

Entity-aware ELMo: Learning Contextual Entity Representation for Entity Disambiguation

Abstract

We present a new local entity disambiguation system. The key to our system is a novel approach for learning entity representations. In our approach we learn an entity aware extension of Embedding for Language Model (ELMo) which we call Entity-ELMo (E-ELMo). Given a paragraph containing one or more named entity mentions, each mention is first defined as a function of the entire paragraph (including other mentions), then they predict the referent entities. Utilizing E-ELMo for local entity disambiguation, we outperform all of the state-of-the-art local and global models on the popular benchmarks by improving about 0.5% on micro average accuracy for AIDA test-b with Yago candidate set. The evaluation setup of the training data and candidate set are the same as our baselines for fair comparison.

1 Introduction

Named Entity Disambiguation (NED) is an essential task in natural language processing that resolves mentions in a document to their referent entities in a Knowledge Base (KB). A notable differentiating factor between NED systems is whether the global joint inference is used to resolve all the mentions in the same document collectively, separating the local and global NED models.

Local NED systems disambiguate a mention individually by utilizing the local compatibility between the mention (and its textual context) and its candidate entities. Global NED models further consider the global coherence between assigned entities via structured prediction. Empirically, they have been shown to consistently outperform local models as they capture long-range document-wise information. Nevertheless, in this work, we demonstrate that a simple local model, when equipped with the right context and entity representation, can achieve competitive, even superior performance compared to the state-of-the-art global models.

There are several limitations with the existing local models [He *et al.*, 2013; Sun *et al.*, 2015; Hu *et al.*, 2015; Yamada *et al.*, 2016; Yamada *et al.*, 2017; Ganea and Hofmann, 2017; Sil *et al.*, 2018; Le and Titov, 2018] in capturing contextual

dependencies between mention and candidate entities. Yamada *et al.* [Yamada *et al.*, 2016; Yamada *et al.*, 2017] and Ganea *et al.* [Ganea and Hofmann, 2017] learn entity representations by encoding the co-occurrence statistics between the entity and the words in their context. Each word is considered individually and hence its syntactic and semantic roles in the sentence are overlooked.

In another effort, Sil *et al.* [Sil *et al.*, 2018] propose a model in which a recurrent neural network is utilized to compare the context of the query mention with the canonical pages of all candidate entities. To make their approach practical, their method is limited to only consider the first paragraph in each canonical page, which may not provide sufficient information for representing an entity throughout a large corpus. A convolutional architecture has also been explored for encoding sentence-level contextual information [Sun *et al.*, 2015]. However its fixed-length windows limit the scope of the context.

Recently, Embedding for Language Model (ELMo) [Peters *et al.*, 2018] is introduced to produce context-sensitive representation of words as a function of the entire sentence. Although ELMo produces context-sensitive representations for words in a sentence, its learning objective is unaware of the entities. For instance in the phrase “*Jordan as a member of the Tar Heels’ national championship team,*” the language model predicts ambiguous mention *Jordan*, instead of the entity *Michael Jordan*.

There are many virtues in ELMo which make it a suitable choice to be used for learning entity representations: 1) Each token is represented as a function of its surrounding context via bi-directional RNNs which can potentially capture the dependencies between a mention and all the surrounding context as well as other named entity mentions in the context. 2) Deep layers of ELMo capture syntactic and semantic dependencies both of which are required for NED, and 3) ELMo is trainable on un-annotated corpora which is an important means to transfer information.

In this work we introduce a novel approach for learning contextual entity representations by learning an entity-aware extension of ELMo, which is surprisingly effective. The learning mechanism which we call E-ELMo trains the language model to predict the grounded entity when encountering its mentions, as opposed to the words in the mentions. This modification affords us context-rich entity representa-

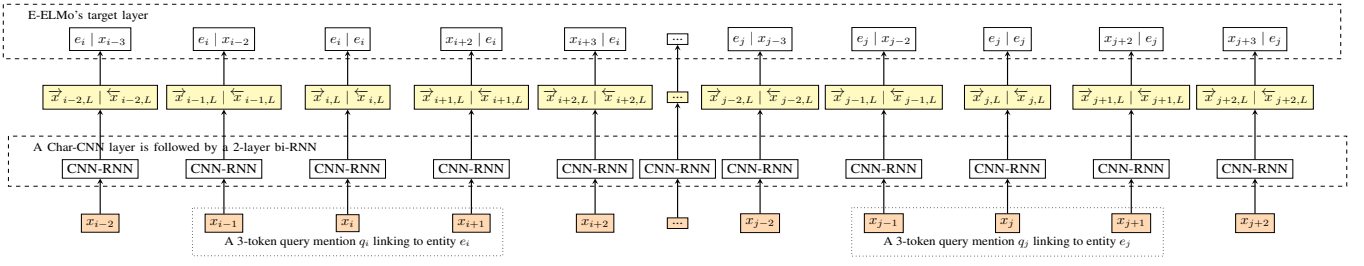


Figure 1: Bidirectional language model predicts different targets in ELMo and E-ELMo.

tions that are well-suited for disambiguation of the named entities with just local contexts. Incorporating E-ELMo’s representations into a very simple local model, we achieve a superior or competitive performance on popular benchmarks compared to the state-of-the-art global models. Here we need to mention that in a similar work [Yamada *et al.*, 2020] (published few days after this work) the effectiveness of incorporating entities into BERT [Devlin *et al.*, 2019] is experimentally analysed and verified for entity linking.

2 Entity-ELMo (E-ELMo)

Here we first briefly review ELMo (Sec 2.1) and then explain E-ELMo (Sec 2.2). As an example, we will consider an instance paragraph, which is a sequence of T tokens containing two mentions with three tokens each, as shown in Figure 1. Here mentions $[x_{i-1}, x_i, x_{i+1}]$ and $[x_{j-1}, x_j, x_{j+1}]$ refer to entities e_i and e_j respectively. Tokens x_{i-2} and x_{i+2} are the preceding and succeeding tokens of the first mention respectively. Note that a paragraph might include any number of mentions, each may contain any number of tokens.

2.1 ELMo Review

For a given sequence, ELMo produces word representations on top of a 2-layer bi-RNN with character convolutions as input. For each direction, ELMo first computes a context-independent representation for each token at position k by applying a character-based CNN. It then passes the token representations through a 2-layer LSTMs. As a result, each LSTM layer outputs a context-dependent representation $\vec{x}_{k,j}$ and $\overleftarrow{x}_{k,j}$ for layers $j \in \{1, \dots, L\}$ for the forward and backward directions respectively. The outputs of the last layer i.e. $\vec{x}_{k,L}$ and $\overleftarrow{x}_{k,L}$ are given to a Softmax layer to predict the next and previous tokens x_{k+1} and x_{k-1} respectively.

ELMo’s objective is to jointly maximize the log likelihood of the forward and backward passes:

$$l_{ELMo} = \sum_{k=1}^n \log p(x_k | \vec{x}_{k-1,L}, \vec{\Theta}_{LM}, \Theta_S, \Theta_x) + \sum_{k=1}^n \log p(x_k | \overleftarrow{x}_{k+1,L}, \overleftarrow{\Theta}_{LM}, \Theta_S, \Theta_x)$$

where $\vec{\Theta}_{LM}$, $\overleftarrow{\Theta}_{LM}$, Θ_S and Θ_x are the parameters for the forward and backward bi-RNNs, the Softmax and the char-CNN layers respectively.

2.2 E-ELMo Model

As shown in Figure 1, E-ELMo is in fact ELMo with entities incorporated in the target layer, replacing the targets from the mention words to the grounded entity. In particular, the target for position $k \in I_i = \{i-2, i-1, i\}$ for the forward and $k \in J_i = \{i, i+1, i+2\}$ for the backward directions should be entity e_i . The log likelihood objective of the E-ELMo is the sum of the log likelihood for both words and entities as follows:

$$l_{E-ELMo} = l_w + l_e \quad (1)$$

where

$$l_e = \sum_i \left(\sum_{k-1 \in I_i} \log p(e_i | \vec{x}_{k-1,L}, \vec{\Theta}_{LM}, \Theta_E, \Theta_x) + \sum_{k+1 \in J_i} \log p(e_i | \overleftarrow{x}_{k+1,L}, \overleftarrow{\Theta}_{LM}, \Theta_E, \Theta_x) \right)$$

where l_w is equal to l_{ELMo} minus the terms predicting entities and Θ_E is the entity parameters. Note that it is important to optimize the entity vectors on the unit sphere to yield qualitative embeddings for NED. To maximize Eq. 1, we consider several different configurations of E-ELMo:

Config-a: We freeze all parameters except for Θ_E .

Config-b: All the parameters are fine-tuned.

Config-c: All the parameters are fine tuned but we redefine Eq. 1 to be $l_{E-ELMo} = l_e$.

3 Local Entity Disambiguation Model

To evaluate the effectiveness of the learned entity representations, we consider a simple local entity disambiguation model to rank the candidate entities. Given a query mention $[x_{i-1}, x_i, x_{i+1}]$ with its local context $x_1, \dots, x_{i-2}, [x_{i-1}, x_i, x_{i+1}], x_{i+2}, \dots, x_T$, the context is fed to the E-ELMo corresponding to one of the config-a, b and c to produce $\vec{x}_{k,L}$ and $\overleftarrow{x}_{k,L}$ for all positions $k \in 1 \dots T$. The context representation for the mention is given by concatenating $\vec{f}_c = \frac{1}{|I_i|} \sum_{k \in I_i} (\vec{x}_{k,L})$ and $\overleftarrow{f}_c = \frac{1}{|J_i|} \sum_{k \in J_i} (\overleftarrow{x}_{k,L})$.

We also utilize the following basic features introduced by prior work [Yamada *et al.*, 2016; Yamada *et al.*, 2017; Sil *et al.*, 2018; Ganea and Hofmann, 2017; Le and Titov, 2018]. Prior Compatibility (f_p): We consider $p(e|m)$, the prior probability that an entity e is linked to a mention string m as prior evidence. String Matching (f_s): We use ten lexical features $f_s[1] \dots f_s[10]$ listed in [Shahbazi *et al.*, 2018] to

capture the lexical similarity between the query mention and the surface string of the entity.

Our local model first transforms the scalar features $f_p, f_s[1] \dots f_s[10]$ through a bin layer [Shahbazi *et al.*, 2018; Sil *et al.*, 2018] to project each feature f to a higher dimensional \hat{f} . Given a scalar variable x , the binning layer projects x to dimension d as follows:

$$\mathbf{p} = [e^{-(\epsilon_1 \|x-x_1\|)^2} \dots e^{-(\epsilon_d \|x-x_d\|)^2}]$$

Binning introduces parameters ϵ_i and x_i for $i \in \{1 \dots d\}$ to project x to vector \mathbf{p} . The concatenated feature $[\hat{f}_p; \hat{f}_s[1] \dots \hat{f}_s[10]; \vec{f}_c; \overleftarrow{f}_c]$ and $\Theta_E[e]$, the representation for e learned by E-ELMo, are then given to a 2-layer feed forward neural network with Relu activation to compute the final score for entity e .

4 Experiments

4.1 Entity disambiguation datasets

We evaluate our NED on the following benchmarks:

- AIDA-CoNLL [Hoffart *et al.*, 2011]: This dataset which is one of the biggest dataset for NED, contains training (AIDA-train), validation (AIDA-A) and test (AIDA-B) sets.
- MSNBC (MSB), AQUAINT (AQ) and ACE2004 (ACE) datasets [Guo and Barbosa, 2016]
- WNED-WIKI (WW) and WNED-CWEB (CWEB): These datasets are bigger and are built from the ClueWeb and Wikipedia corpora by [Guo and Barbosa, 2016; Gabrilovich *et al.*, 2013]
- TAC 2010 [Hoffart *et al.*, 2011] and TAC 2010 [Ji *et al.*, 2010]: This dataset is very popular with the baselines for both local and global model.

To follow similar setup to our baselines [Ganea and Hofmann, 2017; Le and Titov, 2018] we use the train, test, validation splits and also candidate sets and also prior values $p(e|m)$ for all the the datasets from [Ganea and Hofmann, 2017]¹

Table 1 shows the statistics of the datasets:

Dataset	#mens	#docs	#mens in doc	cands recall
AIDA-train	18448	946	19.5	-
AIDA-A (valid)	4791	216	22.1	96.9%
AIDA-B (test)	4485	231	19.4	98.2%
MSNBC	656	20	32.8	98.5%
AQUAINT	727	50	14.5	94.2%
ACE2004	257	36	7.1	90.6%
WNED-CWEB	11154	320	34.8	91.1%
WNED-WIKI	6821	320	21.3	92%
TAC-2010(test)	1020	1013	1	93%

Table 1: Datasets used in our NED

¹Available at <https://github.com/dalab/deep-ed>

4.2 Training E-ELMo

We follow the same experimental setup as our baselines. We train E-ELMo on a subset φ of Wikipedia corpus extracted by [Ganea and Hofmann, 2017] and also used by [Le and Titov, 2018]. We initialize $\vec{\Theta}_{LM}, \overleftarrow{\Theta}_{LM}, \Theta_s$, and Θ_x using the original ELMo pre-trained on 5B tokens. We initialize Θ_E for each entity to be the average of its title tokens in Θ_s . The number of negative samples for both words and entities is set to 8192. Training of E-ELMo is via AdaGrad [Duchi *et al.*, 2011] with a learning rate 0.1 for 10 epochs.

4.3 Training the Local Model

The binning layer in our local model projects a scalar f_p to 15-d vector and each $f_s[k]$ to 10-d vector. The size of $\vec{f}_p, \overleftarrow{f}_c, \theta_e$ are each 512. We use the cross-entropy loss function and dropout of 0.7 on the feed forward neural network. Training is done by ADAM [Kingma and Ba, 2014] with learning rate 0.001.

4.4 Entity Disambiguation Results

We first evaluate our NED systems on the two most commonly used benchmarks: AIDA-CoNLL [Hoffart *et al.*, 2011] and TAC 2010 [Ji *et al.*, 2010]. We follow the same setup as our baselines [Ganea and Hofmann, 2017; Le and Titov, 2018] and use the same train, test, validation splits, the same candidate sets and the prior feature values $p(e|m)$. Please see Appendix A.2 for detailed information about the datasets used in the experiments. For AIDA-CoNLL, the existing literature has considered two different candidate sets: (aida-Yago), which is extracted and used by [Ganea and Hofmann, 2017; Le and Titov, 2018] and (aida-HP): a less ambiguous one extracted by [Perschina *et al.*, 2015]. In our experiments we evaluate our methods using both candidate sets.

Table 3 presents the results of our models as well as baseline local and global models on AIDA-CoNLL and TAC 2010. Our baselines include a large number of state-of-the-art local (seven) and global (eight) models proposed in recent years. In addition, we also consider an alternative NED system based on the original ELMo. In particular, this model, referred to as ELMo_o, is identical to our local NED model except that E-ELMo is replaced with the original ELMo and each entity e is represented by averaging the sentence-level representations of all query mentions linking to e in Wikipedia using the original ELMo.

We first note that the proposed models are highly competitive, achieving substantial improvements over prior local methods, and even outperforming prior state-of-the-art global methods on both benchmarks. Comparing the performance of E-ELMo_a with the other two variants, we observe that fine tuning the language model with the entity-aware objective tends to further improve the performance.

We also note that the performance of the baseline model ELMo_o is substantially lower on both datasets, suggesting that the performance gain achieved by our models is not simply due to improved word representations from ELMo. Rather, it is critical to integrate entities into the language model to learn useful representation for NED.

models	MSB	AQ	ACE	CWEB	WW
global: [Guo and Barbosa, 2016]	92	87	88	77	84.5
global: [Ganea and Hofmann, 2017]	93.7 ± 0.1	88.5 ± 0.4	88.5 ± 0.3	77.9 ± 0.1	77.5 ± 0.1
global: [Le and Titov, 2018]	93.9 ± 0.2	88.3 ± 0.6	89.9 ± 0.8	77.5 ± 0.1	78.0 ± 0.1
local: E-ELMo _b	92.3 ± 0.1	90.1 ± 0.3	88.7 ± 0.1	78.4 ± 0.2	79.8 ± 0.2
local: E-ELMo _c	92.0 ± 0.1	89.6 ± 0.1	87.6 ± 0.1	77.5 ± 0.3	78.4 ± 0.1

Table 2: Results on five out-domain test sets

models	aida (HP)	aida (YAGO+KB)	tac (KB)
local models			
[Francis-Landau <i>et al.</i> , 2016]	85.5	-	-
[Sil and Florian, 2016]	86.2	-	78.6
[Yamada <i>et al.</i> , 2016]	90.9	87.2	84.6
[Yamada <i>et al.</i> , 2017]	94.7	-	87.7
[Ganea and Hofmann, 2017]	-	88.8	-
[Sil <i>et al.</i> , 2018]	94.0	-	87.4
[Shahbazi <i>et al.</i> , 2018]	90.89	-	85.73
global models			
[Pershina <i>et al.</i> , 2015]	91.8	-	-
[Chisholm and Hachey, 2015]	-	88.7	-
[Globerson <i>et al.</i> , 2016]	92.7	91.0	87.2
[Yamada <i>et al.</i> , 2016]	93.1	91.5	85.5
[Guo and Barbosa, 2016]	-	89.0	-
[Ganea and Hofmann, 2017]	-	92.22 ± 0.14	-
[Shahbazi <i>et al.</i> , 2018]	94.44	-	87.9
[Le and Titov, 2018]	-	93.07 ± 0.27	-
this work (local)			
ELMo _c (baseline)	86.2 ± 0.16	84.01 ± 0.19	81.12
E-ELMo _a	95.41 ± 0.11	92.07 ± 0.16	88.15
E-ELMo _b	96.22 ± 0.20	93.1 ± 0.22	87.36
E-ELMo _c	96.24 ± 0.12	93.46 ± 0.14	88.27
ablation (-prior -lexical features)			
E-ELMo _b	-	92.12	83.0
E-ELMo _c	-	92.30	84.4
[Ganea and Hofmann, 2017] -	-	86.34	69.28

Table 3: Results on TAC 2010 and AIDA test-b

The ablation rows in Table 3 present the performance of our models versus the local attention based models [Ganea and Hofmann, 2017] without using the prior and lexical features. As shown in the table our models see only slight reduction in performance when prior and lexical features are removed. In contrast, the local attention based model suffered a substantial performance dip. This suggests that the contextual dependencies captured by E-ELMo are significantly richer, possibly because E-ELMo not only captures syntactic and semantic dependencies but also captures lexical and prior dependencies between the mention and the entity with the help of the char CNN in the first layer.

Table 2 presents the performance of E-ELMo_b compared to state-of-the-art global models on five additional open domain datasets. Because these datasets are not widely used in prior studies, we can only compare to three prior methods, all of which are global methods. As can be seen from the table, E-ELMo_b is very competitive compared to the prior state-of-the-art results achieved by global models on these datasets as well.

4.5 Analysis of Results

We take a closer look at the performance of our method in comparison with a popular global model [Ganea and Hofmann, 2017] (referred to as M_g) on the AIDA-CoNLL dataset (YAGO+KB), focusing on the 4400 test queries with gold in the candidate set. In particular we are interested in how entity

frequencies and the number of entities in the same document influence the performance of our local model in comparison with the global model.

As shown in Table 4, E-ELMo_c has substantially higher per-

entity e	# mens	E-ELMo _b	E-ELMo _c	M_g
freq in Wiki				
1-10	459	94.54	95.42	91.93
11-50	503	93.04	92.84	92.44
≥51	3438	94.50	96.24	94.21
#ents in doc				
1-4	90	96.66	96.68	93.33
5-9	341	97.06	97.94	94.42
10-19	849	95.64	96.46	92.08
≥20	3120	93.62	95.32	94.15

Table 4: Comparing E-ELMo_c and global M_g .

formance compared to M_g in dealing with entities with low frequency in the Wikipedia, suggesting that by unifying the entity representations with the word representation through E-ELMo, we can learn effective representations for rare entities. It is also observed that E-ELMo_c performs consistently well on documents with different number of mentions, but the performance gap with m_g is small for documents with over 20 mentions. This is consistent with expectation because such documents tend to benefit more from the global models.

5 Conclusions

We introduced a novel approach for learning deep contextual entity representation by learning an entity aware extension of ELMo called E-ELMo. We also proposed a local entity disambiguation model which utilizes E-ELMo as its key component. The results demonstrate that our local model with very basic features achieves the best reported performance on AIDA-CoNLL and TAC-2010 with an improvement of about .5% over the latest global model on AIDA-CoNLL. The model is also competitive to the global models on open domain datasets.

References

- [Chisholm and Hachey, 2015] Andrew Chisholm and Ben Hachey. Entity disambiguation with web links. *Transactions of the Association of Computational Linguistics*, page 3:145–156, 2015.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *In Proceedings of the 2019 Conference of the*

- North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1:4171–4186, 2019.
- [Duchi *et al.*, 2011] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.
- [Francis-Landau *et al.*, 2016] Matthew Francis-Landau, Greg Durrett, and Dan Klein. semantic similarity for entity linking with convolutional neural networks. *Annual Conference of the North American Chapter of the Association for Computational Linguistics, NAACL*, 2016.
- [Gabrilovich *et al.*, 2013] Evgeniy Gabrilovich, Michael Ringgaard, and Amarnag Subramanya. Facc1: Freebase annotation of cluweb corpora. 2013.
- [Ganea and Hofmann, 2017] Octavian-Eugen Ganea and Thomas Hofmann. Deep joint entity disambiguation with local neural attention. *In Proc. of Empirical Methods in Natural Language Processing*, 2017.
- [Globerson *et al.*, 2016] Amir Globerson, Nevena Lazic, Soumen Chakrabarti, Amarnag Subramanya, Michael Ringgaard, and Fernando Pereira. Collective entity resolution with multi-focal attention. *In Proc. of Association for Computational Linguistics, ACL*, 2016.
- [Guo and Barbosa, 2016] Zhaochen Guo and Denilson Barbosa. Robust named entity disambiguation with random walks. *Semantic Web*, 2016.
- [He *et al.*, 2013] Zhengyan He, Shujie Liu, Mu Li, Ming Zhou, Longkai Zhang, and Houfeng Wang. Learning entity representation for entity disambiguation. *Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL*, 2013.
- [Hoffart *et al.*, 2011] Johannes Hoffart, Mohamed Amir Yosef, and Ilaria Bordino. Robust disambiguation of named entities in text. *In Proc. of Empirical Methods in Natural Language Processing, EMNLP*, 2011.
- [Hu *et al.*, 2015] Zhiting Hu, Poyao Huang, Yuntian Deng, Yingkai Gao, and Eric P. Xing. Entity hierarchy embedding. *53rd Annual Meeting of the Association for Computational Linguistics*, 2015.
- [Ji *et al.*, 2010] Heng Ji, Ralph Grishman, Hoa Trang Dang, Kira Griffitt, and Joe Ellis. Overview of the tac 2010 knowledge base population track. *In Proc. of the 3rd Text Analysis Conference, TAC*, 2010.
- [Kingma and Ba, 2014] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014.
- [Le and Titov, 2018] Phong Le and Ivan Titov. Improving entity linking by modeling latent relations between mentions. *56th Annual Meeting of the Association for Computational Linguistics*, 2018.
- [Perschina *et al.*, 2015] Maria Perschina, Yifan He, and Ralph Grishman. Personalized page rank for named entity disambiguation. *Annual Conference of the North American Chapter of the Association for Computational Linguistics, NAACL*, 2015.
- [Peters *et al.*, 2018] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. *Annual Conference of the North American Chapter of the Association for Computational Linguistics, NAACL*, 2018.
- [Shahbazi *et al.*, 2018] Hamed Shahbazi, Xiaoli Z. Fern, Reza Ghaeini, Chao Ma, Rasha Obeidat, and Prasad Tadepalli. Joint neural entity disambiguation with output space search. *The 27th International Conference on Computational Linguistics, COLING*, 2018.
- [Sil and Florian, 2016] Avirup Sil and Radu Florian. Towards language independent named entity linking. *Annual Meeting of the Association for Computational Linguistics: System Demonstrations, ACL*, 2016.
- [Sil *et al.*, 2018] Avirup Sil, Gourab Kundu, Radu Florian, and Wael Hamza. Neural cross-lingual entity linking. *Thirty-Second AAAI Conference on Artificial Intelligence, AAAI*, 2018.
- [Sun *et al.*, 2015] Yaming Sun, Lei Lin, Duyu Tang, Nan Yang, Zhenzhou Ji, and Xiaolong Wang. Modeling mention, context and entity with neural networks for entity disambiguation. *International Joint Conference on Artificial Intelligence, IJCAI*, 2015.
- [Yamada *et al.*, 2016] Ikuya Yamada, Hiroyuki Shindo, Hideaki Takeda, and Yoshiyasu Takefuji. Joint learning of the embedding of words and entities for named entity disambiguation. *International Conference on Computational Linguistics, COLING*, 2016.
- [Yamada *et al.*, 2017] Ikuya Yamada, Hiroyuki Shindo, Hideaki Takeda, and Yoshiyasu Takefuji. Learning distributed representations of texts and entities from knowledge base. *CoRR abs/1705.02494*, 2017.
- [Yamada *et al.*, 2020] Ikuya Yamada, Koki Washio, Hiroyuki Shindo, and Yuji Matsumoto. Global entity disambiguation with pretrained contextualized embeddings of words and entities. *arXiv preprint arXiv:1909.00426*, 2020.