; Entity resolution; Digital libraries and archives;

KEYWORDS

Wikification, Scientific article, Mention Extraction, Entity linking, Metapath, Citation network, Author publication network

ACM Reference format:

Abhik Jana, Sruthi Mooriyath, Animesh Mukherjee, and Pawan Goyal. 2016. WikiM: Metapaths based Wikification of Scientific Abstracts. In *Proceedings of ACM Conference, Washington, DC, USA, July 2017 (Conference'17)*, 10 pages.
DOI: 10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Wikipedia (introduced in 2001) is an online encyclopedia that has evolved as the largest repository of collaboratively curated encyclopedic knowledge having millions of articles, in more than 200 languages. The reliable and refreshed knowledge base of Wikipedia makes it a very popular (Alexa Rank – 7¹) source of knowledge. Consequently, there is an increasing research drive to utilize this knowledge base for better interpretation of terms and expressions in a given text. Wikification, one such usage of Wikipedia introduced by Mihalcea and Csomai [21], is the process of identifying important phrases in a given text (*mention extraction*), and linking each of them to appropriate Wikipedia articles depending on their context of appearances (*entity linking*).

What is available? There has been a substantial amount of work in the literature which focuses on the entity disambiguation and linking task [2, 6, 7, 23–26, 31]. There has also been a lot of literature that contributes to the whole process of end-to-end linking. While some researchers focus on wikification of standard text [15, 16, 20], there have been some efforts to wikify microblog text as well [4, 11, 12, 14]. Wikifier [5] is one such tool for wikification which adopts Integer Linear Programming (ILP) formulation of wikification that incorporates the entity-relation inference problem. Another attempt has been made by Yosef et al. [32], where they propose a graph-based system AIDA, a framework and an online tool for entity detection and disambiguation. Moro et al. [22] presented Babelfy, a graph-based approach which exploits the semantic network structure for entity linking and word sense disambiguation. TagMe is also one of the software systems in this area presented by Paolo and Ugo [9] which can annotate any short, poorly composed fragments of text achieving high on-the-fly accuracy. Thus, the research on entity linking targets on a broad range of text including newswire, spoken dialogues, blogs/microblogs and web documents in multiple languages.

¹<http://www.alex.com/siteinfo/www.wikipedia.org>

What is lacking? Despite such a huge effort by the research community, there have been very few attempts to wikify scientific documents. Most of these studies are made specifically for the bio-medical documents. Some mention extraction tasks like human gene name normalization [8, 10, 13, 30], discovery of scientific names from text [1] as well as the entity linking task in this domain [33] have been attempted by the researchers recently. Some efforts have been made in the geo-science literature [19] as well. These approaches, however, are domain specific and do not go beyond mention detection.

Motivation: Scientific articles are one of the most influential, authorized medium of communication among researchers allowing for appropriate dissemination of knowledge. To make this transmission effective, the terms or phrases used in a particular scientific article need to be understood by the readers without putting too much time or effort. The need is further accelerated as the number of scientific articles are in millions, growing at a rate of approximately 3% annually [3], and there is no well-accepted self-sufficient scientific vocabulary. It would be worthwhile to have an automated system which will do this task of entity linking on scientific articles. Abstract of a scientific article provides a glimpse of the whole article in a nutshell. It is a usual practice among the researchers to go through the abstract first to assess the importance and relevance of the article to their interest. Sometimes the availability of full scientific articles is also under certain terms and conditions (e.g., individual or institutional subscriptions), but the abstracts are publicly available. Thus researchers find it always better to perceive the main idea from the abstract before requesting access from the competent authority for the full article. Further, if a researcher wants to step into a new area of research, wikification could be very helpful to get acquainted with the new terminologies. In order to validate our hypothesis that wikified abstracts help in making better judgments as to whether to read the whole article, we conduct a survey where the wikified and the non-wikified versions of 50 random abstracts from different domains are shown to 10 researchers and 10 post-graduate students. Subsequently, they are asked to judge whether the wikified version of the abstract helped them in forming an overview of the scientific article, better than the non-wikified version. As per their responses, wikification was found to be helpful for 72% cases; this shows the importance of wikification in this problem context. Further, when the researchers were given an article abstract from a domain different from their areas of research, they voted for the wikified version in 100% cases, which straightaway supports our hypothesis. Therefore in this work, we focus on the wikification of scientific article abstracts which we believe should immensely help the readers to decide whether they would at all wish to access the full article (which might have a cost/subscription associated).

We tried this wikification task on scientific abstracts using TagMe [9] that leads to poor results for mention extraction and entity linking with precision of 30.56% and 58.91% respectively. Similarly, other baselines like AIDA [32] and Wikifier [5] also do not perform well, achieving precision values of 8.54%-23.88% and 6%-19.33% respectively for mention extraction and entity linking tasks. This constitutes our motivation behind proposing a novel approach to wikify scientific articles that makes use of a variety of information specific

to scientific texts. In the recent literature there has been burgeoning research where metapaths are closely studied in heterogeneous information network for several mining tasks like similarity search, relationship prediction, clustering etc. [18, 27, 29]. Metapath based similarity [28] forms a common base for any network-based similarity that captures the semantics of peer similarity. Motivated by this stream of literature, we attempt to apply metapath based approach for entity linking which is a natural choice in this case as the citation and author publication networks are perfect sources for construction of meaningful and relevant metapaths.

How we contribute? The main contributions of this paper are summarized as follows:

- This is the first attempt to wikify scientific articles, not specific to any domain. Our method – WikiM – can wikify any scientific article abstract for which we have the underlying citation network, author publication network etc.
- We exploit metapaths between scientific document collections to address the entity linking problem. To extract different types of metapaths, we use the citation network and the author publication network.
- We perform extensive evaluations over 3 different datasets. We find that the proposed system is able to improve upon the existing baselines and also gives good performance across various datasets consistently. We outperform most of the competing state-of-the-art techniques by a large margin arriving at precision values of **72.42%** and **73.80%** respectively for mention extraction and entity linking phase over a gold Standard dataset from the ACL Anthology Network. In order to establish the robustness and consistency of our scheme, we wikify three other datasets and get precision values of **63.41%-94.03%** and **67.67%-73.29%** respectively for the mention extraction and the link disambiguation phase. We make available the code, the results of WikiM and other baselines as well as the gold standard annotations².

2 METHODOLOGY

A two step approach is proposed for wikification of scientific documents. The mentions are the important phrases in the scientific document that are potential terms for wikification. As the first step, given a scientific abstract d , we extract a set of important mentions $M = \{m_1, m_2, \dots, m_n\}$ as described in Section 2.1. In the second step as described in Section 2.2, for each mention m , we extract a list of candidate entities (Wikipedia links) C with surface form similar to m from Wikipedia and rank them according to the similarity scores calculated by metapath-based approach. Finally, we select the entity with the highest score as the appropriate entity for linking.

2.1 Mention extraction

The first step in our approach is to identify important mentions from a scientific document abstract d . After performing tokenization, text normalization is done specific to scientific documents to remove author names and year of publications; for example, Blei et al., 2003; Griffiths, 2004 etc. The next step is to apply POS tagging and identify a set F of textual fragments that contain sequences of

²<https://tinyurl.com/ko84whm>

nouns or adjectives of length up to three. Moreover, overlapping text fragments are handled by giving preference to the larger length fragment if it exists. For example, we do not recognize *language* and *modeling* separately, for *language modeling*.

For each textual fragment $f \in F$, we determine the linking validity corresponding to a decision as to whether it is matched with any Wikipedia concept. We have used python library ‘wikipedia’ which wraps the MediaWiki API for this purpose. If there is no wikipedia article having similar surface form as the textual fragment we are looking for, the library function would return an error message. Only the fragments (a single- or multi-word expression) with positive validity values are taken as the candidate set of mentions M . Then the set of single-word and multi-word mentions are ranked based on the tf-idf score.

For each abstract d , the number of mentions $|M|$ selected by the algorithm depends on the number of sentences in the abstract ($nS(d)$) as follows:

$$|M| = \begin{cases} 4 & nS(d) \leq 2 \\ 8 & 2 < nS(d) \leq 4 \\ 12 & nS(d) > 4 \end{cases}$$

These numbers are chosen after going through a few scientific abstracts manually.

2.2 Entity linking

The second step in our approach is to link the extracted mentions to appropriate Wikipedia entities. The detailed procedure of entity linking is described next.

2.2.1 Candidate entity generation & ranking. For each mention $m \in M$, identified in the previous step, we collect all the Wikipedia entities with surface forms similar to the mention’s surface form, and consider them as candidates $e \in E$ for the mention³.

First, we assign a confidence score to all the candidate entities by using the cosine similarity between the document abstract d and the summary of e ’s Wikipedia article (WS_e). Based on these confidence scores, we prepare a ranked list of all the candidate Wikipedia articles for the given mention m . When the difference of the confidence scores between the highest and the second highest candidate thus produced is less than a given threshold (TH_{cs}), we mark this mapping to be ambiguous at this stage and use a metapath approach for further disambiguation; otherwise we choose the highest ranked candidate. In other words, a difference below the threshold indicates that we are not confident enough with the most agreeable candidate obtained from the cosine similarity method and therefore, resort to the metapath based approach.

2.2.2 Candidate entity ranking using metapaths. As mentioned in Section 2.2.1, whenever the cosine similarity based approach is not very discriminative (less than TH_{cs}) for link disambiguation, we attempt to use a larger context for the particular entity mention. The concept of metapaths is used to identify documents related to the abstract d . These related documents are used to add more context in the process of choosing the most agreeable e from the set E for the given entity mention m . Next, we provide a

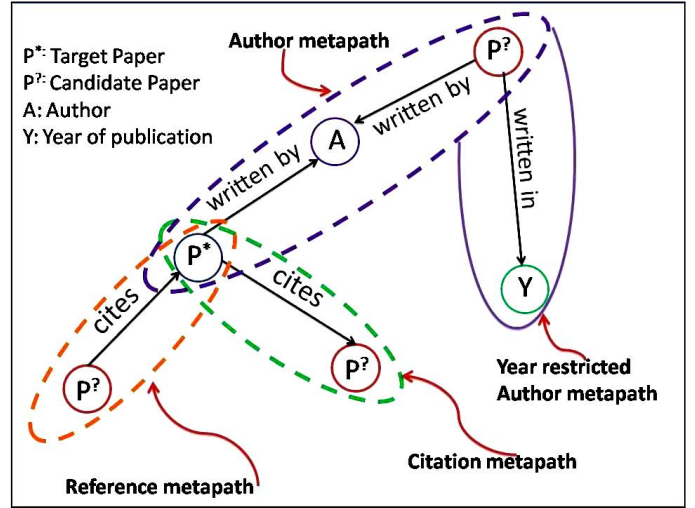


Figure 1: Metapaths in merged Citation and Author Publication Network. The network contains three types of objects: Paper (P), Author (A), Year (Y).

brief overview of the concept of metapaths and how it can be used in the context of wikification of scientific documents.

A metapath is a path composed of sequence of relations between different object types defined over a heterogeneous network structure. We mainly focus on two types of heterogeneous network - citation network and author publication network. Citation network is a network where each node represents a scientific article and edges are of two types - ‘cites’, ‘cited by’. On the other hand, author publication network is a network where nodes are of types ‘author’, ‘scientific article’, ‘year of publication’ etc. and edges are of types ‘written by’, ‘written in’ etc. In our experimental settings, we construct a network as shown in Figure 1, which is the merged version of both the citation network and the author publication network. We use the ACL Anthology Network-2013 dataset to prepare this network, which consists of more than 20000 paper (P) nodes, 17000 author (A) nodes and 40 year (Y) nodes. In this network on an average a paper gets 8 citations, and cites 6 papers; an author writes 3 papers on average, with maximum number of papers written by an author being 161; around 500 papers get published on average per year while the maximum number of papers published per year can go up to 1832. We use this network to define three *basic metapaths* - Citation (C), Reference (R) and Author (A), each of which represents a semantic relation, as follows:

Author metapath (A): $P^* \xrightarrow{w} A \xleftarrow{w} P^?$

Reference metapath (R): $P^? \xrightarrow{c} P^*$

Citation metapath (C): $P^* \xrightarrow{c} P^?$

where P^* is the target paper node in the graph and $P^?$ is the candidate same-author, cited or referenced paper node. For instance, author metapaths try to expand the document using other papers written by the same author A.

³We have used python library ‘wikipedia’ which wraps MediaWiki API for this purpose

| Sample Acronyms | Expansion | Number of dab entries and sections in dab page |
|-----------------|-----------------------------|--|
| IR | Information Retrieval | 44 (5 sections & 2 subsections) |
| MT | Machine Translation | 83 (11 sections & 4 subsections) |
| TC | Text Categorization | 59 (12 sections) |
| DP | Detector of Presuppositions | 81 (10 sections) |

Table 1: Some representative acronym mentions with long dab pages.

Further, we propose one *restricted metapath* as follows, to help choosing only semantically related context-

$$P^* \xrightarrow{w} A \xleftarrow{w'} P^? \xrightarrow{in} Y^*$$

which indicates that the target paper’s (P^*) similar author (A) papers are relevant for consideration only if the paper ($P^?$) is written within a backward window of some years (Y). Such restricted metapaths can be exploited to detect more semantically related contexts. In addition, we incorporate abstracts from the metapaths only if their cosine similarity with the target abstract is greater than a threshold $TH_{relevance}$.

We append the abstracts obtained using metapaths to the target abstract, thus enhancing its scientific context. We then provide a score to each $e \in E$; we take the maximum intersection of possible n -grams of this enhanced context of d with e ’s summary page, and the candidate with maximum intersection is taken as the correct disambiguated candidate entity for that mention m .

2.2.3 Further enhancements based on types of mentions.

On further analysis, we find that many mentions in scientific abstracts are acronyms. As per the Wikipedia style of disambiguation pages (“dab pages”), long disambiguation pages should be grouped into subject sections, and even subsections. These sections (and subsections) should typically be in alphabetical order, e.g.: ‘arts and entertainment’; ‘business’; ‘government and politics’; ‘places’; ‘science and technology’. In terms of organization of the Wikipedia dab pages, we find that acronyms generally tend to have long disambiguation pages. Some examples of acronyms are noted in Table 1. We find that taking intersection from the extended abstract from metapaths sometimes leads to an inappropriate entity linking; therefore, a certain weight needs to be given to the original abstract. On the other hand, non-acronyms tend to have long dab pages (with sections and subsections) very rarely, and thus this problem does not arise.

Thus we propose a separate approach to deal with the acronyms. After finding cosine similarity between the document abstract d and the summary of e ’s Wikipedia article, we incorporate abstracts from metapaths into a new scoring function $Score(n)$, which uses a linear interpolation for the cosine similarity of the original abstract and the expanded context using metapaths as follows:

$$Score(A, B, C) = \alpha * f(A, B) + (1 - \alpha) * f(C, B)$$

where A, B & C are n -dimensional multinomial vectors of d, WS_e and $(d + MP(d))$ respectively where $MP(d) =$

| Mention | Link of Wikipedia articles |
|-----------------------------|---|
| web server | https://en.wikipedia.org/Web_server |
| MT | https://en.wikipedia.org/wiki/Machine_translation |
| natural language processing | https://en.wikipedia.org/wiki/Natural_language_processing |
| WSD | https://en.wikipedia.org/wiki/Word-sense_disambiguation |

Table 2: Examples of annotated mentions in the gold standard dataset.

$\{abstract(metapaths(d))\}$; $f()$ is the cosine similarity function, and α is a constant. The top ranked candidate from the new scoring function $Score()$ is taken as the correct disambiguated link for that mention. As described in Section 3, for acronyms, going with this new scoring function produces better results for majority of the cases.

3 EXPERIMENTAL RESULTS AND DISCUSSIONS

We have broadly two phases in our approach - mention extraction and entity linking cum link disambiguation. First we conduct experiments to evaluate these two phases individually, and then we measure the full system’s performance by considering both the phases together. We evaluate the performance of this system using the standard precision-recall metrics and compare it with various baselines. We compare the approach with mainly three baselines - TagMe [9], AIDA [32] and Wikifier [5], among which TagMe turns out to be the most competitive baseline as per the evaluation results. We also tried Babelfy [22] but it disambiguates all the non-stop keywords (rather than just the relevant keywords). Thus we do not include the evaluation results from Babelfy. As discussed previously, other existing systems, like [12, 14] wikify microblog texts specifically tweets, whereas the baselines that we use in the experimental settings are able to wikify any type of short texts composed of few tens of terms, which makes them applicable for this purpose. The evaluation results for each of the two phases, and comparisons with the baselines are described next.

Gold standard: To evaluate the efficiency of our system WikiM, we prepare the gold-standard data from 2013 ACL Anthology Network (AAN) dataset. 50 random scientific article abstracts are taken from AAN dataset. These 50 abstracts are given to each of the 15 annotators with computational linguistics background, where they are asked to find out the terms or phrases from the abstracts to wikify and then link those terms or phrases with a Wikipedia article, without posing any limit to the number of abstracts they can wikify. On an average each annotator annotated 17 abstracts. Table 2 presents some of these annotated mentions from the gold standard dataset.

Then we use the union⁴ of annotations from multiple annotators in order to aggregate. These manually annotated 50 random scientific article abstracts constitute our gold standard dataset. The full

⁴Taking intersection to aggregate, leads to very few (2-3) mentions per document.

| Abstract Category | # abstracts in the dataset | Avg # single-word mentions (Gold standard/WikiM) | Avg # Multi-word mentions (Gold standard/WikiM) |
|---------------------|----------------------------|--|---|
| 4-mention abstract | 6 | 3.16/2.83 | 1.83/1.17 |
| 8-mention abstract | 25 | 5.72/4.5 | 2.72/3.5 |
| 12-mention abstract | 19 | 7.05/8.67 | 2.58/3.33 |

Table 3: Statistics of single-word and multi-word mentions in the gold standard dataset and in the result of WikiM

gold standard dataset contains an average of 8.5 mentions per abstract. Table 3 gives the statistics of both the single and multi-word mentions in the dataset⁵. In this setting, where each abstract can be wikified by any number of annotators, each of whom can provide any number of mentions, computing the inter-annotator agreement is not relevant. We use 10 random abstracts from this gold standard dataset as the validation dataset for entity linking phase to set the parameter (TH_{cs} , $TH_{relevance}$, α) values. The sensitivity analysis of these parameters and the corresponding performance measures are reported later. We compare the performance of WikiM along with all the baselines against the rest 40 abstracts from the gold standard dataset. Note that we have also tuned the baselines’ parameter values wherever it is possible using validation dataset. For example, TagMe has a parameter ϵ , which can be used to fine tune the disambiguation process. As per the TagMe documentation, a higher value favors the most-common topics (e.g., in tweets, where context may not be much reliable), whereas a lower value uses more context information for disambiguation. TagMe gives the best performance for $\epsilon = 0.2$.

Mention extraction: Statistics from the evaluation of mention extraction phase are presented in Table 4. We see that our mention extraction approach boosts the precision and recall of the system from the range of 8.54%-30.56% to 72.42% and 2.42%-39.03% to 72.1% respectively. Even though our mention extraction approach is very simple, the reason for better results of mention extraction phase could be that we compute idf from the corpus of scientific abstracts only. This clearly shows the advantage of using a system, specifically built for scientific articles. Keyphrase or mention extraction snapshots of all the baselines for a representative abstract are given in Figure 2. It shows that AIDA mainly chooses acronyms as the mentions to wikify, which is true for Wikifier as well, whereas TagMe also chooses other potential mentions besides acronyms. The results from WikiM are shown in Figure 2(d). We see that our approach chooses more appropriate mentions, e.g., ‘scalable’, ‘grid computing’, ‘deployment’, ‘UDDI’ etc., compared to other baselines including both acronyms and non-acronyms. We also test and compare our algorithm on benchmark dataset used in the SemEval-2010 Task-5 on keyword extraction from scientific articles [17] with other state-of-the-art baselines. Even though we are using simple tf-idf based ranking approach for mention extraction, we see comparable

⁵In the statistics acronyms are considered to be single-word mentions

| Method | Precision | Recall | F-Measure |
|----------|---------------|--------------|---------------|
| AIDA | 8.54% | 2.42% | 3.62% |
| Wikifier | 23.88% | 5.76% | 8.31% |
| TagMe | 30.56% | 39.03% | 32.65% |
| WikiM | 72.42% | 72.1% | 71.52% |

Table 4: Evaluation of Mention Extraction w.r.t. AAN 2013 dataset.

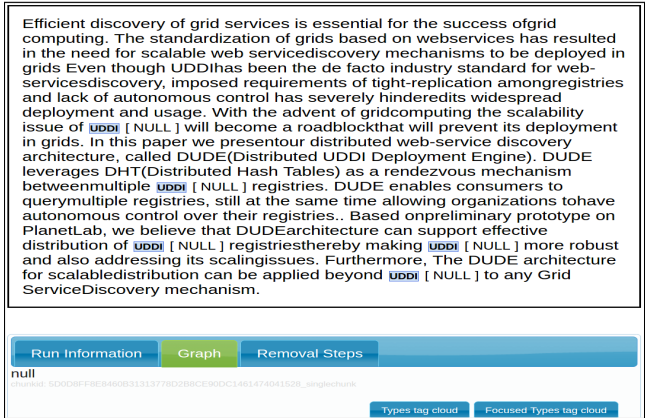
| Method | Link Precision |
|---------------------------------------|----------------|
| AIDA | 6% |
| Wikifier | 19.33% |
| TagMe | 58.91% |
| WikiM - CR metapaths | 69.41% |
| WikiM - Author metapaths | 71.4% |
| WikiM - CRA metapaths | 73% |
| WikiM - Year restricted CRA metapaths | 73.80% |

Table 5: Evaluation of Entity Linking w.r.t. AAN 2013 data set.

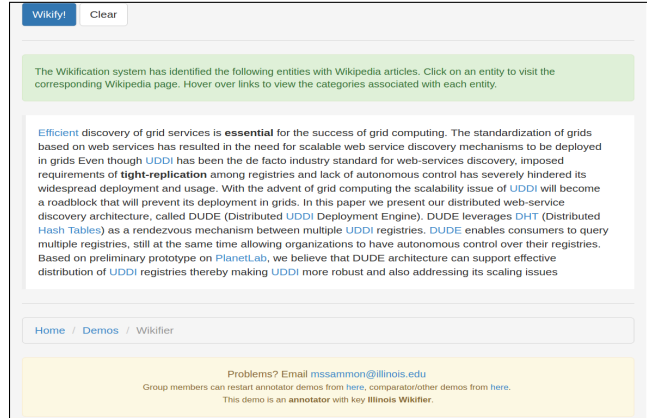
performance in terms of precision and recall with other baselines. As the main focus of our work is entity linking, we do not explore further to improve the mention extraction algorithm which gives an otherwise decent performance.

Entity linking: The comparative evaluation of entity linking is shown in Table 5, evaluated only on those mentions, that are adjudged to be extracted correctly by the system because of the availability of gold standard. The result showcases the link precision (only for true positives) for all these baselines, of which TagMe is again found to be the most competitive one. Note that our system gives significant improvement over TagMe for link precision as well.

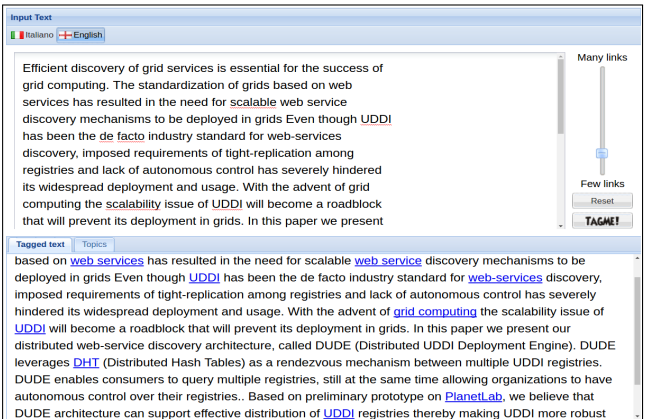
The relatively low performance of the baselines demonstrates that relying on prior popularity within a small document abstract alone may not suffice for the wikification of scientific documents. For instance, it is difficult to link the mention ‘WSD’ from a small abstract, to ‘Word Sense Disambiguation’ without much contextual information. Our approach uses the metapath information from citation and author publication networks, which adds to the context of the seed document. Thus, the chances of linking the mention ‘WSD’ to the most popular concept ‘World Sousveillance Day’ by mistake are by far reduced. We can see that in comparison to all the baselines, our method using author metapaths relying on all the paper abstracts written by the same author as that of the seed paper, achieves comparable performance. For example as shown in Figure 3, the mention ‘tagging’ in a scientific point of view is correctly linked to ‘tag (metadata)’ since other papers written by the same author share similar contexts. We can also see that incorporating citation and reference metapaths into author metapaths provides further gains, indicating the effectiveness of enhanced context using these metapaths. We achieve marginal improvement while using year restricted CRA metapath (with a back-window of 5 years), over the baselines, showing that incorporating related abstracts in the context coupled with timeliness is beneficial as well.



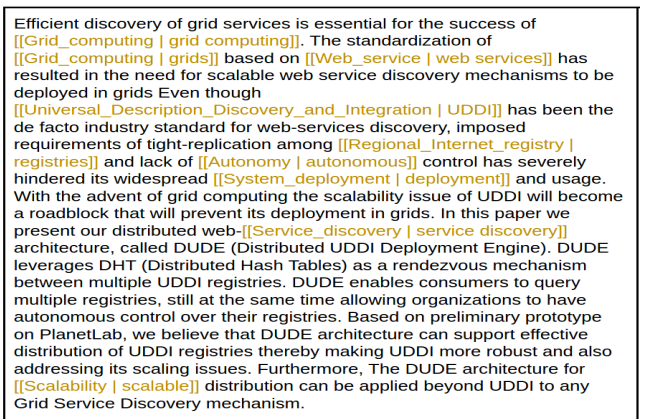
(a) AIDA



(b) Wikifier



(c) TagMe



(d) WikiM

Figure 2: Mention extraction by other baselines and WikiM.

| Baseline | Mention | Link of the gold standard Wikipedia article | Link provided by baseline |
|----------|-----------|---|---|
| AIDA | NLP | https://en.wikipedia.org/wiki/Natural_language_processing | https://en.wikipedia.org/wiki/Nonlinear_programming |
| | internet | https://en.wikipedia.org/wiki/Internet | https://en.wikipedia.org/wiki/YouTube |
| Wikifier | NLG | https://en.wikipedia.org/wiki/Natural_language_generation | https://en.wikipedia.org/wiki/Dutch_guilder |
| | ELS | https://en.wikipedia.org/wiki/ELS_Language_Centers | https://en.wikipedia.org/wiki/London_Buses_route_ELW |
| TagMe | IR | https://en.wikipedia.org/wiki/Information_retrieval | https://en.wikipedia.org/wiki/Infrared |
| | integrate | https://en.wikipedia.org/wiki/System_integration | https://en.wikipedia.org/wiki/Globalization |

Table 6: Sample cases where baselines provide erroneous entity linking.

Cases where the baselines are linking to a wrong entity page, are shown in Table 6. It shows that, in almost all the cases the baselines fail to utilize the context where the mention exists. The probable reason of WikiM’s better performance could be taking into account more contexts by incorporating the abstracts from different metapaths in order to comprehend the context of the mention.

Links given for the mention ‘signal’ (a representative example) by TagMe and WikiM, are shown in Figure 4 and Figure 5 respectively. The correct meaning of the mention ‘signal’ in the context of the

given abstract is ‘signal related to speech or audio’. Since Wikipedia is having quite a few dab page entries for the mention ‘signal’, it is very tricky to choose between them. However, from the snapshots given here, it can be well adjudged that our approach works better than TagMe. TagMe provides ‘signal (electrical engineering)’ as the correct link while there are more appropriate entries in Wikipedia. As Figure 5 shows, WikiM produces the exact meaningful link with respect to the context of the current document. Other baselines such as AIDA and Wikifier do not provide an annotation for the



Figure 3: Links given by various metapaths for the mention ‘Tagging’.

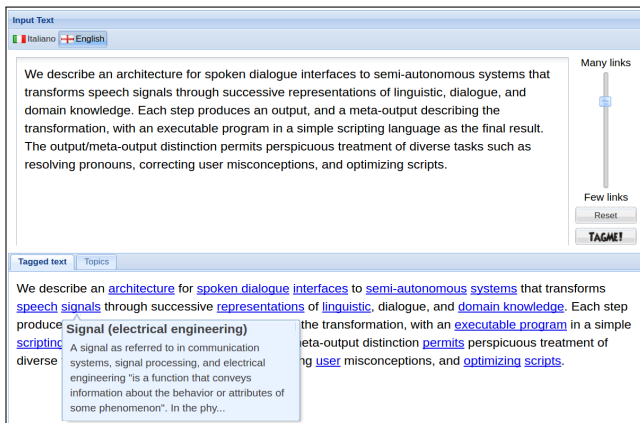


Figure 4: Linking mention ‘Signal’ by TagMe.

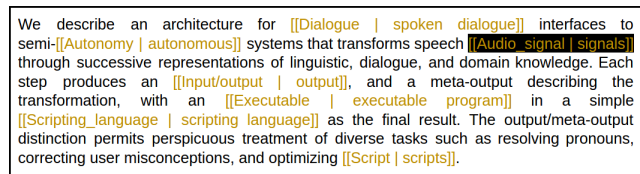


Figure 5: Linking mention ‘Signal’ by WikiM.

| Algorithm | Link Precision |
|------------------------------|----------------|
| Same Algorithm for both | 63.92% |
| Separate algorithms for both | 67.08% |

Table 7: (Validation set) Effect of acronyms & non-acronyms distinction on entity linking.

mention ‘signal’ for this representative example, and therefore are not shown here.

Analysis of the effect of the parameters: We attempt to analyze the importance of various parameters used in our approach using a validation set of 10 random abstracts from the gold standard

| Value of TH_{cs} | Link Precision |
|--|----------------|
| without threshold (always consider metapath) | 66% |
| 0.02 | 59% |
| 0.04 | 62% |
| 0.06 | 67.08% |
| 0.08 | 59% |
| never consider metapath | 63.33% |

Table 8: (Validation set) Effect of variation of parameter TH_{cs} on link precision.

| Value of $TH_{relevance}$ | Link Precision |
|--|----------------|
| 0 (take all the abstracts from metapath) | 64.4% |
| 0.2 | 57.5% |
| 0.4 | 67.08% |
| 0.6 | 60.6% |
| 0.8 | 60.4% |

Table 9: (Validation set) Effect of variation of parameter $TH_{relevance}$ on link precision.

dataset, which we also use to tune the parameters. First, we analyze the effect of using **separate algorithms for acronyms and non-acronyms**. This affects only the link precision. For the AAN dataset, 17 out of 50 document abstracts contained acronyms as mentions in the gold standard and 30% of those acronyms do not have their full forms in the abstract. As shown in Table 7, using separate algorithms for acronyms and non-acronyms gives a link precision of 67.08%, as compared to 63.92% obtained using the same algorithm for both while tested on the validation set. Next, to analyze the effect of the TH_{cs} (difference between the cosine similarity score of top two entities), we conduct experiments with and without keeping TH_{cs} , along with various values for this threshold on the validation set. The results are presented in Table 8. The results show that while the performance is not very sensitive to TH_{cs} , it gives a small improvement and a value of 0.06 gives the best performance. If we do not keep the threshold TH_{cs} , which means that irrespective of the difference of the top two candidate Wikipages’ confidence score, metapath based approach would be taken for further processing, the link precision drops by a small margin. We also see that the link precision drops significantly to 63.3% when metapaths are not considered at all. We can see the similar effect for the variation of parameter $TH_{relevance}$ in Table 9. It shows that the results are quite sensitive, and a value of 0.4 yields the best performance on the validation set.

Similarly, to see the effect of the parameter α in the scoring function for acronyms, we experiment with various values of α as noted in Table 10. Our system gives the best performance for the α value of 0.6.

Using this validation dataset, we set the parameters of the system ($TH_{cs}, TH_{relevance}, \alpha$) to (0.06, 0.4, 0.6) for all the experiments.

Performance dependence on metapath counts: We further analyze the effectiveness of varied proportion of added metapaths. We divide the papers into low, medium and high zones, based on the count of citations. The results presented in Table 11 imply that link precision for the papers from high zone and mid zone increases if we use year restricted methapaths as it reduces the diversity of

| Values for α | Link Precision |
|---------------------|----------------|
| $\alpha = 0.5$ | 64% |
| $\alpha = 0.6$ | 67.08% |
| $\alpha = 0.7$ | 63% |

Table 10: (Validation set) Effect of variation of parameter α on link precision.

| Method | Low (< 5) | Med ($\geq 5, < 10$) | High (> 10) |
|-------------------------------|-----------|------------------------|-------------|
| CRA Metapaths | 73.07% | 64.96% | 75.95% |
| Year restricted CRA Metapaths | 71.75% | 67.81% | 86.67% |

Table 11: Studying the effect of number of citations: citation-zone based evaluation (link precision) of entity linking. The citation-zones are indicated in the parenthesis.

| Method | Overall Recall |
|-----------------------|----------------|
| AIDA | 1% |
| Wikifier | 4.1% |
| TagMe | 25.76 % |
| WikiM - CRA Metapaths | 52.75% |

Table 12: Comparison of WikiM with all the baselines for the full system recall

context. But year restricted CRA metapath approach does not help improve the link precision of low zoned papers as the number of contexts from the metapath are already very few in this case.

Full system recall: A mention and a candidate pair $\langle m, e \rangle$ is considered to be correct if and only if m is linkable and e is its correct candidate concept. Following this definition of correctness, we measure the recall of the full system. The results in Table 12 confirm that our system’s performance is significantly better than all the baselines in terms of overall effectiveness, outperforming the most competing baseline TagMe by a significant margin.

Single annotator result: In order to verify that the comparison results are not due to some bias in our experimental settings, we compare the performance of WikiM and the most competitive baseline TagMe on single annotator results as well. The average measures computed by taking each individual annotator’s annotations as ground truth are given in Table 13. We also see that out of 15 annotators, WikiM does better than TagMe in 12, 11, 12, 9, 10 cases for the measures presented in Table 13 respectively. The consistency of WikiM’s good performance in this experiment shows that the gold standard creation based on the union of the annotations from 15 annotators does not make the evaluation procedure biased.

Error analysis: Even though we achieve significant performance boost over the baselines, we investigate further to point out the erroneous cases for which WikiM links the mention to a wrong Wikipedia entity page and try to find out the possible reasons for the same. As discussed earlier, in order to collect all the Wikipedia entities with surface form similar to the mention’s surface form, we have used python library ‘wikipedia’ which wraps the MediaWiki

API. For this purpose, we use `wikipedia.page()` function which takes as argument a word and returns either the Wikipedia article with the same surface form as the word or a list of Wikipedia articles present in that word’s disambiguation page. Nevertheless, this function has some limitations which leads to poor performance for some cases. Some example cases are shown in Table 14, where for a particular mention both the disambiguation page as well as the page having the same surface form exist. The function in the wikipedia library returns only the page with the same surface form for a given mention, causing no candidate entities to be disambiguated, which leads to linking those mentions wrongly. There are also some other cases, where even though the library function returns a set of Wikipedia articles to disambiguate for a given mention, the gold standard article is not in that set and hence ends in a wrong entity linking. Some of the examples are shown in Table 15. In future, we plan to get rid of these errors by pre-processing the wikidump ourselves instead of using the MediaWiki API.

Further, while we see that for most of the cases extending the target abstract by taking relevant abstracts from the metapath helps in link disambiguation, there are a few cases where this broadening of context leads to wrong links. Table 16 shows some of those examples. The probable reason could be that in all these cases the number of relevant abstracts from the metapath is too large to be informative. The possible solution could be to put another constraint in terms of the importance of a paper while adding abstracts from the metapaths along with textual relevance. The importance of an article could be measured by its citation count, PageRank in the citation graph etc. Note that, this step of incorporating importance to improve the performance of WikiM would however lead to the introduction of one more parameter to the system.

Evaluation of top cited documents from AAN 2013 data set:

To investigate if having many metapaths has any adverse effect on the performance of the system, a dataset consisting of the top 50 cited articles (should usually have a large number of metapaths) is taken from AAN 2013 dataset. Table 17 gives the keyword precision and link precision (for true positives alone) values for the top-50 cited documents from the this dataset. The results of WikiM for each of these document abstracts are evaluated by three annotators independently. For each wikified mention in an abstract, the annotators are asked to choose among the following three options: i). keyword correct and link correct, ii). keyword correct but link incorrect, iii). keyword incorrect. Thus, the link precision is computed only if the annotator responded using one of the first two options. We report the keyword precision and link precision based on three evaluation criteria: **majority decision** is taken from the agreement of at least two out of the three annotators. **Macro-averaged precision** is calculated by first taking the average precision of each abstract (which in turn is calculated as the fraction of annotators who agree to a particular case statement for each keyword, averaged over all the keywords in that abstract) and then taking the average of these. **Micro-averaged precision** results are calculated by computing precision for each wikified mention (fraction of annotators agreeing) and taking an average of all the mentions in all the 50 abstracts. We see that even in this dataset, link precision achieved by WikiM is around 70% for all the three evaluation criteria. The keyword precision is close to 90%

| Method | Precision (Mention Extraction) | Recall (Mention Extraction) | F-Measure (Mention Extraction) | Link Precision | Full System Recall |
|--------|--------------------------------|-----------------------------|--------------------------------|----------------|--------------------|
| TagMe | 18% | 42% | 23% | 54% | 35% |
| WikiM | 26% | 68% | 34% | 64% | 50% |

Table 13: Comparison of WikiM with TagMe while taking individual single annotators as ground truth. Mean values have been reported over 15 annotators.

| Mention | Link of Wikipedia article with same surface form | Link of disambiguation page |
|---------|--|---|
| Java | https://en.wikipedia.org/wiki/Java [an island of Indonesia] | https://en.wikipedia.org/wiki/Java_(disambiguation) |
| Tree | https://en.wikipedia.org/wiki/Tree [a perennial plant with an elongated stem] | https://en.wikipedia.org/wiki/Tree_(disambiguation) |

Table 14: Example mentions having a disambiguation page but the library returns another page (erroneous) with the same surface form.

| Mention | Link of gold standard Wikipedia article | Link of wikipedia article by WikiM |
|---------|---|---|
| modify | https://en.wikipedia.org/wiki/Editing | https://en.wikipedia.org/wiki/Modding |
| precise | https://en.wikipedia.org/wiki/Accuracy_and_precision | https://en.wikipedia.org/wiki/Promise |

Table 15: Example cases where gold standard Wikipedia article is not in the set of articles to be disambiguated returned by 'wikipedia' library

| Mention | Link of gold standard wikipedia article | Link of wikipedia article by WikiM | Number of abstracts in the metapath |
|----------------|---|---|-------------------------------------|
| schemes | https://en.wikipedia.org/wiki/Scheme_(linguistics) | https://en.wikipedia.org/wiki/Scheme_(Mathematics) | 72 |
| transcriptions | https://en.wikipedia.org/wiki/Orthographic_transcriptions | https://en.wikipedia.org/wiki/Transcription_(linguistics) | 38 |
| romanian | https://en.wikipedia.org/wiki/Romanian_language | https://en.wikipedia.org/wiki/Romania | 33 |

Table 16: Example cases where too many abstracts from metapaths lead to wrong entity linking.

using any of the evaluation criteria. Thus, the performance of this

| Method | Mention Extraction Precision | Link Precision |
|-------------------|------------------------------|----------------|
| Majority Decision | 94.03% | 73.23% |
| Macro-Averaging | 89.36% | 71.11% |
| Micro-Averaging | 88.81% | 69.53% |

Table 17: Evaluation results for Top Cited articles from AAN 2013 data set.

| Method | Mention Extraction Precision | Link Precision |
|-------------------|------------------------------|----------------|
| Majority Decision | 67.09% | 67.67% |
| Macro-Averaging | 65.98% | 71.8% |
| Micro-Averaging | 63.41% | 67.71% |

Table 18: Evaluation results for data mining articles from MAS data set.

| Method | Mention Extraction Precision | Link Precision |
|-------------------|------------------------------|----------------|
| Majority Decision | 86.6% | 73.29% |
| Macro-Averaging | 82.37% | 69.55% |
| Micro-Averaging | 85.04% | 69.33% |

Table 19: Evaluation results for bio-medical data set.

system is quite consistent even on the dataset with many metapaths.

Evaluation of data mining documents from MAS dataset: To further verify the consistency of our approach, we test its performance on a different dataset of scientific articles. 50 scientific abstracts are taken from Microsoft Academic Search (MAS) dataset⁶ in the domain of 'data mining' and are wikified using WikiM. The evaluation framework is similar to the one used for the top cited articles from AAN dataset. Each of these abstracts is evaluated by three annotators independently. The results presented in Table 18 show that even when we move on to the domain of 'data mining', the performance of the system is quite consistent. The link precision is at least 67% using any of the evaluation criteria which is consistent with the evaluation results of the gold standard data. As per the majority decision, the keyword precision is 67.67%. Since we use Wikipedia as the knowledge base, the links provided by WikiM confine to the existing Wikipedia entries. There are cases where all the currently existing candidate pages in Wikipedia are inappropriate for the mention with respect to the context of the document under observation. For instance, mentions such as 'splits' from the AAN data set and 'scientific data' from data mining domain of MAS data set, do not have valid Wikipedia pages for their corresponding meaning; thus, the links proposed by WikiM are labeled as incorrect by the annotators.

⁶<https://academic.microsoft.com/>

Evaluation of bio-medical dataset: We further test our system on a completely different dataset of scientific articles. 50 scientific article abstracts taken from the bio-medical dataset⁷ are wikified using WikiM. According to the results presented in Table 19, we see that moving into completely different domain does not affect the consistency of WikiM much.

4 CONCLUSION

This paper addresses a novel wikification approach for scientific articles. We use a tf-idf based approach to find out important terms to wikify. Then we facilitate the step of ranking the candidate links with metapaths extracted from citation and author publication networks of scientific articles. Experimental results show that the proposed approach helps to significantly improve the performance of the wikification task on scientific articles. The performance of our system is tested across various datasets, and the results are found to be consistent. We plan to do similar experiments on a larger population and more variety of abstracts in future to further strengthen the necessity of wikification. Immediate future work should focus on some additional methods to detect the mentions more precisely, which may lead to overall improved entire system's performance. Our future plan is to incorporate the metapath based approach for mention extraction phase as well, and study the scope of metapath based joint mention extraction and entity linking system for this purpose.

REFERENCES

- [1] Lakshmi Manohar Akella, Catherine N Norton, and Holly Miller. 2012. NetiNeti: discovery of scientific names from text using machine learning methods. *BMC bioinformatics* 13, 1 (2012), 1.
- [2] Ayman Alhelbawy and Robert J Gaizauskas. 2014. Graph Ranking for Collective Named Entity Disambiguation. In *ACL (2)*, 75–80.
- [3] Lutz Bornmann and Rüdiger Mutz. 2015. Growth rates of modern science: a bibliometric analysis based on the number of publications and cited references. *Journal of the Association for Information Science and Technology* 66, 11 (2015), 2215–2222.
- [4] Taylor Cassidy, Heng Ji, Lev-Arie Ratino, Arkaitz Zubiaga, and Hongzhao Huang. 2012. Analysis and Enhancement of Wikification for Microblogs with Context Expansion. In *COLING*, Vol. 12, 441–456.
- [5] Xiao Cheng and Dan Roth. 2013. Relational inference for wikification. *Urbana* 51 (2013), 61801.
- [6] Alexandre Davis, Adriano Veloso, Altigran S Da Silva, Wagner Meira Jr, and Alberto HF Laender. 2012. Named entity disambiguation in streaming data. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*. Association for Computational Linguistics, 815–824.
- [7] Gianluca Demartini, Djellel Eddine Difallah, and Philippe Cudré-Mauroux. 2012. ZenCrowd: leveraging probabilistic reasoning and crowdsourcing techniques for large-scale entity linking. In *Proceedings of the 21st international conference on World Wide Web*. ACM, 469–478.
- [8] Haw-ren Fang, Kevin Murphy, Yang Jin, Jessica S Kim, and Peter S White. 2006. Human gene name normalization using text matching with automatically extracted synonym dictionaries. In *Proceedings of the Workshop on Linking Natural Language Processing and Biology: Towards Deeper Biological Literature Analysis*. Association for Computational Linguistics, 41–48.
- [9] Paolo Ferragina and Ugo Scaiella. 2010. Tagme: on-the-fly annotation of short text fragments (by wikipedia entities). In *Proceedings of the 19th ACM international conference on Information and knowledge management*. ACM, 1625–1628.
- [10] Matthias Frisch, Bernward Klocke, Manuela Haltmeier, and Kornelie Frech. 2009. LitInspector: literature and signal transduction pathway mining in PubMed abstracts. *Nucleic acids research* 37, suppl 2 (2009), W135–W140.
- [11] Yegin Genc, Yasuaki Sakamoto, and Jeffrey V Nickerson. 2011. Discovering context: classifying tweets through a semantic transform based on wikipedia.

- In *Foundations of augmented cognition. Directing the future of adaptive systems*. Springer, 484–492.
- [12] Stephen Guo, Ming-Wei Chang, and Emre Kiciman. 2013. To Link or Not to Link? A Study on End-to-End Tweet Entity Linking. In *HLT-NAACL*, 1020–1030.
 - [13] Lynette Hirschman, Marc Colosimo, Alexander Morgan, and Alexander Yeh. 2005. Overview of BioCreAtIvE task 1B: normalized gene lists. *BMC bioinformatics* 6, 1 (2005), 1.
 - [14] Hongzhao Huang, Yunbo Cao, Xiaojiang Huang, Heng Ji, and Chin-Yew Lin. 2014. Collective Tweet Wikification based on Semi-supervised Graph Regularization. In *ACL (1)*, 380–390.
 - [15] Heng Ji, Ralph Grishman, Hoa Trang Dang, Kira Griffitt, and Joe Ellis. 2010. Overview of the TAC 2010 knowledge base population track. In *Third Text Analysis Conference (TAC 2010)*, Vol. 3, 3–3.
 - [16] Heng Ji, Joel Nothman, and Ben Hachey. 2014. Overview of tac-kbp2014 entity discovery and linking tasks. In *Proc. Text Analysis Conference (TAC2014)*.
 - [17] Su Nam Kim, Olena Medelyan, Min-Yen Kan, and Timothy Baldwin. 2010. SemEval-2010 Task 5: Automatic Keyphrase Extraction from Scientific Articles. In *Proceedings of the 5th International Workshop on Semantic Evaluation (SemEval '10)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 21–26. <http://dl.acm.org/citation.cfm?id=1859664.1859668>
 - [18] Xiaozhong Liu, Yingying Yu, Chun Guo, and Yizhou Sun. 2014. Meta-Path-Based Ranking with Pseudo Relevance Feedback on Heterogeneous Graph for Citation Recommendation. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management (CIKM '14)*. ACM, New York, NY, USA, 121–130. DOI : <http://dx.doi.org/10.1145/2661829.2661965>
 - [19] Xiaogang Ma. 2015. Illuminate Knowledge Elements in Geoscience Literature. In *2015 AGU Fall Meeting*. Agu.
 - [20] Paul McNamee and Hoa Trang Dang. 2009. Overview of the tac 2009 knowledge base population track. In *Text Analysis Conference (TAC)*, Vol. 17, 111–113.
 - [21] Rada Mihalcea and Andras Csomai. 2007. Wikify!: linking documents to encyclopedic knowledge. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*. ACM, 233–242.
 - [22] Andrea Moro, Alessandro Raganato, and Roberto Navigli. 2014. Entity linking meets word sense disambiguation: a unified approach. *Transactions of the Association for Computational Linguistics* 2 (2014), 231–244.
 - [23] Xiaoman Pan, Taylor Cassidy, Ulf Hermjakob, Heng Ji, and Kevin Knight. 2015. Unsupervised entity linking with abstract meaning representation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics—Human Language Technologies*.
 - [24] Maria Pershina, Yifan He, and Ralph Grishman. 2015. Personalized Page Rank for named entity disambiguation. In *Proc. 2015 Annual Conference of the North American Chapter of the ACL, NAACL HLT*, Vol. 14, 238–243.
 - [25] Lev Ratino, Dan Roth, Doug Downey, and Mike Anderson. 2011. Local and global algorithms for disambiguation to wikipedia. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*. Association for Computational Linguistics, 1375–1384.
 - [26] Avirup Sil, Ernest Cronin, Penghai Nie, Yinfei Yang, Ana-Maria Popescu, and Alexander Yates. 2012. Linking named entities to any database. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. Association for Computational Linguistics, 116–127.
 - [27] Yizhou Sun and Jiawei Han. 2012. *Mining Heterogeneous Information Networks: Principles and Methodologies*. Morgan & Claypool Publishers.
 - [28] Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S. Yu, and Tianyi Wu. 2011. Pathsim: Meta path-based top-k similarity search in heterogeneous information networks. In *In VLDB* 11.
 - [29] Yizhou Sun, Brandon Norick, Jiawei Han, Xifeng Yan, Philip S. Yu, and Xiao Yu. 2013. PathSelClus: Integrating Meta-Path Selection with User-Guided Object Clustering in Heterogeneous Information Networks. *ACM Trans. Knowl. Discov. Data* 7, 3, Article 11 (Sept. 2013), 23 pages. DOI : <http://dx.doi.org/10.1145/2500492>
 - [30] Yu Usami, Han-Cheol Cho, Naoaki Okazaki, and Jun'ichi Tsujii. 2011. Automatic acquisition of huge training data for bio-medical named entity recognition. In *Proceedings of BioNLP 2011 Workshop*. Association for Computational Linguistics, 65–73.
 - [31] Chi Wang, Kaushik Chakrabarti, Tao Cheng, and Surajit Chaudhuri. 2012. Targeted disambiguation of ad-hoc, homogeneous sets of named entities. In *Proceedings of the 21st international conference on World Wide Web*. ACM, 719–728.
 - [32] Mohamed Amir Yosef, Johannes Hoffart, Ilaria Bordino, Marc Spaniol, and Gerhard Weikum. 2011. Aida: An online tool for accurate disambiguation of named entities in text and tables. *Proceedings of the VLDB Endowment* 4, 12 (2011), 1450–1453.
 - [33] Jin G Zheng, Daniel Howson, Boliang Zhang, Juergen Hahn, Deborah McGuinness, James Hendler, and Heng Ji. 2015. Entity linking for biomedical literature. *BMC medical informatics and decision making* 15, Suppl 1 (2015), S4.

⁷It contains more than 1.1 million Bio-medical articles, downloaded from http://www.ncbi.nlm.nih.gov/pmc/tools/ftp/#XML_for_Data_Mining