Temporal Evolution of Entity Relatedness using Wikipedia and DBpedia

Narumol Prangnawarat and Conor Hayes

Insight Centre for Data Analytics, National University of Ireland, Galway
{narumol.prangnawarat,conor.hayes}@insight-centre.org

Abstract. Entity relatedness is a task that is required in many applications such as entity disambiguation and clustering. Although there are many works on entity relatedness, most of them do not consider temporal aspects and thus, it is unclear how entity relatedness changes over time. This paper attempts to address this gap by showing how entity relatedness develops over time with graph-based approaches using Wikipedia and DBpedia as well as how transient links and stable links, by which is meant links that do not persist over different times and links that persist over time respectively, affect the relatedness of the entities. We show that different versions of the knowledge base at different times give different accuracy on the KORE dataset, and show how using multiple versions of the knowledge base through time can increase the accuracy of entity relatedness over a version of the knowledge graph from a single time point.

1 Introduction

Many researches have made use of Wikipedia and DBpedia for relatedness and similarity tasks. However, most approaches use only the current information but not temporal information. For example, the entities *Mobile* and *Camera* may have been less related in the past but may be more related at present. This work shows how semantic relatedness develops over time using graph-based approaches over the Wikipedia network as well as how transient links, links that do not persist over different times, and stable links, links that persist over time, affect the relatedness of the entities.

We hypothesise that using graph-based approaches on the Wikipedia network provides higher accuracy in term of relatedness score to the ground truth data in comparison to text-based approaches. We take each Wikipedia article as an entity, for example, the article on *Semantic similarity*[1] corresponds to the entity *Semantic_similarity*. Although the term *article* and *entity* may be referred to interchangeably in this work, *article* refers to Wikipedia article and *entity* refers to a single concept in the semantic network. Wikipedia users provide Wikipedia page links within articles, so we can make use of the provided links as entities. We assume that entities which are closely related to an entity, a Wikipedia article in

---

[1] https://en.wikipedia.org/wiki/Semantic_similarity

this case, are mentioned in that Wikipedia article. Hence, closely related entities share the same adjacent nodes in the Wikipedia article link network. We make use of the temporal Wikipedia article link network to demonstrate the evolution of entity relatedness and how transient or stable links affect the relatedness of the entities. We analyse different models of aggregated graphs, which integrate networks at various times into the same graph as well as time-varying graphs which are the series of the networks at each time step.

We first show that the proposed graph-based approach outperforms the text-based approach in terms of the accuracy of relatedness score. Then, we present our proposed similarity method, which outperforms the baseline graph-based similarity methods in term of relatedness score accuracy over various different network models. We also show the evolution of relatedness as well as how transient links and stable links effect the relatedness using graph-based similarity.

2   Related Works

2.1   Semantic Relatedness

Semantic relatedness works have been carried out for words (natural language texts such as common nouns and verbs) and entities (concepts in semantic networks such as companies, people and places). The approaches used for semantic relatedness includes corpus-based or text-based approaches as well as structure-based or graph-based approaches. Wikipedia and DBpedia have been widely used as resources to find semantic relatedness. However, most approaches use only the current information without temporal aspects. One of the well known system is WikiRelated [13] introduced by Strube and Ponzetto. WikiRelated use the structure of Wikipedia links and categories to compute the relatedness between concepts.

Gabrilovich and Markovitch proposed Explicit Semantic Analysis (ESA) [5] to compute semantic relatedness of natural language texts using high-dimensional vectors of concepts derived from Wikipedia. DiSER [1], presented by Aggarwal and Buitelaar, improves ESA by using annotated entities in Wikipedia.

Leal et al. [9] proposed a novel approach for computing semantic relatedness as a measure of proximity using paths on DBpedia graph. Hulpus et al. [8] provided a path-based semantic relatedness using DBpedia and Freebase as well as presenting the use in word and entity disambiguation.

Radinsky et al. proposed a new semantic relatedness model, Temporal Semantic Analysis (TSA) [12]. TSA computes relatedness between concepts by analysing the time series between words and find the correlation over time. Although this work make use of the temporal aspect to find relatedness between words, it does not show how the relatedness evolve over time.

2.2   Time-aware Wikipedia

There have been a number of works in the area of time-aware Wikipedia analysis. The authors have primarily focused on content and statistical analysis.

WikiChanges [10] presents a Wikipedia article's revision timeline in real time as a web application. The application presents the number of article edits over time in day and month granularity. The authors also provide the extension script for embedding a revision activity to Wikipedia. Ceroni et. al. [4] introduced a temporal aspect for capturing entity evolution in Wikipedia. They provided time-based framework for temporal information retrieval in Wikipedia as well as statistical analysis such as the number of daily edits, Wikipedia pages having edits and top Wikipedia pages that has been changed over time. Whiting et. al. [15] presented Wikipedia temporal characteristics, such as topic coverage, time expressions, temporal links, page edit frequency and page views, to find how the knowledge can be exploited in time-aware research. However, in-degree and out-degree are the only networks properties that are discussed in the paper.

Recent research from Bairi et. al. [2] presents statistics of categories and articles, such as the number of articles, the number of links and the number of categories, comparing between the Wikipedia instance in October 2012 and the Wikipedia instance in June 2014. The authors also analysed the Wikipedia category hierarchy as a graph and provided statistics of the category graph such as number of cycles and and the cycle length of the two Wikipedia instances. Contropedia [3] identifies when and which topics have been most controversial in Wikipedia article using Wikipedia links as the representation of the topics. However, the approach also focus on the content changed in the article.

Our work make use of the changes in Wikipedia to analyse evolution of entity relatedness. We show how entity relatedness develops over time using graph based approaches over Wikipedia network as well as how transient links and stable links affect the relatedness of the entities.

3 Dataset

The Wikipedia data can be downloaded from the Wikimedia download page[2], where dumps are extracted twice a month. A major limitation of data availability from Wikipedia dumps is that the oldest available Wikipedia dump at the time of our experiment was from 20 August 2016.

We treat each Wikipedia article as an entity. Each Wikipedia article has user provided Wikipedia article links which link to the corresponding Wikipedia articles. We extract Wikipedia links within Wikipedia articles from each Wikipedia dump. Figure 1 shows an example of a part of Wikipedia article link network of the *DBpedia* article[3]. The part of *DBpedia* article contains links to *Database*, *Structured_content*, *Wikipedia*, *World_Wide_Web*, *Semantic_Query*, *Tim_Berners-Lee*, *Dataset*, and *Linked_Data*. The article links of all articles in Wikipedia construct the full Wikipedia article link network.

As the oldest Wikipedia dumps available for download on Wikimedia Downloads page is Wikipedia dump on 20 August 2016 at the time we conduct this

---

[2] https://dumps.wikimedia.org/
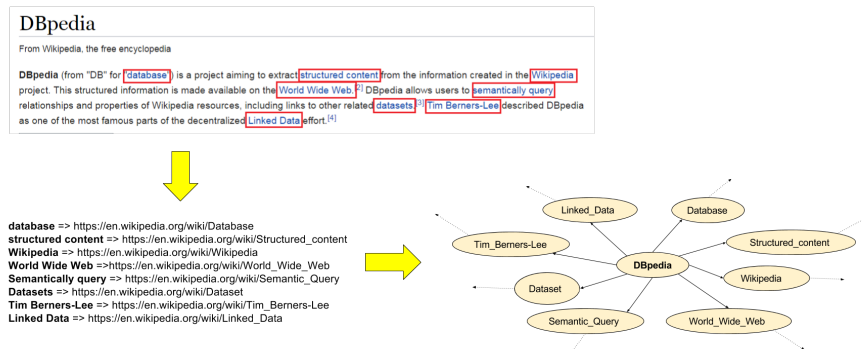[3] https://en.wikipedia.org/wiki/DBpedia

Fig. 1: An example of a part of Wikipedia article link network or the *DBpedia* article

experiment, in order to get older Wikipedia link data, we obtained data from DBpedia[4] which is an open knowledge base extracted from Wikipedia and other Wikimedia projects. We use the page links datasets which are the relationship of article links within each article, the same information as the Wikipedia links we extracted from Wikipedia. Each DBpedia concept corresponds to a Wikipedia article. For example, the DBpedia concept `http://dbpedia.org/resource/Semantic_similarity` corresponds to the article *Semantic similarity*[5]. We refer to both the *Semantic similarity* article and the DBpedia concept `http://dbpedia.org/resource/Semantic_similarity` as the entity *Semantic_similarity*. We make use of page links dataset from DBpedia to construct a series of Wikipedia article link networks for each year from 2007 to 2016. The DBpedia versions used are shown in Table 1.

4 Methodology

We take each Wikipedia article as an entity. We assume that entities which are closely related to an entity, a Wikipedia article in this case, are mentioned in that Wikipedia article. Hence, closely related entities share the same adjacent nodes in the Wikipedia article link network. However, there might be some articles that link to a lot of the articles that might not be semantically related. For example, the main page, which is the landing page for featured articles and news, is not semantically related to the articles it links to. On the other hand, articles that do not have many links to other pages might have more semantically relation to their links. Because of this reason, we apply weights to the relationships to penalise the articles that link to many other unrelated articles.

---

[4] `http://wiki.dbpedia.org`

[5] `https://en.wikipedia.org/wiki/Semantic_similarity`

Table 1: DBpedia versions

Dataset	DBpedia version	Wikipedia Dumps
DBpedia 2016	DBpedia 2016-04	Mar-2016/Apr-2016
DBpedia 2015	DBpedia 2015-04	Feb-2015/Mar-2015
DBpedia 2014	DBpedia 2014	02-May-2014
DBpedia 2013	DBpedia 3.9	03-Apr-2013
DBpedia 2012	DBpedia 3.8	01-Jun-2012
DBpedia 2011	DBpedia 3.7	22-Jul-2011
DBpedia 2010	DBpedia 3.6	11-Oct-2010
DBpedia 2009	DBpedia 3.4	20-Sep-2009
DBpedia 2008	DBpedia 3.2	08-Oct-2008
DBpedia 2007	DBpedia 3.0rc	23-Oct-2007

First, we introduce different models we used to represent Wikipedia article link networks. Then, we explain our approach we used to find relatedness over the proposed models.

4.1 Models

To reflect how relatedness changes over time, we represent temporal information from Wikipedia article link networks in two different ways. One is as time-varying graphs which are a series of Wikipedia article link snapshots at each time step. We use each version of datasets to construct a network as each snapshot. Another is as an aggregated graph, which aggregates all networks at each time step together with time information as weights.

Time-Varying Graphs Given an article a for a corresponding entity, the series of networks G_S^a at the set of time $T = \{1, .., n\}$ is constructed as a set of time step graphs $\{G_1^a, G_2^a, ..., G_n^a\}$. Each graph $G_t^a = (V_t^a, E_t^a)$ represents a snapshot of the 2-hop ego network of the article links around an entity a at the time t, where V_t^a is a set of nodes where each node represents a Wikipedia entities that have links with a or have links with the nodes that are adjacent to a at time t and E_t^a is a set of edges where each edge e_{ijt} is an internal link between Wikipedia entities i and j at time t. In other words, $v \in V_t^a$ if at time $t, e_{avt} \in E_t^a$ or $e_{abt} \in E_t^a \wedge e_{bvt} \in E_t^a$. Figure 2 demonstrate the example of a 2-hop ego network around the entity a at time t. We construct a series of 2-hop ego networks of Wikipedia article links over time around each seed entity that we are interested in.

Aggregated Graphs We create different models of aggregated graphs as follows.

Intersection model
Given a Wikipedia article a for a corresponding entity, $G_I^a = (V_I^a, E_I^a)$, V_I^a is a
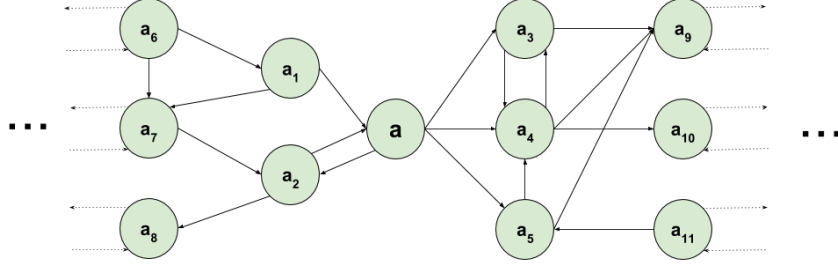
Fig. 2: An example of a 2-hop ego networks around the entity a at time t

set of nodes where each node represents a Wikipedia article that have links with a or have links with the nodes that are adjacent to a at all time points and E_I^a is a set of edges where each edge e_{ij} is an internal link between Wikipedia entities i and j which appear at all time. In other words, $v \in V_I^a$ if $v \in V_1^a \cap V_2^a \cap ... \cap V_n^a$ and $e_{ij} \in E_I^a$ if $e \in E_1^a \cap E_2^a \cap ... \cap E_n^a$ for $[1..n] \in T$.

Union model
Given a Wikipedia article a for a corresponding entity, $G_U^a = (V_U^a, E_U^a)$, V_U^a a set of nodes where each node represents a Wikipedia article that have links with a or have links with the nodes that are adjacent to a at any time in T and E_U^a is a set of edges where each edge e_{ij} is an internal link between Wikipedia entities i and j. In other words, $v \in V_U^a$ if $v \in V_1^a \cup V_2^a \cup ... \cup V_n^a$ and $e_{ij} \in E_U^a$ if $e_{ij} \in E_1^a \cup E_2^a \cup ... \cup E_n^a$ for $[1..n] \in T$.

4.2 Proposed Extended Jaccard Similarity

Jaccard similarity coefficient measures similarity between two objects using binary attributes. Given object a and b with a vector of features A and B respectively, Jaccard similarity coefficient of a and b, $J(a, b)$ is computed as the following equation.

$$J(a, b) = \frac{|A \cap B|}{|A \cup B|}$$

Taking adjacent nodes of an entity a in the Wikipedia article link network as the features of the entity a, Jaccard similarity coefficient can reflect our assumption that closely related entities share the same adjacent links in Wikipedia article links network. However, Jaccard similarity coefficient cannot take into account non-binary features. As we discussed before, there might be some article that links to a lot of the pages that might not be semantically related so we want to apply weights to the relationships to penalise pages that link to many

unrelated pages. PageRank [11] is used to rank the importance of nodes. It was originally created to rank web pages in World Wide Web network in Google search engine. We use the same idea to apply to our Wikipedia article link network. An article that mentions a lot of other articles may not be semantically related to the articles that they link to. On the other hand, an article that is mentioned in a lot of articles may just be a general article that is not semantically related to them. We make use of Tanimoto similarity [14] to extend Jaccard similarity using reciprocal PageRank, which is 1 divided by PageRank score, as weights to find similarity between each entity in the network. The underlying assumption is that articles with lower PageRank scores might have more semantically relation to their links as they only link to fewer articles that are really related to them. Given A is a vector of reciprocal PageRank score of the articles having links with an entity a in the 2-hop ego network around an entity a and B is a vector of reciprocal PageRank score of the articles having links with an entity b in the 2-hop ego network around an entity b, the relatedness between two entities a and b is computed as:

$$R(a,b) = \frac{A \cdot B}{|A|^2 + |B|^2 - A \cdot B}$$

We apply the extended Jaccard similarity with reciprocal PageRank to our models, time-varying graphs and aggregated graphs.

5 Experiment Results

We conducted experiments and evaluated with the KORE [6] dataset. The KORE dataset has been created to measure relatedness between named entities. It consists of 420 related entity pairs from a selected set of 21 seed entities from the YAGO2 [7] knowledge base from 4 different domains, which are 5 entities from IT companies, 5 entities from Hollywood celebrities, 5 entities from video games, 5 entities from television series, and one singleton entity. Each of the entities has 20 ranked related entities. All entities in the KORE dataset corresponds to entities in our Wikipedia article link networks. We use Spearman Correlation to compare the relatedness scores from each approach with the scores from the KORE dataset. As the KORE dataset provides only the ranking but not the score, we assume that the highest entity has a score of 20 and each subsequent entity has a score 1 lower.

5.1 Proposed Extended Jaccard Similarity results

For each DBpedia version stated above, we constructed the series of networks of entities seeding the entities from the KORE dataset as described in Section 4.1.

The experiments were conducted to compare two different perspectives. In the first evaluation perspective, we performed experiments to show that the proposed graph-based approach outperforms the text-based approach in terms of the relatedness scores. We used Term Frequency-Inverse Document Frequency

(TF-IDF) based similarity as the text-base baseline to compared to the proposed extended Jaccard similarity. In the second evaluation perspective, we performed experiments to show that our proposed link-based extended Jaccard similarity with Reciprocal PageRank outperforms the baseline approaches in terms of the relatedness score accuracy. We used Jaccard similarity as the baseline for the graph-based approach to compare to our extended Jaccard similarity. We analysed variations of features for Jaccard methods by considering only direct predecessors of the nodes which are the entities that have links to the nodes, only direct successors of the nodes which are the entities that have links from the nodes, and both direct predecessors and direct successors of the nodes which are entities that have links to or from the nodes.

We performed TF-IDF based similarity on three different Wikipedia text revisions. One is the revision at the time when YAGO2 was created (17-Aug-2010) which is used to constructed the KORE dataset. The second one is the revision at the dump time of DBpedia 2009 dataset (20-Sep-2009) and the last one is the revision at the dump time of DBpedia 2010 dataset (11-Oct-2010). We performed Spearman Correlation to evaluate with the gold standard dataset, the KORE dataset, as described previously. The Spearman Correlations of the 3 different datasets compared to the KORE gold standard are shown in Table 2. We can see that the result of the data acquired at the time when YAGO2 was created has the highest correlation as the same information is captured at that time.

Table 2: Spearman Correlations to the KORE gold standard with the text-based approach for different Wikipedia dumps

Version	Correlation
Wikipedia at the time when YAGO2 was created	**0.503212**
Wikipedia at Dbpedia 2009 version dump time	0.491440
Wikipedia at Dbpedia 2010 version dump time	0.502987

We compared the text-based approach over each version of dataset to the graph-based approaches. In this section, we focused on DBpedia 2009 dataset and DBpedia 2010 dataset as they are the closest snapshots to the Wikipedia dump from 2010-08-17 which is used to constructed YAGO2 using by the KORE dataset. We found that the graph-based approaches outperform the result from the text-based approach in term of accuracy of relatedness score as shown in Table 3. The Spearman Correlation to the KORE dataset from our Extended Jaccard with Reciprocal PageRank is statistically significantly better than TF-IDF based similarity (p-value < 0.05) on both datasets. Moreover, the results show that our proposed extended Jaccard with reciprocal PageRank gives a better accuracy of relatedness score than the baseline Jaccard methods. The Spearman Correlation to the KORE dataset from our Extended Jaccard with Reciprocal

PageRank is statistically significantly better than the baseline Jaccard methods (p-value < 0.05) on DBpedia 2009.

Table 3: Spearman Correlations to the KORE gold standard comparing text-based approach with graph-based approaches of data from 2009 and 2010

Method/Dataset	DBpedia 2009	DBpedia 2010
TF-IDF	0.491440	0.502987
Jaccard (direct predecessors)	0.568509	0.568026
Jaccard (direct successors)	0.511564	0.559898
Jaccard (both direct predecessors and direct successors)	0.578706	0.585535
Extended Jaccard with Reciprocal PageRank	**0.604112**	**0.599284**

5.2 Evolution of Relatedness

We analysed the change of correlations to the KORE gold standard in different datasets to demonstrate how relatedness progress over subsequent years. We found that the dataset from the network in 2009 and 2010 got the highest results. This is because the KORE dataset is constructed using YAGO2 that use the Wikipedia dump from 2010-08-17 [7], which is in between DBpedia version 2009 (2009-09-20) and 2010 (2010-10-11). Figure 3 shows the comparison of Spearman Correlation result of different methods from each dataset.
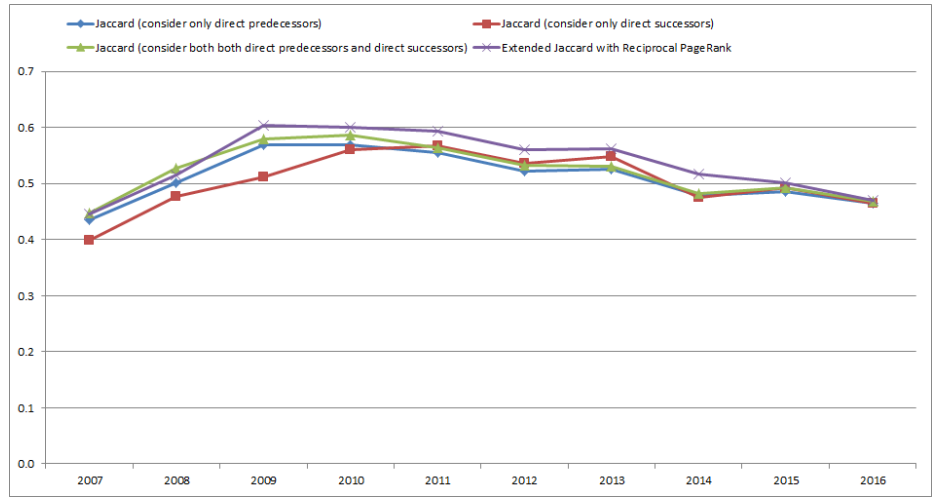

Fig. 3: The comparison of Spearman Correlations to the KORE gold standard of different graph-based methods from each DBpedia dataset
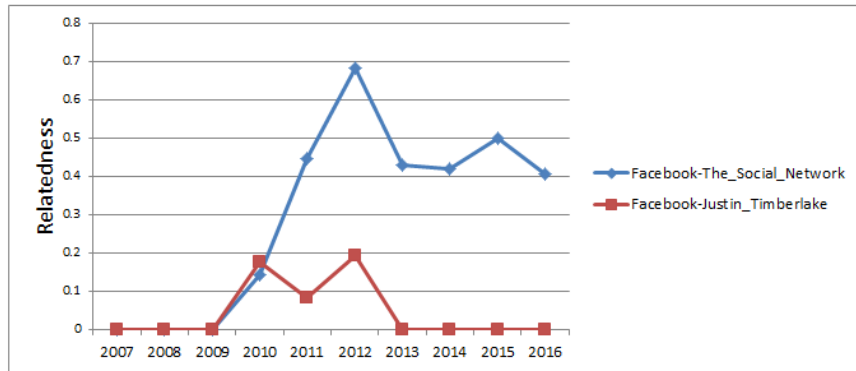
Fig. 4: Evolution of relatedness between *Facebook* and *The_Social_Network* in comparison to *Facebook* and *Justin_Timberlake*

We can see from the result that the most updated knowledge bases will not give the highest correlation if the ground truth data is created in different time. This is because the relatedness score varies according to the relatedness of the entities at that time. For instance, *Facebook* and *Justin_Timberlake* has a higher relatedness score in 2010 as *Justin_Timberlake* stared in *The_Social_Network* movie, the film portrays the founding of *Facebook*, and the score faded after that as shown in Figure 4.
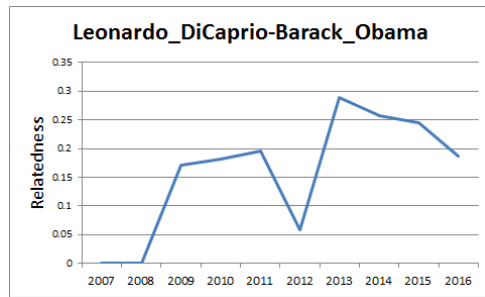

Fig. 5: Evolution of relatedness between *Leonardo_DiCaprio* and *Barack_Obama*

As shown in Figure 5, *Leonardo_DiCaprio* became related to *Barack_Obama* after 2008 as he supported Barack Obama's presidential campaign in the 2008 election and became more related again after 2012 as his support to Obama's

2012 campaign[6,7]. As the events occurred the end of the years, the changes in relatedness are not shown in these versions (2008 and 2012) but appear in the next versions (2009 and 2013) instead.

We also did a qualitative analysis for the entities which are not in the KORE dataset as they become more related to the seed entities after the dataset was created. For instance, *Tim_Cook* started to have high relatedness with *Apple_Inc.* since 2011 when he become the CEO of the company. Figure 6 shows the relatedness of *Apple_Inc.* and *Tim_Cook* in comparison to *Apple_Inc.* and *Steve_Jobs*.
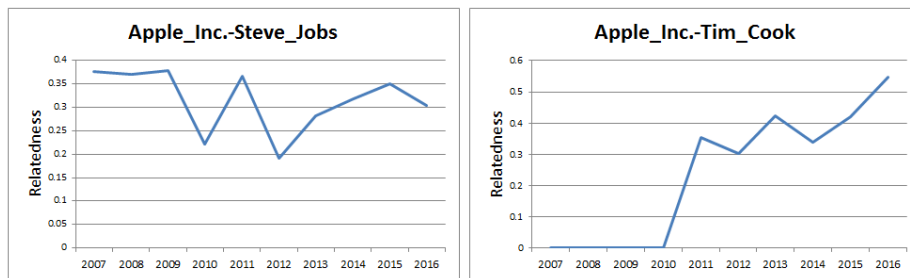

Fig. 6: Evolution of relatedness between *Apple_Inc.* and *Tim_Cook* in comparison to *Apple_Inc.* and *Steve_Jobs*
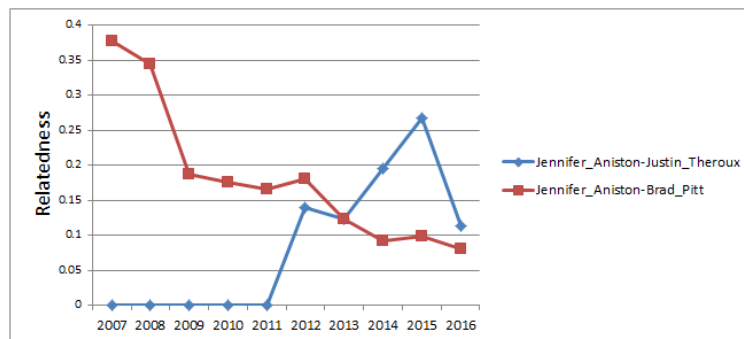

Fig. 7: Evolution of relatedness between *Jennifer_Aniston* and *Justin_Theroux* in comparison to *Jennifer_Aniston* and *Brad_Pitt*

---

[6] https://en.wikipedia.org/wiki/Leonardo_DiCaprio

[7] http://www.hollywoodreporter.com/news/julianne-moore-leonardo-dicaprio-obama-video-375796

Another example is the relatedness between *Jennifer_Aniston* and *Justin_-Theroux*[8] which became higher from 2011 when they started dating[9]. While the relatedness between *Jennifer_Aniston* and *Brad_Pitt* stayed high for a while after they divorced in 2005[10] and then faded as shown in Figure 7.

We performed the Spearman Correlation between different datasets and found that the correlations of the entity relatedness are higher to the dataset versions that are closer in time to themselves and lower when the time is more different as shown in Table 4. This shows that the entity relatedness gradually change over time.

Table 4: Correlations of the entity relatedness between different DBpedia datasets

Dataset	2008	2009	2010	2011	2012	2013	2014	2015	2016
2007	0.7141	0.6054	0.4542	0.4762	0.4356	0.4430	0.3856	0.3780	0.3828
2008		0.7963	0.5748	0.5445	0.5170	0.5488	0.4780	0.4587	0.4620
2009			0.6796	0.6334	0.5976	0.6032	0.5532	0.5515	0.5273
2010				0.7578	0.6648	0.6692	0.5800	0.5415	0.5188
2011					0.8225	0.8186	0.7220	0.6552	0.6385
2012						0.9300	0.8425	0.7498	0.7078
2013							0.8591	0.7700	0.7282
2014								0.8826	0.8242
2015									0.9140

5.3   Relatedness with Aggregated Time Information

In order to find how *transient links*, links that do not persist over different times, and *stable links*, links that persist over time, affect the relatedness of the entities, we constructed different models of aggregated graphs as described in Section 4.1. We then applied variations of Jaccard methods described previously and our proposed extended Jaccard similarity with Reciprocal PageRank over the models.

Table 5 shows the comparison of Spearman Correlations to the KORE gold standard of different methods from each dataset over time-varying graphs and aggregated graphs. $\star, *, **$ and $***$ show that the results are significantly lower than the union model with the Extended Jaccard with Reciprocal PageRank where p-value < 0.1, p-value < 0.05, p-value < 0.01 and p-value < 0.001 respectively. We use P for direct predecessors, S for direct successors, P+S for both direct predecessors and RP for reciprocal PageRank. We can see that aggregating temporal information as a union graph with the Extended Jaccard with

---

[8] https://en.wikipedia.org/wiki/Justin_Theroux

[9] http://people.com/celebrity/jennifer-aniston-justin-theroux-engaged/

[10] http://people.com/celebrity/week-ahead-brad-jen-finalize-divorce/

Table 5: The comparison of Spearman Correlations to the KORE gold standard of different methods from each dataset over time-varying graphs and aggregated graphs. $\star$,$*$,$**$ and $***$ mean the results are significantly lower than the union model with the Extended Jaccard with Reciprocal PageRank where p-value $<$ 0.1, p-value $<$ 0.05, p-value $<$ 0.01 and p-value $<$ 0.001 respectively.

	Jaccard (P)	Jaccard (S)	Jaccard (P+S)	Extended Jaccard RP
2007	0.434629 $***$	0.398682 $***$	0.447841 $***$	0.444859 $***$
2008	0.501336 $**$	0.476366 $**$	0.526640 $*$	0.514211 $*$
2009	0.568509 $\star$	0.511564 $**$	0.578706 $\star$	0.604112
2010	0.568026 $**$	0.559898 $*$	0.585535 $**$	0.599284 $\star$
2011	0.554493 $**$	0.566468 $*$	0.563957 $*$	0.592812 $\star$
2012	0.522038 $**$	0.536122 $**$	0.533181 $**$	0.559762 $*$
2013	0.524990 $**$	0.547175 $*$	0.529726 $**$	0.562687 $*$
2014	0.478078 $***$	0.474560 $***$	0.481921 $***$	0.516497 $**$
2015	0.485979 $***$	0.491777 $**$	0.492041 $***$	0.500776 $**$
2016	0.463793 $***$	0.464086 $***$	0.467977 $***$	0.469434 $***$
Intersection	0.360936 $***$	0.351557 $***$	0.361019 $***$	0.352376 $***$
Union	0.567262 $*$	0.501396 $**$	0.535696 $**$	**0.641174**

Reciprocal PageRank gives more relatedness accuracy than the results from each dataset from time-varying graphs. Intersection graph gives the least relatedness accuracy to the KORE dataset but it represents entities that are strongly related to each other at all time. For instance, the top 3 entities with the highest relatedness scores from the Intersection model with the Extended Jaccard with Reciprocal PageRank of *Apple_Inc.* are *Steve_Jobs*, *IPhone* and *IMac*.

6   Conclusions

We have shown that our proposed graph-based extended Jaccard similarity with reciprocal PageRank outperforms the baseline text-based approach and graph-based Jaccard methods. Our relatedness score from DBpedia 2009 and DBpedia 2010, which is the time period when the KORE dataset was created, are the most correlated to the KORE dataset and using the most recent version of DBpedia loses a lot of accuracy. This shows that even in a short space of time the evaluations made by the annotators of the KORE dataset have become outdated. However, we show that by aggregating temporal information as one graph, the accuracy of relatedness is better than any other results from each dataset, demonstrating the value of considering not just the most recent version of a semantic graph but also temporal information when performing entity relatedness.

References

1. Aggarwal, N., Buitelaar, P.: Wikipedia-based Distributional Semantics for Entity Relatedness. In: Association for the Advancement of Artificial Intelligence (AAAI) Fall Symposium. AAAI-2014 (2014)
2. Bairi, R.B., Carman, M., Ramakrishnan, G.: On the Evolution of Wikipedia: Dynamics of Categories and Articles. In: Ninth International AAAI Conference on Web and Social Media (Apr 2015)
3. Borra, E., Weltevrede, E., Ciuccarelli, P., Kaltenbrunner, A., Laniado, D., Magni, G., Mauri, M., Rogers, R., Venturini, T.: Societal controversies in wikipedia articles. In: Proceedings of the 33rd annual ACM conference on human factors in computing systems. pp. 193–196. ACM (2015)
4. Ceroni, A., Georgescu, M., Gadiraju, U., Naini, K.D., Fisichella, M.: Information Evolution in Wikipedia. In: Proceedings of The International Symposium on Open Collaboration. pp. 24:1–24:10. OpenSym '14, ACM, New York, NY, USA (2014)
5. Gabrilovich, E., Markovitch, S.: Computing Semantic Relatedness Using Wikipedia-based Explicit Semantic Analysis. In: Proceedings of the 20th International Joint Conference on Artifical Intelligence. pp. 1606–1611. IJCAI'07, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2007)
6. Hoffart, J., Seufert, S., Nguyen, D.B., Theobald, M., Weikum, G.: Kore: Keyphrase overlap relatedness for entity disambiguation. In: Proceedings of the 21st ACM International Conference on Information and Knowledge Management. pp. 545–554. CIKM '12, ACM, New York, NY, USA (2012)
7. Hoffart, J., Suchanek, F.M., Berberich, K., Weikum, G.: YAGO2: A Spatially and Temporally Enhanced Knowledge Base from Wikipedia. Artificial Intelligence 194, 28 – 61 (2013)
8. Hulpus, I., Prangnawarat, N., Hayes, C.: Path-Based Semantic Relatedness on Linked Data and Its Use to Word and Entity Disambiguation. In: The Semantic Web - ISWC 2015. pp. 442–457. Springer, Cham (Oct 2015)
9. Leal, J.P., Rodrigues, V., Queirós, R.: Computing Semantic Relatedness using DBPedia. In: Simões, A., Queirós, R., da Cruz, D. (eds.) 1st Symposium on Languages, Applications and Technologies. OpenAccess Series in Informatics (OASIcs), vol. 21, pp. 133–147. Dagstuhl, Germany (2012)
10. Nunes, S., Ribeiro, C., David, G.: WikiChanges: Exposing Wikipedia Revision Activity. In: Proceedings of the 4th International Symposium on Wikis. pp. 25:1–25:4. WikiSym '08, ACM, New York, NY, USA (2008)
11. Page, L., Brin, S., Motwani, R., Winograd, T.: The pagerank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab (November 1999)
12. Radinsky, K., Agichtein, E., Gabrilovich, E., Markovitch, S.: A Word at a Time: Computing Word Relatedness Using Temporal Semantic Analysis. In: Proceedings of the 20th International Conference on World Wide Web. pp. 337–346. WWW '11, ACM, New York, NY, USA (2011)
13. Strube, M., Ponzetto, S.P.: WikiRelate! Computing Semantic Relatedness Using Wikipedia. In: Proceedings of the 21st National Conference on Artificial Intelligence - Volume 2. pp. 1419–1424. AAAI'06, Boston, Massachusetts (2006)
14. Tanimoto, T.: An Elementary Mathematical Theory of Classification and Prediction. International Business Machines Corporation (1958)
15. Whiting, S., Jose, J., Alonso, O.: Wikipedia As a Time Machine. In: Proceedings of the 23rd International Conference on World Wide Web. pp. 857–862. WWW '14 Companion, ACM, New York, NY, USA (2014)