Fact Checking in Heterogeneous Information Networks

Baoxu Shi Tim Weninger
Department of Computer Science and Engineering
University of Notre Dame
{bshi, tweninge}@nd.edu

ABSTRACT

Traditional fact checking by experts and analysts cannot keep pace with the volume of newly created information. It is important and necessary, therefore, to enhance our ability to computationally determine whether some statement of fact is true or false. e view this problem as a link-prediction task in a knowledge graph, and show that a new model of the top *discriminative meta paths* is able to understand the meaning of some statement and accurately determine its veracity. e evaluate our approach by examining thousands of claims related to history, geography, biology, and politics using public, million node knowledge graphs extracted from ikipedia and SemMedDB. Not only does our approach significantly outperform related models, we also find that the discriminative path model is easily interpretable and provides sensible reasons for the final determination.

1. INTRODUCTION

Misinformation in media and communication creates a situation in which opposing assertions of fact compete for attention. Although much of the information presented on the eb is a good resource, its accuracy cannot be guaranteed. In order to avoid being fooled by false assertions, it is necessary to separate fact from fiction and to assess the credibility of an information source.

ith that goal in mind, we present a method for fact checking in knowledge graphs. Given a statement $\mathcal{S}$ in format (subject, predicate, object), (Chicago, capitalOf, Illinois) for example, our approach mines discriminative paths that alternatively define the generalized statement (US city, capitalOf, US states) and uses the mined rules to evaluate the veracity of statement $\mathcal{S}$.

Unlike existing models which rely on connectivity, specificity, or human annotated relations, the proposed method simulates how experienced human fact-checkers examine a statement; namely fact-checkers will first attempt to *understand* the generalized version of a given type of statements using prior knowledge, and then validate the specific statement by applying their understandings to it. Returning to the "*Chicago is the capital of Illinois*" example, a fact checker, as well as our model, will learn alternative definitions of capitalOf, such as "*a US city is likely to be the capital of a US state if an state agency with jurisdiction in that state has its headquarters in that city*", from its knowledge base and validate the statement based on these definitions.

To summarize, we show that we can leverage a collection of factual statements for automatic fact checking. Based on the principles underlying link prediction, similarity search and network closure, we computationally gauge the truthfulness of an assertion by mining heterogeneous connectivity patterns within a network of factual statements, *e.g.*, ikipedia and SemMedDB. Our current work focuses on determining the validity of factual assertions from simple, well-formed RDF (subject, predicate, object) statements; the related problems of information extraction, claim identification, and compound assertion answering are generally built on top of this central task.

2. METHODOLOGY

e define the fact checking problem as a link prediction task in typed knowledge graph $\mathcal{G}$. Checking statement $\mathcal{S} = (\mathsf{s}, \mathsf{p}, \mathsf{o})$ is equivalent to predicting the existence of edge $\mathsf{s} \xrightarrow{\mathsf{p}} \mathsf{o}$ in $\mathcal{G}$. e perform this task by mining alternative paths $\mathbf{P}$ between entities of the same type as s and o through the knowledge graph $\mathcal{G}$.

In order to mine the alternative definition $\mathbf{P}$, we need to determine the search space containing the statements that are similar to $\mathcal{S}$. Unlike association rule mining which discards the type information of the entities [3] or topology-based approaches that look only at distances between the subject and the object [1, 2, 5, 4], in this work we use both entity type and predicate type information to test the validity of the proposed statement.

e first find the entity types ψ of s and o, then construct positive node-pair set $\mathbf{T}^+$ and negative node-pair set $\mathbf{T}^-$, where $\mathbf{T}^+ = \{(u,v)|u \xrightarrow{\mathsf{p}} v \in \mathcal{G}\}$, $\mathbf{T}^- = \{(u,v)|u \xrightarrow{\mathsf{p}} v \notin \mathcal{G}\}$, $\psi(u) = \psi(\mathsf{s})$, and $\psi(v) = \psi(\mathsf{o})$. Because knowledge graphs are typically incomplete, *i.e.*, all possible truthful facts are not listed, $\mathbf{T}^+$ and $\mathbf{T}^-$ can also be generated externally if $\psi(\mathsf{s})$, $\psi(\mathsf{o})$, or p is missing in G. After we compute $\mathbf{T}^+$ and $\mathbf{T}^-$, we perform a DFS-like graph traversal algorithm to retrieve meta path sets $\mathbf{\Pi}^+$ and $\mathbf{\Pi}^-$ corresponding to the positive and negative end points.

For example, given the statement of fact (Chicago, capitalOf, Illinois), which means "*Chicago is the capital of Illinois*,", we find that Chicago is a city-entity, and Illinois is a state-entity. The positive node-pair $\mathbf{T}^+$ contains other meta paths that match city⇝state like, Sacramento⇝California and Albany⇝New ork, etc., which are connected by capitalOf, as well as negative meta paths like Los Angeles⇝California and New ork City⇝New ork, which are not capital-state pairs.

This gives us a large set of positive and negative meta paths between the entities of the same type as the subject and entities of the same type as the object. e then extract the top-k meta paths $\mathbf{P}$

WWW'16 Companion, April 11–15, 2016, Montréal, Québec, Canada.
ACM 978-1-4503-4144-8/16/04.
http://dx.doi.org/10.1145/2872518.2889354.

Table 1: **Result of fact checking tasks. The score is the area under ROC curve score computed by logistic regression with 10-fold cross validation.**

Algorithm	CapitalOf	Company CEO	N T Bestseller	US Civil ar	US Vice-President	Disease	Cell
Adamic/Adar	0.387	0.665	0.650	0.642	0.795	0.671	0.755
Semantic Proximity	0.706	0.614	0.641	0.582	0.805	0.871	0.840
Preferential Attachment	0.396	0.498	0.526	0.599	0.474	0.563	0.755
Katz	0.370	0.600	0.623	0.585	0.791	0.763	0.832
SimRank	0.553	**0.824**	**0.695**	0.685	0.912	0.809	0.749
AMIE	0.550	0.669	0.520	0.659	0.987	0.889	0.898
Personalized PageRank	0.535	0.579	0.529	0.488	0.683	0.827	0.885
Path-Constrained Random alk	0.550	0.542	0.486	0.488	0.672	0.911	0.765
Discriminative Path Count	**0.920**	0.747	0.664	**0.749**	**0.993**	**0.941**	**0.928**

Table 2: Discriminative path mined by the proposed method. † means missed by AMIE.

Task	Entity Type	Predicate Path	Entity Type
US Vice-President	{Person}	<president^{-1}> <successor> <successor, president^{-1}>†	{Person}
CaoitalOf	{City}	<headquarter^{-1}, jurisdiction^{-1}>† <location^{-1}, jurisdiction^{-1}>† <isPartOf>†	{State}

from $\mathbf{\Pi}^{+} \cup \mathbf{\Pi}^{-}$ by calculating the information gain of each meta path. This effectively prunes the irrelevant meta paths that exist in both positive and negative path set, and results in a training set with positive and negative subjects and objects as training instances and the top-k meta paths as the features where the cell values are the number of paths in $\mathcal{G}$ that separate the (subject, object) pair via the meta path. e use logistic regression to train the model, but any standard numeric classifier will work.

Most interestingly, the top-k paths found by this method **provide a human interpretable, intuitive explanation about the meaning of the fact being analyzed**, and can usually describe why the stated fact is true or false by investigating the regression variables.

3. EXPERIMENTS

To test the performance and the generality of the proposed model, we evaluate our model on two large knowledge bases, DBpedia and SemMedDB, with seven different fact checking tasks related to history, geography, biology and politics.

To test the ability of our method to validate missing facts, we remove all edges labelled by the given predicate p from the given statement (s, p, o). e also set the true/false label ratio to 20/80 to simulate real-world fact checking scenarios where the proportion of potential false statements are significantly larger than true statements. All experiments are performed using 10-fold cross validation. The source code is available at `http://github.com/nddsg/KGMiner`.

In Tab. 1 we compare the proposed fact-checking algorithm with six link prediction models, one fact checking model (Semantic Proximity) [2], and an association rule mining model (AMIE) [3].

The proposed method outperforms other methods in most tasks. For certain tasks where the path variation is low and entity connectivity is high, *e.g.*, book authors and company CEOs, there are relatively few alternate paths that are suitable defining the given statement. As a result the connectivity based methods will have a more competitive performance.

Table 2 shows the discriminative paths we mined for CapitalOf and US Vice-President. The proposed method successfully discovers discriminative paths that alternatively define the given statement. One interesting finding is the predicate path <successor, president^{-1}>, which reveals that the US constitution allows for the possibility to replace one vice president with another – a little known, yet valid part of the definition of US vice-president, and makes our model perform almost perfectly for that task.

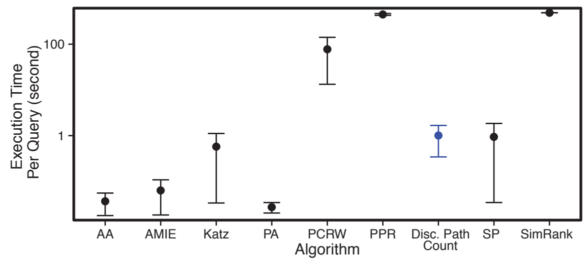

Figure 1: Time consumption of each algorithm. Point represents the average feature generation time of one query. Error bars represent 95% confidence interval over the seven tasks. Lower is better. The execution time of AMIE does not include the time ($\approx 4,000$ hours in total) spent on rule mining.

The efficiency of the methods shown in Tab. 1 are compared in Fig. 1, where the y-axis is on a log scale. e find that our method is comparable to the simple topology methods, and can analyze knowledge base facts in about a second on average.

4. CONCLUSIONS

e presented a fact checking framework for knowledge graphs that discovers the definition of a given statement of fact by extracting discriminative paths from the knowledge graph, and uses the model to validate the truthfulness of the given statement.

To evaluate the proposed method, we checked the veracity of several thousand statements across seven different tasks on DBpedia and SemMedDB. e found that our framework was the all around best in terms of fact-checking performance and has a running time similar to existing models. Finally, we showed that the proposed framework can discover interpretable and informative discriminative paths that are missed by other methods.

5. ACKNOWLEDGEMENTS

This work is sponsored by an AFOSR grant FA9550-15-1-0003, and a John Templeton Foundation grant FP053369-M.

6. REFERENCES

[1] L. A. Adamic and E. Adar. Friends and neighbors on the eb. *Social Networks*, 25(3), 2003.

[2] G. L. Ciampaglia, P. Shiralkar, L. M. Rocha, J. Bollen, F. Menczer, and A. Flammini. Computational fact checking from knowledge networks. *PLoS ONE*, 10(6), 2015.

[3] L. A. Galárraga, C. Teflioudi, K. Hose, and F. Suchanek. AMIE: association rule mining under incomplete evidence in ontological knowledge bases. In *WWW*, 2013.

[4] T. H. Haveliwala. Topic-sensitive pagerank. In *WWW*, 2002.

[5] L. Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1), 1953.