

WikiRef: Wikilinks as a route to recommending appropriate references for scientific Wikipedia pages

Abhik Jana
IIT Kharagpur, India
abhik.jana@iitkgp.ac.in

Pranjal Kanojiya
IIT Kharagpur, India
QUALCOMM India Private Limited
pranjal989091@gmail.com

Pawan Goyal
IIT Kharagpur, India
pawang@cse.iitkgp.ac.in

Animesh Mukherjee
IIT Kharagpur, India
animeshm@gmail.com

Abstract

The exponential increase in the usage of Wikipedia as a key source of scientific knowledge among the researchers is making it absolutely necessary to metamorphose this knowledge repository into an integral and self-contained source of information for direct utilization. Unfortunately, the references which support the content of each Wikipedia entity page, are far from complete. Why are the reference section ill-formed for most Wikipedia pages? Is this section edited as frequently as the other sections of a page? Can there be appropriate surrogates that can automatically enhance the reference section? In this paper, we propose a novel two step approach – **WikiRef** – that (i) leverages the wikilinks present in a scientific Wikipedia target page and, thereby, (ii) recommends highly relevant references to be included in that target page appropriately and automatically borrowed from the reference section of the wikilinks. In the first step, we build a classifier to ascertain whether a wikilink is a potential source of reference or not. In the following step, we recommend references to the target page from the reference section of the wikilinks that are classified as potential sources of references in the first step. We perform an extensive evaluation of our approach on datasets from two different domains – Computer Science and Physics. For Computer Science we achieve a notably good performance with a *precision@1* of 0.44 for reference recommendation as opposed to 0.38 obtained from the most competitive baseline. For the Physics dataset, we obtain a similar performance boost of 10% with respect to the most competitive baseline.

1 Introduction

Wikipedia, the largest online encyclopedia, is a collaborative work of 33,778,487 users and 1,210 admins, with 5,662,889 English articles and a total of 45,132,517 articles in more than 250 languages as of June 7, 2018¹. Introduced in 2001, it has become one of the most visited websites having global Alexa Rank – 5². Each Wikipedia article contains a rich collection of concepts along with a set of hyperlinks associating crucial terms to other wiki pages (termed as wikilinks), which makes the article more comprehensive. In addition, the ability to edit the entity pages easily with proper source citations and to view the edits reflected in the pages timely are some of the key reasons behind Wikipedia’s growth and popularity. As a consequence, Wikipedia has become a destination of all kinds of information about

This work is licenced under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>

¹https://en.wikipedia.org/wiki/Wikipedia:Size_of_Wikipedia

²<http://www.alexa.com/siteinfo/www.wikipedia.org>

entities and events. Further, Wikipedia is being widely utilized by many applications focused on entity linking and disambiguation (Moro et al., 2014; Hachey et al., 2013; Cucerzan, 2007; Hoffart et al., 2011), named entity disambiguation (Cucerzan, 2007; Barrena et al., 2015), semantic similarity (Gabrilovich and Markovitch, 2007; Wu and Giles, 2015), information extraction (Wu and Weld, 2010) etc.

What is lacking? Nevertheless, there are different competing opinions among the researchers regarding the reliability, integrity and usage of Wikipedia as an authentic source of scientific information. While many researchers consider Wikipedia as a valuable resource, others question the accuracy, comprehensiveness and completeness of the articles. For instance, according to Gorman (2007), Wikipedia is an “unethical resource unworthy of our respect” whereas Orłowski (2005) questions that the technology which makes Wikipedia possible is not a substitute for expert editors, and poor writing will remain unavoidable without any expertise. Chesney (2006) conducted an experiment to evaluate the credibility of Wikipedia and found that the studies focusing only on history articles provide mixed evidence concerning its accuracy. The general solution is to post-process the edits done by the users. In order to ensure the accuracy and quality of writing of articles, Wikipedia has imposed a new policy of requiring senior editors to approve changes to articles.

Even then, the completeness of the article is not fully ensured. Not all Wikipedia pages referring to entities (entity pages) are comprehensive: relevant information can either be missing or added with a delay. In addition, the frequency of editing of different sections in a given Wikipedia article varies extensively. For instance, frequency of editing text is higher than frequency of editing wikilinks or references; also add/edit references is considered as the hardest task³. This opens up the requirement for an automated system which increases the number of relevant references by either processing the content of the target article itself or by appropriately collating references from alternative sources so as to make the overall information content of the article more complete.

State-of-the-art: Efforts have been made by the researchers to populate Wikipedia pages automatically. Sauper and Barzilay (2009) propose an approach for populating Wikipedia pages with content coming from external sources by automatically generating whole entity pages for specific entity classes. Taneva and Weikum (2013) propose an approach to generate novel summaries from the external information which could be added to Wikipedia entity pages. Balog and Ramampiaro (2013) and Balog et al. (2013) try to recommend news citations for an entity in Wikipedia. West et al. (2014) focuses on the problem of knowledge base completion, through question answering and tries to complete missing facts in Freebase based on templates. All these works significantly rely on high quality input sources which are utilized to extract textual facts for Wikipedia page population. In one of the recent works, Fetahu et al. (2015) propose a two-stage supervised approach for suggesting news articles corresponding to entity pages. Attempt has also been made to introduce appropriate wikilinks automatically (İkikat et al., 2015). In a similar line, Raganato et al. (2016) present the automatic construction and evaluation of a Semantically Enriched Wikipedia (SEW) in which the overall number of wikilinks has been more than tripled. Efforts have also been made to generate Wikipedia articles of named entities in a semi-supervised framework (Pochampally et al., 2016).

Motivation: Recently, with the increasing number of scientific articles and with the need for getting an overall view of a particular scientific topic within a very less amount of time, researcher’s tendency of referring to the Wikipedia pages instead of going through a set of very specific scientific articles is massively increasing. Given that, such a huge community of researchers rely on Wikipedia for scientific information, it becomes absolutely necessary to deeply focus on the completeness of the Wikipedia articles which talk about some scientific topics.

In the following we enlist some initial observations to show that there is a huge scope of improvement of the scientific Wikipedia articles through the inclusion of a proper set of references. Figure 1 illustrates the snapshot of the reference section of the ‘Bigram’ page⁴ in the year 2018. There are only five references present in the page; very relevant scientific articles like, “Statistical Identification of Language” (Dunning, 1994), “Foundations of Statistical Natural Language Processing” (Manning and

³https://www.mediawiki.org/wiki/Editing_Tasks_Survey

⁴<https://en.wikipedia.org/wiki/Bigram>

Schütze, 1999), etc. are missing, without which the article seems to be incomplete. Also from the survey conducted by wikimedia ⁵ it has been observed that adding or editing the reference section is harder compared to adding or editing simple text or a wikilink which therefore leads to much less frequent changes of the reference section. In order to further confirm this trend, we conduct an analysis on the edit history of 1120 Wikipedia articles from the Computer Science category and observe that till Jan, 2017 on average around 65% edits are in the text content, 32% are wikilink edits and strikingly only 1% are reference edits. Rest correspond to table, template and category edits etc. We present some example Wikipedia articles along with the statistic of their edit history in Table 1. Clearly, all these observations point to the requirement of an automated reference recommendation system to keep the reference section of a scientific Wikipedia article up-to-date.




Figure 1: Snapshot of Wikipedia ‘Bigram’ in the year 2018.

| Wikipage title | Start date | End date | #(Total edits) | #(Wikilink edits) | #(Ref. edits) |
|--------------------|------------|----------|----------------|-------------------|---------------|
| Connection Machine | 7/8/02 | 6/1/17 | 242 | 98 | 3 |
| Fuzzy Logic | 15/4/02 | 29/1/17 | 1343 | 390 | 5 |
| Javascript | 5/2/02 | 31/1/17 | 4141 | 1146 | 69 |

Table 1: Representative scientific Wikipedia articles with detailed statistic of their edit history. Start date indicates the date of creation of Wikipage and end date indicates the time point till which we perform the analysis of edit history.

What we envisage? Motivated by the fact that wikilink edits are less hard and more frequent compared to reference edits, a natural question that arises is that whether it is possible to develop an automatic procedure to add a set of references whenever there is a wikilink edit on a target page. In this paper, we investigate the possibility of enriching the reference section of a scientific Wikipedia article using the wikilinks present in the article as a surrogate. In particular, *we introduce the novel idea of reference inheritance from the source pages pointed to by the wikilinks present on the target page to populate the reference section of the target page with relevant references.*

For this study, we consider 3842 target Wikipedia pages (till June, 2017) from the Computer Science domain. In order to show that our method is quite generic, we also consider an additional dataset comprising 2871 target Wikipedia pages (till February, 2018) from the Physics domain.

We propose a novel two step approach as follows. In the first step, *we determine the potential wikilinks from which references could be inherited.* In the second step, *we recommend the most relevant references from the potential wikilinks obtained in the first step.*

Key contributions: In summary, we make the following contributions:

a) *Novel problem formulation:* We formulate and address the problem of automatically populating the reference section by inheriting the references from the wikilinks present in a target scientific Wikipedia article.

b) *The two step approach:* We propose a two-step reference recommendation system and show its advantage over simple baselines that do not implement the first step of potential wikilink selection from which

⁵https://commons.wikimedia.org/wiki/File:WMF_editing_tasks_survey_July_2015.pdf

references could be inherited.

c) *Evaluation*: We perform extensive experiments to evaluate our system. We evaluate in two different ways - first, by comparing our recommendations with the references already present on the target page (acting as a ground-truth) and second, through human judgement where experts are tasked to judge the appropriateness of the recommended references. For the automatic evaluation we obtain a *precision@1* of 0.44 for the Computer Science dataset compared to 0.38 obtained for the most competitive baseline. For the Physics dataset we obtain a similar performance boost of 10% over the best baseline. Human judgement experiments show that the average Spearman’s rank correlation coefficient (ρ) between the ranked list returned by WikiRef and the one produced by averaging the responses of the annotators is **0.203** compared to 0.168 obtained for the most competitive baseline⁶.

2 Problem formulation

Let us assume that we have a scientific Wikipedia article A (i.e., target) whose reference section we wish to populate using a set of relevant references. As an input we use only the text content and wikilinks present in A . Suppose, A has a set of n wikilinks B_1, B_2, \dots, B_n which in this case simply correspond to the n Wikipedia source articles from which A could possibly inherit the references. We assume that while the reference sections of the individual source pages themselves may not be “well populated”, together the reference sections of all the source articles may provide a potential set of references, relevant to A . We divide this task in two major steps. First, we obtain the wikilinks (i.e., a set of B s) that are appropriate for inheriting references from. We pose this as a binary classification problem thus dividing the set of B s in those that are relevant (appropriate for inheriting references) vs those that are not. Second, we prepare a ranked list of k references taken from each of the correctly classified (i.e., relevant) wikilinks which could be added to the reference section of A . We pose this second step as a “learning-to-rank” problem and use various features to prepare the final list of recommendations.

3 Dataset

We use datasets from two different domains – Computer Science (CS) and Physics (PH) – to evaluate our system. Note that while CS is our primary focus (owing to the background of the authors) we use the PH dataset to show that our method is quite generic.

CS dataset: In order to prepare the CS dataset we crawl a set of 3842 target Wikipedia pages till June, 2017. The topics we span in this crawl are information retrieval, machine learning, automata theory, graph theory etc. In addition, we also crawl all the source pages corresponding to all the wikilinks present in the 3,842 target articles. The total number of wikipages thus crawled is 121,154.

PH dataset: For the PH dataset we crawl a total of 2,871 target articles spanning topics like gravitation, mechanics, motion etc. The source pages pointed to by the wikilinks present on these target articles together make the total crawl size of 107,332 pages.

4 Building WikiRef

Recall, that our approach works in two steps. We present the details of each of these steps below which together refers to our system – WikiRef.

4.1 Classification of wikilinks (Step - I)

As discussed in Section 2, we have a target Wikipedia article A with n wikilinks (B_1, B_2, \dots, B_n) present in it. Our goal in this step is to ascertain how appropriate a source page pointed to by a wikilink B_i is for inheriting references. To estimate this, we measure the extent of similarity between A and the page pointed to by B_i . The hypothesis is that the higher this similarity, the more relevant is the source page to A and thus higher is the propensity of reference inheritance. We define various notions of similarity that act as features to a binary classifier which classifies if the source is relevant or not.

⁶We perform the human judgement experiment only for the CS domain because of the background of the contributing authors of this paper.

- **Tf-idf similarity of article summaries (TIS):** We represent the summaries (first paragraphs) of A and the source page pointed to by B_i as tf-idf vectors and compute the cosine similarity between them. Note that, inverse document frequency (idf) is calculated only from the dataset we have prepared.
- **Outlinks similarity (OS):** Outlinks of a Wikipedia article (say A) is the set of Wikipedia articles that A hyperlinks to. We compute outlinks similarity as the Jaccard overlap between the outlinks of A and the source page pointed to by B_i .
- **Inlinks similarity (IS):** Inlinks of a Wikipedia article (say A) is the set of Wikipedia articles that hyperlink to the page A . We compute inlinks similarity as the Jaccard overlap between the inlinks of A and the source page pointed to by B_i . Note that, for computing inlinks we consider the articles present only in our dataset.
- **Out sentence similarity (OSS):** Here we consider the common outlinks of A and the source page pointed to by B_i . We collate all the sentences in A where these common outlinks occur, to prepare a document OS_A . Similarly, we collate all sentences in the page pointed to by B_i where these common outlinks occur to form a document OS_{B_i} . Now we represent OS_A and OS_{B_i} as tf-idf vectors and compute cosine similarity between them.
- **In sentence similarity (ISS):** This is same as OSS with the exception that here we construct the documents based on the sentences where the common inlinks of A and the page pointed to by B_i occur.

Additional deep neural representations: Apart from these five similarity measures we use additional variants of TIS, OSS and ISS where instead of representing a text document as a tf-idf vector we represent it as a document vector obtained using an encoder based on a bi-directional LSTM architecture with max pooling, trained on the Stanford Natural Language Inference (SNLI) dataset (Bowman et al., 2015) as proposed by Conneau et al. (2017). Note that, for applying this architecture we use the GloVe vectors trained on Common Crawl data (840B tokens)⁷ as seeds for representing words in a document. We name these variants of TIS, OSS and ISS respectively as vector similarity of article summaries (VS), out-sentence vector similarity (OSVS) and in-sentence vector similarity (ISVS) respectively.

We plug in all these features in classifiers like SVM, Random Forest, Logistic Regression etc. We observe that Random Forest performs the best among these and therefore report the results for this classifier only. This step automatically identifies the wikilinks that are potential candidates for reference inheritance.

4.2 Recommending the exact list of references (Step - II)

As an outcome of the previous step, we have a target Wikipedia article A and a wikilink B_i which is classified as either relevant for reference inheritance or not. Let us assume that for a B_i that is classified as relevant in the first step, the source page pointed to by B_i has m references R_1, R_2, \dots, R_m . However, all of these m references might not be appropriate to inherit. Therefore, the task here is to produce a ranked list of k references which A can inherit depending on the relevance of the reference with respect to the content of A . In order to estimate the relevance of a reference (say R_j) we propose the following features.

- **Similarity between the citation context of the reference R_j in the source page pointed to by B_i and the citation context of B_i in A :** Citation context represents the sentence in which a reference or wikilink gets cited. We represent the citation contexts as simple tf-idf vectors and compute the cosine similarity between the tf-idf vectors of the citation context of R_j in the source page pointed to by B_i and the citation context of B_i in A . We term this feature as **F1-TI**. We also represent citation contexts as sentence vectors using the bi-LSTM architecture proposed by Conneau et al. (2017) (already discussed earlier) and compute the cosine similarity of these vector representations of the citation contexts of R_j in the source page pointed to by B_i and the citation context of B_i in A . We call this feature **F1-Vec**.
- **Similarity between the title of the reference R_j in the source page pointed to by B_i and the citation context of B_i in A :** As the title of any reference contains the most important clue about the

⁷<https://nlp.stanford.edu/projects/glove/>

reference itself, we represent the title of the reference R_j in the source page pointed to by B_i and citation context of B_i in A as tf-idf vectors and compute the cosine similarity between them. We call this feature **F2-TI**⁸.

We use – SVMRank (Joachims, 2006) – a standard learning-to-rank framework and plug in the above three features to obtain a ranked list of k references to recommend. We report the performance our approach for different values of k in Section 5.

5 Evaluation

We first individually evaluate Step - I and Step - II and subsequently report the performance of WikiRef as a whole. We evaluate the performance using the standard precision-recall metrics and compare it with various baseline approaches that we ourselves propose since there is no known work in the literature that could serve as a suitable baseline for this task.

5.1 Gold standard dataset

As discussed in section 3 we have a total of 3,842 target Wikipedia pages from the Computer Science (CS) domain and 2,871 target Wikipedia pages from the Physics (PH) domain for which we aim to come up with a list of references. We consider the existing references in the target pages (i.e., the A pages) at the time-point of crawling as the gold standard references to be evaluated against.

5.2 Baselines

There have been several reference recommendation systems (Huang et al., 2015; Caragea et al., 2013; Huang et al., 2014; Tang et al., 2014; Ren et al., 2014; Meng et al., 2013; Huang et al., 2012) for scientific articles but all of them either use the text of the cited paper or the underlying citation network which makes these systems inappropriate to be used as baselines to compare with WikiRef. Therefore, we propose some standard baselines to ascertain the importance of each step in our approach.

Baseline I: Here, Step - I of our approach is skipped and we assume that all the wikilinks present in the target Wikipedia article are relevant for reference inheritance. In Step - II we rank the references of each wikilink only on the basis of tf-idf similarity between references' citation context in the page pointed to by the wikilink and the wikilink's citation context in target Wikipedia article.

Baseline II: Here again, Step - I of our approach is skipped and we assume that all the wikilinks present in the target Wikipedia article are relevant for reference inheritance. In Step - II we rank the references of each wikilink only on the basis of tf-idf similarity between the references' title and the wikilink's citation context in the target Wikipedia article.

Baseline III: Step - I is fully retained. In Step - II we rank the references of each wikilink only on the basis of tf-idf similarity between references' citation context in wikilink's content and wikilink's citation context in target Wikipedia article (F1-TI).

Baseline IV: Step - I is fully retained. In Step - II we rank the references of each wikilink only on the basis of tf-idf similarity between references' title and wikilink's citation context in the target Wikipedia article (F2-TI).

Baseline V: Step - I is fully retained. In Step - II we rank the references of each wikilink only on the basis of cosine similarity between sentence vector representation of references' citation context in the page pointed to by the wikilink and the wikilink's citation context in the target Wikipedia article (F1-Vec).

We compare the performance of these baselines with WikiRef in section 5.5.

5.3 Performance analysis of Step - I

We use 70% of the Wikipedia pages for training and rest 30% of them for testing using the gold standard dataset. However, we observe from the gold standard dataset that on average only 10% wikilinks are suitable for reference inheritance while the rest 90% are irrelevant. To tackle this imbalance in the

⁸We do not use the corresponding deep vectors since the reference title are in many cases very small and do not produce appropriate vectors thus negatively affecting the overall performance of our system.

dataset, we follow the under-sampling technique based on the repeated edited nearest neighbour method proposed by Lemaitre et al. (2017).

First, in order to understand the importance of the proposed features we perform a χ -square test for the classification. We report the ranking of the features in Table 2. Subsequently, we plug the features according to this rank from top to bottom one by one into the classifier; the performance obtained thereby are presented in Table 3. We observe that appending the features one by one according to the χ -square test ranking helps to improve the overall performance of the classifier except the last feature which causes no improvements.

| Rank | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
|----------------|-----|-----|-----|----|------|----|------|----|
| Feature | TIS | OSS | ISS | OS | OSVS | VS | ISVS | IS |

Table 2: χ -square test ranking of features for Step - I (wikilink classification).

| Features Used | Precision | Recall | F-Measure |
|--|------------------|---------------|------------------|
| TIS | 0.11 | 0.90 | 0.20 |
| TIS, OSS | 0.16 | 0.36 | 0.20 |
| TIS, OSS, ISS | 0.39 | 0.37 | 0.33 |
| TIS, OSS, ISS, OS | 0.47 | 0.44 | 0.40 |
| TIS, OSS, ISS, OS, OSVS | 0.48 | 0.42 | 0.40 |
| TIS, OSS, ISS, OS, OSVS, VS | 0.49 | 0.42 | 0.40 |
| TIS, OSS, ISS, OS, OSVS, VS, ISVS | 0.50 | 0.45 | 0.42 |
| TIS, OSS, ISS, OS, OSVS, VS, ISVS, IS | 0.50 | 0.45 | 0.42 |

Table 3: Step - I (wikilink classification) evaluation with incremental feature set.

Some of the representative Wikipedia articles from the CS dataset along with the relevant/irrelevant wikilinks for reference inheritance as correctly predicted by our classifier are noted in Table 4. We observe that the relevant wikilinks are semantically more close to the target Wikipedia articles compared to the irrelevant ones, thus making the former ones as more probable candidates for reference inheritance.

| Target Wikipedia article | Relevant wikilinks | Irrelevant wikilinks |
|---------------------------------|--|-------------------------------------|
| Feature (machine learning) | Machine learning, Perceptron | Computer vision, Speech recognition |
| Vanishing gradient problem | Recurrent neural network, Back-propagation | OPTICS algorithm, Cluster analysis |
| Sensitivity and specificity | True positive rate, Precision and recall | Clinical research, Airport security |

Table 4: Representative results of wikilink classification for the CS dataset.

5.4 Performance analysis of Step - II

Table 5 reports the performance of our approach for various values of k . Considering the difficulty of the task we observe that our system achieves very good performance in terms of precision even for higher values of k . Some of the representative recommended references along with the target Wikipedia articles are shown in Table 6. The correct references in the table signify the references which are actually present in the target Wikipedia article’s reference section (i.e., in the gold standard dataset). We also observe that the incorrect references which are actually not present in the target Wikipedia article’s reference section are not adjudged as relevant by our system. As we cannot compare these references with the gold standard, we perform manual evaluation for only these references in order to analyze how suitable these references are for the target Wikipedia article (see Section 5.6).

| k | Precision | Recall | F-Measure |
|----------|------------------|---------------|------------------|
| 1 | 0.44 | 0.21 | 0.28 |
| 2 | 0.41 | 0.23 | 0.30 |
| 3 | 0.40 | 0.26 | 0.30 |
| 4 | 0.37 | 0.27 | 0.31 |
| 5 | 0.34 | 0.30 | 0.31 |
| 10 | 0.25 | 0.35 | 0.30 |

Table 5: Evaluation of Step - II for different values of k .

| Target Wikipedia article | Relevant wikilink | Correct reference | Incorrect reference |
|-----------------------------|----------------------|---------------------------------|---|
| Heap (data structure) | Binary heap | Introduction to Algorithms | Min-max heaps and generalized priority queues |
| Sensitivity and specificity | Precision and recall | An Introduction to ROC Analysis | Advanced Data Mining Techniques |

Table 6: Representative results of Step - II for the CS dataset.

5.5 WikiRef vs the baselines

The comparisons of the performances of different baselines with WikiRef are noted in Table 7. The table shows that for all the values of k the performance of BL-III and BL-IV is significantly better than BL-I and BL-II respectively leading to the fact that the role of Step - I is very crucial in this context and, in particular, helps boost the performance of the full system. In addition, we observe that the performance we get only using unsupervised ranking in step - II based on any one of the three features (used in BL-III, BL-IV and BL-V respectively) described in section 4.2 cannot beat the performance of our approach where we use a supervised approach like SVMRank that draws advantage of all the three features.

| k | Metric | BL-I | BL-II | BL-III | BL-IV | BL-V | WikiRef |
|----|-----------|-------|-------|--------|-------|------|---------|
| 1 | Precision | 0.22 | 0.15 | 0.38 | 0.31 | 0.36 | 0.44 |
| | Recall | 0.04 | 0.05 | 0.18 | 0.09 | 0.10 | 0.21 |
| | F-Measure | 0.068 | 0.075 | 0.20 | 0.11 | 0.14 | 0.28 |
| 3 | Precision | 0.12 | 0.09 | 0.34 | 0.26 | 0.27 | 0.40 |
| | Recall | 0.08 | 0.1 | 0.22 | 0.17 | 0.16 | 0.26 |
| | F-Measure | 0.096 | 0.075 | 0.28 | 0.17 | 0.20 | 0.30 |
| 5 | Precision | 0.1 | 0.07 | 0.31 | 0.23 | 0.25 | 0.34 |
| | Recall | 0.1 | 0.12 | 0.26 | 0.21 | 0.22 | 0.30 |
| | F-Measure | 0.1 | 0.08 | 0.27 | 0.18 | 0.21 | 0.31 |
| 10 | Precision | 0.06 | 0.06 | 0.25 | 0.22 | 0.23 | 0.25 |
| | Recall | 0.19 | 0.19 | 0.30 | 0.25 | 0.26 | 0.35 |
| | F-Measure | 0.09 | 0.09 | 0.26 | 0.20 | 0.24 | 0.30 |

Table 7: Comparison of the performance of proposed baselines along with WikiRef. Best results are in green cells and the most competing baseline results are in blue cells.

5.6 Manual evaluation

The aim of this manual evaluation is to understand the quality of the recommendations returned by WikiRef for the cases which are not present in the gold standard dataset. In the survey, we provide a total of 25 Wikipedia pages from the Computer Science (CS) dataset along with five recommended references which are not already present in the gold standard and ask the annotators to choose the appropriate references which should be there in the reference section of the given Wikipedia article to make it more resourceful. A sample evaluation page can be seen in this link⁹. The list of scientific Wikipedia articles that we use for manual evaluation are noted in Table 8. 10 participants with Computer Science background have taken part in this survey. First, we compute absolute performance of our system which shows on an average, the fraction of references actually found to be relevant by the annotators. We compute the performance score per Wikipage as the fraction of the proposed references found to be relevant by the annotators (averaged over all the annotators). We then compute the average over the full list of 25 Wikipages, obtaining a very good performance score of 0.63 leading to the fact that out of 5 recommended references, close to 3 references are found to be relevant by the annotators.

Next, we also attempt to correlate the ranking of the competing systems with the implicit ranking imposed by the annotators. For each Wikipedia article, we rank the five references depending on the number of participants voting for a particular reference, using standard fractional ranking method. Then we compute the Spearman’s rank correlation coefficient (ρ) between these ranks and the ranked list returned by WikiRef as well as other baselines. Table 9 shows that WikiRef beats the most competitive baseline by a significant margin.

⁹<https://goo.gl/forms/N0rmN5xPRzhXsoCL2>

Algebraic modeling language, Bayesian model of computational anatomy, Object-based language, Recursive ordinal, Graph isomorphism, Recurrence plot, RossFahroo pseudospectral method, Deterministic finite automaton, Clustal, Graph structure theorem, Bidirected graph, Algebraic graph theory, Belt machine, HadwigerNelson problem, Finite-state machine, Two-phase locking, Urquhart graph, Software transactional memory, Radial basis function kernel, Semi-symmetric graph, Algorithm characterizations, Elliott formula, Just-in-time compilation, Abstract family of languages, PAdES.

Table 8: Target Wikipedia pages for manual evaluation.

| Metric | BL-I | BL-II | BL-III | BL-IV | BL-V | WikiRef |
|----------------|-------|-------|--------|-------|-------|---------|
| Average ρ | 0.099 | 0.16 | 0.089 | 0.168 | 0.104 | 0.203 |

Table 9: Human judgement experiments comparing the proposed baselines with WikiRef. The best result is highlighted in a green cell and the most competing baseline result is highlighted in a blue cell.

5.7 Performance analysis for the Physics (PH) dataset

So far, we have reported all the performance figures for the Computer science (CS) dataset. In order to investigate the applicability of our system in other domains, we repeat the experiments on the Physics (PH) dataset as well. The results for both step - I and step - II are noted in Table 10. We use the same eight features (TIS, OSS, ISS, OS, OSVS, VS, ISVS, IS) with Random Forest classifier for step - I and for step - II we again choose the same SVMRank framework fed in with the three features (F1-TI, F1-Vec and F2-TI) as discussed earlier. From Table 11 we observe that for almost all the cases WikiRef performs better than the baselines.

| Step | | Precision | Recall | F-Measure |
|---------|------|-----------|--------|-----------|
| Step-I | | 0.41 | 0.44 | 0.37 |
| Step-II | k=1 | 0.45 | 0.1 | 0.16 |
| | k=2 | 0.41 | 0.13 | 0.19 |
| | k=3 | 0.38 | 0.16 | 0.18 |
| | k=4 | 0.35 | 0.20 | 0.25 |
| | k=5 | 0.32 | 0.22 | 0.23 |
| | k=10 | 0.30 | 0.25 | 0.26 |

Table 10: Performance of WikiRef on the Physics (PH) dataset

| k | Metric | BL-I | BL-II | BL-III | BL-IV | BL-V | WikiRef |
|----|-----------|------|-------|--------|-------|------|---------|
| 1 | Precision | 0.20 | 0.24 | 0.36 | 0.36 | 0.41 | 0.45 |
| | Recall | 0.11 | 0.07 | 0.09 | 0.09 | 0.09 | 0.1 |
| | F-Measure | 0.14 | 0.09 | 0.12 | 0.12 | 0.13 | 0.16 |
| 3 | Precision | 0.18 | 0.18 | 0.28 | 0.26 | 0.29 | 0.38 |
| | Recall | 0.18 | 0.13 | 0.15 | 0.15 | 0.15 | 0.16 |
| | F-Measure | 0.16 | 0.11 | 0.15 | 0.15 | 0.16 | 0.18 |
| 5 | Precision | 0.15 | 0.15 | 0.24 | 0.24 | 0.25 | 0.32 |
| | Recall | 0.22 | 0.18 | 0.19 | 0.19 | 0.18 | 0.22 |
| | F-Measure | 0.16 | 0.12 | 0.16 | 0.16 | 0.16 | 0.23 |
| 10 | Precision | 0.17 | 0.12 | 0.22 | 0.22 | 0.22 | 0.30 |
| | Recall | 0.29 | 0.27 | 0.23 | 0.23 | 0.22 | 0.25 |
| | F-Measure | 0.15 | 0.13 | 0.17 | 0.18 | 0.17 | 0.26 |

Table 11: Comparison of the performance of proposed baselines along with WikiRef for Physics (PH) dataset. Best results are in green cells and the most competing baseline results are in blue cells.

6 Conclusion

This paper presented a novel two-step recommendation system for enhancing the reference section of the Wikipedia entity pages, dealing with scientific concepts by inheriting the references from the wikilinks present in the page itself. In the first stage we obtain relevant wikilinks for an entity page via a supervised

classification approach. In the second stage we recommend a ranked list of references from the relevant wikilinks obtained from the first stage. WikiRef achieves an overall *precision@1* of 0.44 for the CS dataset and 0.45 for the PH dataset. These are very significant improvements over the most competing baselines in both cases. In addition, manual evaluations over 25 wikipages show that WikiRef also recommends the most relevant references that are absent in the gold standard dataset. We have made our code and data publicly available¹⁰.

Immediate future work would be to focus on further evaluations on various datasets including biology and medicine. Also, the proposed approach can be adapted to filter any type of entity pages' reference section. We further plan to prepare an online tool that triggers WikiRef to recommend references as soon as relevant wikilinks are added to a target Wikipedia article.

References

- [Balog and Ramampiaro2013] Krisztian Balog and Heri Ramampiaro. 2013. Cumulative citation recommendation: Classification vs. ranking. In *Proc. of ACM-SIGIR*. ACM.
- [Balog et al.2013] Krisztian Balog, Heri Ramampiaro, Naimdjon Takhirov, and Kjetil Nørkvåg. 2013. Multi-step classification approaches to cumulative citation recommendation. In *Proc. of OAIR*.
- [Barrena et al.2015] Ander Barrena, Aitor Soroa, and Eneko Agirre. 2015. Combining mention context and hyperlinks from wikipedia for named entity disambiguation. In *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics*, pages 101–105.
- [Bowman et al.2015] Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.
- [Caragea et al.2013] Cornelia Caragea, Adrian Silvescu, Prasenjit Mitra, and C Lee Giles. 2013. Can't see the forest for the trees?: a citation recommendation system. In *Proceedings of the 13th ACM/IEEE-CS joint conference on Digital libraries*, pages 111–114. ACM.
- [Chesney2006] Thomas Chesney. 2006. An empirical examination of wikipedia's credibility. *First Monday*.
- [Conneau et al.2017] Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680. Association for Computational Linguistics.
- [Cucerzan2007] Silviu Cucerzan. 2007. Large-scale named entity disambiguation based on wikipedia data. In *EMNLP-CoNLL*.
- [Dunning1994] Ted Dunning. 1994. *Statistical identification of language*. Computing Research Laboratory, New Mexico State University.
- [Fetahu et al.2015] Besnik Fetahu, Katja Markert, and Avishek Anand. 2015. Automated news suggestions for populating wikipedia entity pages. *CIKM*.
- [Gabrilovich and Markovitch2007] Evgeniy Gabrilovich and Shaul Markovitch. 2007. Computing semantic relatedness using wikipedia-based explicit semantic analysis. In *IJCAI*, volume 7, pages 1606–1611.
- [Gorman2007] GE Gorman. 2007. A tale of information ethics and encyclopaedias; or, is wikipedia just another internet scam? *Online Information Review*.
- [Hachey et al.2013] Ben Hachey, Will Radford, Joel Nothman, Matthew Honnibal, and James R Curran. 2013. Evaluating entity linking with wikipedia. *Artificial intelligence*.
- [Hoffart et al.2011] Johannes Hoffart, Mohamed Amir Yosef, Iliaria Bordino, Hagen Fürstenauf, Manfred Pinkal, Marc Spaniol, Bilyana Taneva, Stefan Thater, and Gerhard Weikum. 2011. Robust disambiguation of named entities in text. In *EMNLP*.
- [Huang et al.2012] Wenyi Huang, Saurabh Kataria, Cornelia Caragea, Prasenjit Mitra, C Lee Giles, and Lior Rokach. 2012. Recommending citations: translating papers into references. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 1910–1914. ACM.

¹⁰<https://github.com/KingOfThePirate/Wikiref>

- [Huang et al.2014] Wenyi Huang, Zhaohui Wu, Prasenjit Mitra, and C Lee Giles. 2014. Refseer: A citation recommendation system. In *Proceedings of the 14th ACM/IEEE-CS Joint Conference on Digital Libraries*, pages 371–374. IEEE Press.
- [Huang et al.2015] Wenyi Huang, Zhaohui Wu, Liang Chen, Prasenjit Mitra, and C Lee Giles. 2015. A neural probabilistic model for context based citation recommendation. In *AAAI*, pages 2404–2410.
- [İkikat et al.2015] F Yesjm İkikat, Behire Gürhan, and Banu Dırı. 2015. Automatic linking of wikipedia pages by their semantic similarity. In *Innovations in Intelligent SysTems and Applications (INISTA), 2015 International Symposium on*, pages 1–5. IEEE.
- [Joachims2006] Thorsten Joachims. 2006. Training linear svms in linear time. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 217–226. ACM.
- [Lemaître et al.2017] Guillaume Lemaître, Fernando Nogueira, and Christos K. Aridas. 2017. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *Journal of Machine Learning Research*, 18(17):1–5.
- [Manning and Schütze1999] Christopher D Manning and Hinrich Schütze. 1999. *Foundations of statistical natural language processing*. MIT press.
- [Meng et al.2013] Fanqi Meng, Dehong Gao, Wenjie Li, Xu Sun, and Yuexian Hou. 2013. A unified graph model for personalized query-oriented reference paper recommendation. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 1509–1512. ACM.
- [Moro et al.2014] Andrea Moro, Alessandro Raganato, and Roberto Navigli. 2014. Entity linking meets word sense disambiguation: a unified approach. *TACL*.
- [Orlowski2005] Andrew Orlowski. 2005. Wikipedia science 31% more cronky than britannica’s. *The Register*.
- [Pochampally et al.2016] Yashaswi Pochampally, Kamalakar Karlapalem, and Navya Yarrabelly. 2016. Semi-supervised automatic generation of wikipedia articles for named entities. In *Wiki@ICWSM*.
- [Raganato et al.2016] Alessandro Raganato, Claudio Delli Bovi, and Roberto Navigli. 2016. Automatic construction and evaluation of a large semantically enriched wikipedia. In *IJCAI*, pages 2894–2900.
- [Ren et al.2014] Xiang Ren, Jialu Liu, Xiao Yu, Urvashi Khandelwal, Quanquan Gu, Lidan Wang, and Jiawei Han. 2014. Cluscite: Effective citation recommendation by information network-based clustering. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 821–830. ACM.
- [Sauper and Barzilay2009] Christina Sauper and Regina Barzilay. 2009. Automatically generating wikipedia articles: A structure-aware approach. In *Proc. of ACL-IJCNLP*.
- [Taneva and Weikum2013] Bilyana Taneva and Gerhard Weikum. 2013. Gem-based entity-knowledge maintenance. In *Proc. of ACM-CIKM*, pages 149–158.
- [Tang et al.2014] Xuewei Tang, Xiaojun Wan, and Xun Zhang. 2014. Cross-language context-aware citation recommendation in scientific articles. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval*, pages 817–826. ACM.
- [West et al.2014] Robert West, Evgeniy Gabrilovich, Kevin Murphy, Shaohua Sun, Rahul Gupta, and Dekang Lin. 2014. Knowledge base completion via search-based question answering. In *WWW*.
- [Wu and Giles2015] Zhaohui Wu and C Lee Giles. 2015. Sense-aware semantic analysis: A multi-prototype word representation model using wikipedia. In *AAAI*, pages 2188–2194.
- [Wu and Weld2010] Fei Wu and Daniel S Weld. 2010. Open information extraction using wikipedia. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 118–127. Association for Computational Linguistics.