



On the Temporality of Priors in Entity Linking

Renato Stoffalette João^(✉) 

L3S Research Center, Appelstraße 9A, 30167 Hannover, Germany
joao@L3S.de

Abstract. Entity linking is a fundamental task in natural language processing which deals with the lexical ambiguity in texts. An important component in entity linking approaches is the mention-to-entity *prior* probability. Even though there is a large number of works in entity linking, the existing approaches do not explicitly consider the time aspect, specifically the *temporality* of an entity's prior probability. We posit that this prior probability is temporal in nature and affects the performance of entity linking systems. In this paper we systematically study the effect of the prior on the entity linking performance over the temporal validity of both texts and KBs.

Keywords: Entity linking · Entity disambiguation · Knowledge base

1 Introduction and Motivation

Entity linking is a well studied problem in natural language processing which involves the process of identifying ambiguous entity mentions (i.e persons, locations and organisations) in texts and linking them to their corresponding unique entries in a reference knowledge base. There has been numerous approaches and eventually systems proposing solutions to the task at hand. To mention a few, AIDA [9], Babely [14], WAT [15] and AGDISTIS [18] for example, rely on graph based algorithms and the most recent approaches rely on techniques such as deep neural networks [8] and semantic embeddings [10, 19].

An important component in most approaches is the probability that a mention links to one entity in the knowledge base. The prior probability, as suggested by Fader et al. [4], is a strong indicator to select the correct entity for a given mention, and consequently adopted as a baseline. Computation of this prior is typically done over knowledge sources such as Wikipedia. Wikipedia in fact provides useful features and has grounded several works on entity linking [1, 2, 9, 12, 13].

An entity's popularity is temporally sensitive and may change due to short term events. Fang and Chang [6] noticed the probability of entities mentioned in texts often change across time and location in micro blogs, and in their work they modeled spatio-temporal signals for solving ambiguity of entities. We, on

the other hand, take a macroscopic account of time, where perceivably a larger fraction of mention to entity bindings might not be observable in the short time duration but are only evident over a longer period of time, i.e., over a year. These changes might be then reflected in a reference knowledge base and disambiguation methods can produce different results for a given mention at different times.

When using a 2006 Wikipedia edition as a reference knowledge base for example, the mention *Amazon* shows different candidates as linking destinations, but the most popular one is the entity page referring to *Amazon River*, whilst when using a 2016 Wikipedia edition, the same term leads to the page about the e-commerce company *Amazon.com* as the most popular entity to link to.

In this paper, we systematically study the effect of temporal priors on the disambiguation performance by considering priors computed over snapshots of Wikipedia at different points in time. We also consider benchmarks that contain documents created and annotated at different points in time to better understand the potential change in performance with respect to the temporal priors.

We firstly show that the priors change over time and the overall disambiguation performance using temporal priors show high variability. This strongly indicates that temporal effects should be not only taken into account in (a) building entity linking approaches, but have major implications in (b) evaluation design, when baselines that are trained on temporally distant knowledge sources are compared.

2 Problem Definition

In this section we briefly define the entity linking task as well as describe our methodology.

Consider a document d from a set of documents $D = \{d_1, d_2, \dots, d_n\}$, and a set of mentions $M = \{m_1, m_2, \dots, m_n\}$ extracted from d . The goal of the entity linking is to find a unique identity represented by an entity e from a set of entities $E = \{e_1, e_2, \dots, e_n\}$, with relation to each mention m . The set of entities E is usually extracted from a reference knowledge base KB .

A typical entity linking system generally performs the following steps: (1) mention detection which extracts terms or phrases that may refer to real world entities, and (2) entity disambiguation which selects the corresponding knowledge-base KB entries for each ambiguous mention.

Since we take into account the time effect on the disambiguation task, we now pose entity linking at a specific time t as follows. Given a document $d^t \in D^t$ and a set of mentions $M = \{m_1, m_2, \dots, m_n\}$ from document d^t , the goal of the entity linking at time t is to find the correct mapping entity $e^t \in E^t$ with relation to the mention m . The difference now is that the set of entities E^t is extracted from the reference knowledge base KB at different time periods.

2.1 Candidate Entities Generation and Ranking

As suggested by Fader et al. [4], the entity’s prior probability is a strong indicator to select the correct entity for a given mention. In our case the entity’s prior probability is directly obtained from the Wikipedia corpus. To calculate entities’ probability, we parsed all the articles from a Wikipedia corpus and collected all terms that were inside double square brackets in the Wikipedia articles. $[[\textit{Andy Kirk (footballer)} \mid \textit{Kirk}]]$ for instance, represents a pair of mention and entity where *Kirk* is the mention term displayed in the Wikipedia article and *Andy Kirk (footballer)* is the title of the Wikipedia article corresponding to the real world entity. In this way we created a list of mentions and possible candidate entities according to each Wikipedia snapshot used in this paper.

Table 1. Information about the Wikipedia editions used for mining mention and entities. #Pages refers only to the number of entities’ pages, excluding special pages.

Year	Date	#Pages
2006	30/11/2006	~1.4 M
2008	03/01/2008	~1.9 M
2010	15/03/2010	~2.8 M
2012	02/09/2012	~3.5 M
2014	06/11/2014	~4.1 M
2016	01/07/2016	~5.1 M

The probability of a certain entity e^t given a mention m was only calculated if the entity had a corresponding article inside Wikipedia at time t . Thus, the probability $P(e^t \mid m)$ that a mention m links to a certain entity e^t is given by the number of times the mention m links to the entity e^t over the number of times that m occurs in the whole corpus at time t .

We created dictionaries of mentions and their referring entities ranked by popularity of occurrence for every Wikipedia edition as seen on Table 1. As an example of mention and its ranked candidate entities, in the KB created from the 2016 Wikipedia edition, the mention *Obama* refers in 86.15% of the cases to the president *Barack Obama*, 6.47% to the city *Obama*, *Fukui* in Japan, 1.79% to the genus of planarian species *Obama (genus)*, and so on and so forth.

We filtered out mentions that occurred less than 100 times for simplicity matters in the whole corpus and for every mention we checked whether the referring candidate entities pointed to existing pages inside the Wikipedia corpus at a given time, and only after these steps we calculated the prior probability values the entities.

Our framework supports multiple selection of mention-entity dictionaries created from different KBs based on Wikipedia snapshots from different years.

3 Experiments and Results

3.1 Datasets

In order to evaluate our experiments we employed some data sets that are widely used benchmark datasets for entity linking tasks. *ACE04* is a news corpus introduced by Ratinov et al. [16] and it is a subset from the original ACE co-reference data set [3]. *AIDA/CONLL* is proposed by Hoffart et al. [9] and it is based on the data set from the CONLL 2003 shared task [17]. *AQUAINT50* was created in the work proposed by Milne and Witten [13], and is a subset from the original AQUAINT newswire corpus [7]. *IITB* is a dataset extracted from popular web pages about sports, entertainment, science and technology, and health¹², and it was created in the work proposed by Kulkarni et al. [11]. *MSNBC* was introduced by Cucerzan [2] and contains news documents from 10 MSNBC news categories. Table 2 shows more details about these datasets including the number of documents, documents' publication time, number of annotations as well as the reference knowledge base time.

3.2 Prior Probability Changes

In many entity linking systems, the entity mentions that should be linked are given as the input, hence the number of mentions generated by the systems equals the number of entity mentions that should be linked. For this reason most researchers use accuracy to evaluate their method's performance. Accuracy is a straightforward measure calculated as the number of correctly linked mentions divided by the total number of mentions.

Since we take into account the time variation, we only calculated accuracy over the total number of annotations that persisted across time, i.e. the entities from the ground truth that were also present in every Wikipedia edition used in this paper. Table 3 shows the accuracy calculated on the ground truth datasets using the prior probability model from different time periods. We can observe

Table 2. #Docs is the number of documents. Docs Year is the documents' publication time. #Annotations is the number of annotations (Number of non-NIL annotations). Annot. Year is the reference KB time period where the annotations were taken from.

Dataset	#Docs	Docs Year	#Annotations	Annot. Year
ACE04 [16]	57	2000	257	2010
AIDA/CONLL [9]	231	1996	4.485	2010
AQUAINT50 [13]	50	1998–2000	727	2007
IITB [11]	107	2008	12.099	2008
MSNBC [2]	20	2007	747	2006

¹ <http://news.google.com/>.

² <http://www.espnstar.com/>.

an accuracy change from 77.19% to 82.63% on *ACE04* using models created from Wikipedia 2006 and 2010 editions respectively, from 64.80% to 69.07% for *AQUAINT50* using models from 2006 and 2012 editions, from 64.13% to 68.16% for *AIDA/CONLL* using models from 2008 and 2014, from 46.60% to 49.76% on *IITB* using models from 2014 and 2006, and for *MSNBC* a change from 63.82% to 65.86% using models from Wikipedia 2012 and 2008 editions respectively.

Even though it is out of the scope of this work to spot a temporal trend on the entities changes when using knowledge bases from different time periods, we can clearly see there is some temporal variability which is easily observed by the influence on the accuracy calculated over the ground truth datasets. We observe that a simplistic popularity only based method that takes into account reference *KBs* from different time periods can produce an improvement of 5.4% points in the best case for the *ACE04* dataset and 2.0% points in the worst case for *MSNBC* dataset.

3.3 Comparing Ranked Entities

We detected distinct changes when it comes to entity linking using Wikipedia as a knowledge base. The first case occurs when the entity page title changes but still refers to the same entity in the real world. For example in the 2006 Wikipedia edition the mention *Hillary Clinton* showed higher probability of linking to the referring entity page titled *Hillary Rodham Clinton* and in the 2016 Wikipedia edition, the same mention was most likely to be linked to the entity page titled *Hillary Clinton*. In this case only the entity page title changed but they both refer to the same entity in the real world.

The second case happens when an entity's popularity actually changes over time. For example in the 2006 Wikipedia edition, the mention *Kirk* was most likely to be linked to the entity page titled *James T. Kirk* whereas in the 2016 Wikipedia edition the same mention showed a higher probability of linking to the entity page titled *Andy Kirk (footballer)*.

Another observation is the case when an entity mention that was considered unambiguous in the past and became ambiguous in a newer Wikipedia edition due to the addition of new information to Wikipedia. For example in the 2006 Wikipedia edition the mention *Al Capone* showed a single candidate entity, the north american gangster and businessman *Al Capone*, while in the newer 2016 Wikipedia edition, the same mention showed more candidate entities, including the former one plus a movie, a song, and other figures with the same name.

Top Ranked Entity Changes. Initially we were only concerned with the top ranked candidate entity for each mention. We made comparisons between the dictionaries of mentions from Wikipedia editions 2006 and 2016 and despite the fact of observing 33,531 mentions in the 2006 version and 161,264 mentions in the

2016 version, only 31,123 mentions appeared in both editions. Moreover, when we take into consideration both the ambiguous and unambiguous mentions, in 9.44% of the cases the mentions change their top ranked candidate entities, whilst when removing the unambiguous mentions this number increases to 15.36%. This is mainly due to the fact that most of the unambiguous mentions keep the same entity bindings, even though we spotted cases of mentions that were unambiguous and became ambiguous in a more recent knowledge base.

Table 3. Accuracy of the models on different data sets across different time periods.

Dataset	2006	2008	2010	2012	2014	2016
ACE2004	77.19	81.17	82.63	80.96	80.54	79.49
AIDA/CONLL_testb	61.86	64.13	66.47	67.78	68.16	68.14
AQUAINT50	64.80	68.18	68.92	69.07	67.30	66.86
IITB	49.76	49.43	49.50	47.78	46.60	47.60
MSNBC	65.30	65.86	65.67	63.82	64.56	65.67

Top 5 Entities Changes. In another experiment we wanted to calculate the entities rank correlation. One way to calculate rank correlation for lists that do not have all the element in common, is to ignore the non conjoint elements, but unfortunately this approach is not satisfactory since it throws away information. Hence, a more satisfactory approach, as proposed by Fagin et al. [5], is to treat an element i which appears ranked in list L_1 and does not appear in list L_2 , at position $k+1$ or beyond, considering L_2 's depth is k . This measure was used to assess the changes in the top 5 candidate entities rank positions.

We calculated the rank correlation for 18,727 mentions, since this is the number of mentions that are ambiguous and appears both in the 2006 and 2016 Wikipedia corpus. We normalized our results so the values would lie between [0,1]. Any value close to 0 means total agreement while any value close to 1 means total disagreement. Thus we observed an average value of 0.59 with a variance of 0.05 and a standard deviation of 0.21. We noticed that in 71.98% of the cases the rank correlation values are greater than 0.5. That tells us there is some significant number of changes in the candidate entities' rank's positions. Table 4 shows the mention *Watson* and its top 5 candidate entities together with their respective prior probabilities extracted from two different Wikipedia editions, one from 2006 and one from 2016.

Table 4. A mention example and its top 5 ranked candidate entities captured from two Wikipedia editions.

Mention	Entity	$P(e^t)$	Year
Watson	Doctor Watson	0.146	2006
	James D. Watson	0.130	
	Watson, Australian Capital Territory	0.115	
	Division of Watson	0.076	
	Watson	0.061	
Watson	Watson (computer)	0.068	2016
	Ben Watson (footballer, born July 1985)	0.054	
	Je-Vaughn Watson	0.050	
	Jamie Watson (soccer)	0.047	
	Arthur Watson (footballer, born 1870)	0.043	

4 Conclusions and Future Work

In this work we conducted experiments with different Wikipedia editions and also created an entity linking model that uses the entity’s prior probability calculated over different Wikipedia snapshots. One limitation of previous works is the fact that the systems are trained on a fixed time Wikipedia edition. An entity’s prior probability is temporal in nature, and we have observed in our experiments that mention to entity bindings change over time. We could clearly see some temporal variability which should be taken into account for entity linking system’s evaluations. As future work we plan to extend this paper’s experimental setup and build a ground truth for temporal entity linking as well as try to create an adaptive entity linking system.

References

1. Bunescu, R.C., Pasca, M.: Using encyclopedic knowledge for named entity disambiguation. In: EACL, Trento, Italy, 3–7 April 2006 (2006)
2. Cucerzan, S.: Large-scale named entity disambiguation based on Wikipedia data. In: EMNLP-CoNLL, Prague, Czech Republic, 28–30 June 2007, pp. 708–716 (2007)
3. Doddington, G.R., Mitchell, A., Przybocki, M.A., Ramshaw, L.A., Strassel, S.M., Weischedel, R.M.: The automatic content extraction (ACE) program-tasks, data, and evaluation. In: LREC, Lisbon, Portugal, vol. 2, p. 1 (2004)
4. Fader, A., Soderland, S., Etzioni, O.: Scaling Wikipedia-based named entity disambiguation to arbitrary web text. In: WIKIAI (2009)
5. Fagin, R., Kumar, R., Sivakumar, D.: Comparing top k lists. In: ACM-SIAM, Baltimore, Maryland, USA, 12–14 January 2003, pp. 28–36 (2003)
6. Fang, Y., Chang, M.: Entity linking on microblogs with spatial and temporal signals. *TACL* **2**, 259–272 (2014)

7. Graff, D.: The AQUAINT Corpus of English News Text: [Content Copyright] Portions 1998–2000 New York Times Inc., 1998–2000. Associated Press Inc., 1996–2000 Xinhua News Service. Linguistic Data Consortium (2002)
8. He, Z., Liu, S., Li, M., Zhou, M., Zhang, L., Wang, H.: Learning entity representation for entity disambiguation. In: ACL 2013, Sofia, Bulgaria, 4–9 August 2013, Volume 2: Short Papers, pp. 3–34 (2013)
9. Hoffart, J., et al.: Robust disambiguation of named entities in text. In: EMNLP 2011, Edinburgh, UK, 27–31 July 2011, pp. 782–792 (2011)
10. Huang, H., Heck, L., Ji, H.: Leveraging deep neural networks and knowledge graphs for entity disambiguation. CoRR abs/1504.07678 (2015)
11. Kulkarni, S., Singh, A., Ramakrishnan, G., Chakrabarti, S.: Collective annotation of Wikipedia entities in web text. In: SIGKDD, Paris, France, 28 June–1 July 2009, pp. 457–466 (2009)
12. Mihalcea, R., Csomai, A.: Wikify!: linking documents to encyclopedic knowledge. In: CIKM 2007, Lisbon, Portugal, 6–10 November 2007, pp. 233–242 (2007)
13. Milne, D.N., Witten, I.H.: Learning to link with Wikipedia. In: CIKM 2008, Napa Valley, California, USA, 26–30 October 2008, pp. 509–518 (2008)
14. Moro, A., Raganato, A., Navigli, R.: Entity linking meets word sense disambiguation: a unified approach. *TACL* **2**, 231–244 (2014)
15. Piccinno, F., Ferragina, P.: From TagME to WAT: a new entity annotator. In: ERD 2014, Gold Coast, Queensland, Australia, 11 July 2014, pp. 55–62 (2014)
16. Ratinov, L., Roth, D., Downey, D., Anderson, M.: Local and global algorithms for disambiguation to Wikipedia. In: The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Proceedings of the Conference, Portland, Oregon, USA, 19–24 June 2011, pp. 1375–1384 (2011)
17. Sang, E.F., De Meulder, F.: Introduction to the CoNLL-2003 shared task: language-independent named entity recognition. arXiv preprint cs/0306050 (2003)
18. Usbeck, R., et al.: AGDISTIS - graph-based disambiguation of named entities using linked data. In: Mika, P., et al. (eds.) ISWC 2014. LNCS, vol. 8796, pp. 457–471. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-11964-9_29
19. Zwicklbauer, S., Seifert, C., Granitzer, M.: Robust and collective entity disambiguation through semantic embeddings. In: SIGIR 2016, Pisa, Italy, 17–21 July 2016, pp. 425–434 (2016)