

Improving Entity Linking by Modeling Latent Entity Type Information

Shuang Chen,^{1*} Jinpeng Wang,² Feng Jiang,^{1,3} Chin-Yew Lin²

¹Harbin Institute of Technology, Harbin, China

²Microsoft Research Asia, ³Peng Cheng Laboratory

hitercs@gmail.com, {jinpwa, cyl}@microsoft.com, fjiang@hit.edu.cn

Abstract

Existing state of the art neural entity linking models employ attention-based bag-of-words context model and pre-trained entity embeddings bootstrapped from word embeddings to assess topic level context compatibility. However, the latent entity type information in the immediate context of the mention is neglected, which causes the models often link mentions to incorrect entities with incorrect type. To tackle this problem, we propose to inject latent entity type information into the entity embeddings based on pre-trained BERT. In addition, we integrate a BERT-based entity similarity score into the local context model of a state-of-the-art model to better capture latent entity type information. Our model significantly outperforms the state-of-the-art entity linking models on standard benchmark (AIDA-CoNLL). Detailed experiment analysis demonstrates that our model corrects most of the type errors produced by the direct baseline.

Introduction

Entity Linking (EL) is the task of disambiguating textual mentions to their corresponding entities in a reference knowledge base (e.g., Wikipedia). An accurate entity linking system is crucial for many knowledge related tasks such as question answering (Yih et al. 2015) and information extraction (Hoffmann et al. 2011).

Traditional entity linkers mainly depend on manually designed features to evaluate the local context compatibility and document-level global coherence of referent entities (Cheng and Roth 2013; Durrett and Klein 2014). The design of such features requires entity-specific domain knowledge. These features can not fully capture relevant statistical dependencies and interactions. One recent notable work (Ganea and Hofmann 2017) instead pioneers to rely on pre-trained entity embeddings, learnable context representation and differentiable joint inference stage to learn basic features and their combinations from scratch. Such model design allows to learn useful regularities in an end-to-end fashion and eliminates the need for extensive feature engineering. It also substantially outperforms

*Contribution during internship at Microsoft Research. Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

| Wikipedia Title | Local context score | Golden |
|-------------------|---------------------|--------|
| Milwaukee Brewers | 0.20 | × |
| Milwaukee | 0.11 | √ |

↑
 In Milwaukee, Marc Newfield homered off Jose Parra (5-4)
 (a) leading off the bottom of the 12th as the Brewers rallied for a 5-4 victory over the Minnesota Twins .
 ninth ninth seventh seventh Milwaukee Milwaukee inning
 (b) games games victory Royals Royals win rallied second second
 innings run run run run run Brewers Brewers Minnesota

Figure 1: One error case on AIDA-CoNLL development set of the full model of Ganea and Hofmann (2017). (a) Immediate context; (b) Attended contextual words sorted by attention weights. The preposition “In” is a strong cue predictive of the type of mention “Milwaukee” which is not captured by the local context model.

the traditional methods on standard benchmark (e.g., AIDA-CoNLL). A line of follow-up work (Le and Titov 2018; 2019a; 2019b) investigate potential improvement solution or other task settings based on that.

Such state-of-the-art entity linking models (Ganea and Hofmann 2017; Le and Titov 2018) employ attention-based bag-of-words context model and pre-trained entity embeddings bootstrapped from word embeddings to assess topic level context compatibility. However, the latent entity type information in the immediate context of the mention is neglected. We suspect this may sometimes cause the models link mentions to incorrect entities with incorrect type. To verify this, we conduct error analysis of the well known DeepED¹ model (Ganea and Hofmann 2017) on the development set of AIDA-CoNLL (Hoffart et al. 2011), and found that more than half of their error cases fall into the category of type errors where the predicted entity’s type is different from the golden entity’s type, although some predictive contextual cue for them can be found in their local context. As shown in Fig. 1, the full model of Ganea and Hofmann (2017) incorrectly links the mention “Milwaukee” to the entity MILWAUKEE_BREWERS. However, the prepo-

¹<https://github.com/dalab/deep-ed>

sition “In” is a strong cue predictive of the type (location) of mention “Milwaukee” which is helpful for disambiguation. The reason why the local context model of Ganea and Hofmann (2017) couldn’t capture such apparent cue is two folds. On one hand, the context encoding module adopts a bag-of-words encoding scheme which is position agnostic. As shown in Fig. 1(b), the attention mechanism is helpful for selecting predictive words (e.g. “Milwaukee”, “games” etc.), but does not capture the pattern that the previous word “In” of the mention “Milwaukee” which very likely refers to an entity with location type. On the other hand, the pre-trained entity embedding of Ganea and Hofmann (2017) is not very sensitive to entity types. For example, as shown in Table 8, when we query the most similar entities with the entity STEVE_JOBS, the top one returned entity is APPLE_INC., which is a different type but related at topic level. So it is natural for the model of Ganea and Hofmann (2017) to make type errors when it is trained to fit such entity embeddings.

As is argued in (Zhou et al. 2018), “context consistency is a strong proxy for type compatibility”. Based on this claim, the mention’s immediate context is a proxy of its type. For example, we consider the following context from Wikipedia linking to the entity APPLE in which the mention is replaced with the [MASK] token.

Fruits that tend to be more popular in this area are [MASK], pears, and berries.

By reading the context surrounding the [MASK] token, we can easily determine that the entities fitting this context should be a kind of fruit.

In this paper, we propose to inject latent entity type information into the entity embeddings by modeling the immediate context surrounding the mention. Specifically, we apply pre-trained BERT (Devlin et al. 2019) to represent the entity context and build a shared entity representation by aggregating all the entity contexts linking to the same entity via average pooling. Pre-trained BERT models naturally fit our purpose to represent the entity context surrounding the [MASK] token as it is trained with masked language model objective. What’s more, we integrate a BERT-based entity similarity feature into the local model of Ganea and Hofmann (2017) to better capture entity type information. This can leverage both the pre-trained entity embeddings from BERT and the domain adaption capability of BERT via fine-tuning.

We conduct entity linking experiments on standard benchmark datasets: AIDA-CoNLL and five out-domain test sets. Our model achieves an absolute improvement of 1.32% F1 on AIDA-CoNLL test set and average 0.80% F1 on five out-domain test sets over five different runs. In addition, we conduct detailed experiment analysis on AIDA-CoNLL development set which shows our proposed model can reduce 67.03% type errors of the state-of-the-art model (Ganea and Hofmann 2017) and more than 90% of the remaining type error cases are due to over estimation of prior and global modeling problem which we leave as the further work.

Our contributions can be summarized as follows.

- We show current state-of-the-art (SOTA) neural entity linking models based on attention-based bag-of-words context model often produce type errors and analyze the

possible causes.

- We propose a novel entity embedding method based on pre-trained BERT to better capture latent entity type information.
- We integrate a BERT-based entity similarity into the local model of a SOTA model (Ganea and Hofmann 2017).
- We verify the effectiveness of our model on standard benchmark datasets and achieve significant improvement over the baseline. And the detailed experiment analysis demonstrates that our method truly corrects most of the type errors produced by the baseline.

Background

Entity Linking Problem

Formally, given a document D consisting of a list of entity mentions m_1, \dots, m_n . The goal of an entity linking system is to assign each m_i a KB entity e_i or predict that no corresponding entity in the KB (i.e., $e_i = \text{NIL}$).

Due to potentially very large entity space (e.g. Wikipedia has more than 4 million entities), standard entity linking is often divided into two stages: *candidate generation* which chooses potential candidates $C_i = (e_{i1}, \dots, e_{il_i})$ using a heuristic and *entity disambiguation* which learns to select the best entity from the candidates using a statistical model. In this work, we focus on the second stage *entity disambiguation*. As for entity disambiguation, two different kinds of information can be leveraged: local context compatibility and document-level global coherence which respectively corresponds to the local model and the global model. Next, we introduce the general formulation of entity linking problem with a focus on the well known DeepED model (Ganea and Hofmann 2017).

General Formulation An entity linking model integrating both local and global features can be formulated as a conditional random field. Formally, we can define a scoring function g to evaluate the entity assignment e_1, \dots, e_n to mentions m_1, \dots, m_n in a document D .

$$g(e_1, \dots, e_n | D) = \sum_{i=1}^n \Psi(e_i | D) + \sum_{j \neq i} \Phi(e_i, e_j | D) \quad (1)$$

where the first term scores how well an entity fits its local context and the second one measures the global coherence.

Local Model Following Ganea and Hofmann (2017), we instantiate the local model as an attention model based on pre-trained word and entity embeddings. Specifically, for each mention m_i , a pruned candidate set $C_i = (e_{i1}, \dots, e_{il_i})$ is identified in the candidate generation stage. We compute a local context score for each $e \in C_i$ based on the K -word (in practice, K is set to 100 and stop words are removed.) local context $c = \{w_1, \dots, w_K\}$ surrounding m_i .

$$\Psi_{\text{long}}(e, c) = \mathbf{x}_e^\top \mathbf{B} h(c) \quad (2)$$

where \mathbf{B} is a learnable diagonal matrix, \mathbf{x}_e is the embeddings of entity e , $h(c)$ applies a hard attention mechanism to context words in c to obtain the representation of the context.

Besides, Ganea and Hofmann (2017) combined this context score with the prior $\hat{p}(e|m)$ (computed by mixing mention-entity hyperlink count statistics from Wikipedia, a large Web corpus and YAGO.²) using a two-layer feed-forward neural network in the local model.

$$\Psi(e, m, c) = f(\Psi_{\text{long}}(e, c), \log \hat{p}(e|m)) \quad (3)$$

Global Model The second term in Equation 1 is given by:

$$\Phi(e, e') = \frac{2}{n-1} \mathbf{x}_e^\top \mathbf{C} \mathbf{x}_{e'} \quad (4)$$

where \mathbf{C} is a diagonal matrix. The model defined by Equation 1 is a fully-connected pairwise conditional random field. Exact maximum-a-posteriori inference on this CRF, needed both at training and testing phrase, is NP-hard (Wainwright, Jordan, and others 2008). So they used maximum-product loopy belief propagation (LBP) to estimate the maximum marginal probability

$$\hat{g}_i(e|D) \approx \max_{\substack{e_1, \dots, e_{i-1} \\ e_{i+1}, \dots, e_n}} g(e_1, \dots, e_n|D) \quad (5)$$

for each mention m_i . The final score for m_i is given by:

$$\rho_i(e) = f'(\hat{g}_i(e|D), \hat{p}(e|m_i)) \quad (6)$$

where f' is another two-layer neural network and $\hat{p}(e|m_i)$ is the prior feature.

Related Work

Our work focuses on improving entity linking by capturing latent entity type information with BERT. Specifically, our work related to previous approaches in three aspects.

Entity Embedding The entity linking task is essentially a zero-shot task where the answer of test cases may not exist in the training data.³ So we need to build a shared entity embedding space for all entities which allows neural entity linking models to generalize to both seen and unseen entities during test time. Based on the distributional hypothesis (Harris 1954), an entity is characterized by its contexts. Different methods to characterize an entity’s context result in different information its entity embedding can capture. Previous work (Yamada et al. 2016; Ganea and Hofmann 2017) on learning entity representation are mostly extensions of the embedding methods proposed by (Mikolov et al. 2013). An entity’s context is a bag-of-words representation which mainly captures topic level entity relatedness rather than entity type relatedness. In contrast, we propose a simple method to build entity embeddings directly from pre-trained BERT (Devlin et al. 2019) which can better capture entity type information.

Type Information Previous work attempt to integrate type information into the entity linking task mostly by jointly modeling named entity recognition and entity linking. Specifically, a line of work (Durrett and Klein 2014;

Luo et al. 2015; Nguyen, Theobald, and Weikum 2016) jointly model entity linking and named entity recognition to capture the mutual dependency between them using structured CRF. These methods mainly differ in the design of hand-engineered features. Recently, Martins, Marinho, and Martins (2019) perform multi-task learning using learned features by extending Stack-LSTM (Dyer et al. 2015). However, all of these work rely on extensive annotation of the type of mentions which are difficult to obtain on most of the entity linking datasets. In contrast, based on the assumption that “context consistency is a strong proxy for type compatibility” from Zhou et al. (2018), we propose to model a mention’s immediate context using BERT (Devlin et al. 2019) to capture its contextual latent entity type information.

Applications of BERT Since the advent of the well-known BERT models (Devlin et al. 2019), it has been applied successfully to and has achieved state-of-the-art performance on many NLP tasks. The main challenges which the entity linking task has over other tasks e.g. sentence classification, named entity recognition, where BERT has been applied are: (1) a very large label space, i.e. every mention has many target entities and (2) the zero-shot nature of the entity linking task. Training label embeddings from a small labeled dataset could not generalize to cover unseen entities in test time. To tackle this problem, we introduce a novel method to build entity embeddings from BERT by modeling the immediate context of an entity.

Model

Our model consists of two phrases: (1) Build entity embeddings from BERT (2) Add a BERT-based entity similarity component to the local model. Next we will describe each phrase in the following sections.

Entity Embeddings from BERT

Given lists of mention context⁴ $\{c_{i1}, c_{i2}, \dots, c_{iN}\}$ in Wikipedia for every entity $e_i \in \mathcal{E}$, we build the entity embeddings map $\mathbf{B} : \mathcal{E} \rightarrow \mathbb{R}^d$. Here, the anchor context $c_{ij} = (\text{lctx}_{ij}, m_{ij}, \text{rctx}_{ij})$ where m_{ij} is the mention, lctx_{ij} is the left context and rctx_{ij} is the right context.

Context Representation A mention’s immediate context is a proxy for its type. Here, the mention’s immediate context is a sequence of tokens where the mention m_{ij} is replaced with a single [MASK] token. Then, we represent the immediate entity context by extracting the upper most layer representation of pre-trained BERT (Devlin et al. 2019) corresponding to the [MASK] token.

$$c_{ij} = \text{BERT}(\{\text{lctx}_{ij}, [\text{MASK}], \text{rctx}_{ij}\}) \quad (7)$$

Entity Representation For each entity $e_i \in \mathcal{E}$, we randomly sample at most N anchor contexts $\{c_{i1}, c_{i2}, \dots, c_{iN}\}$ from Wikipedia. Then the entity representation of e_i is computed by aggregating all the context representation

²See Ganea and Hofmann (2017) for more details.

³Only 58.6% answers of test cases in AIDA-CoNLL dataset are existent in the training data.

⁴A mention context of an entity is the surrounding text of an anchor text, i.e. mention, pointing to the entity page in Wikipedia.

$\{c_{i1}, c_{i2}, \dots, c_{iN}\}$ via average pooling.

$$\mathbf{B}_{e_i} = \frac{1}{N} \sum_{j=1}^N \mathbf{c}_{ij} \quad (8)$$

As will be shown in the analysis section, the entity embeddings from BERT better capture entity type information than those from Ganea and Hofmann (2017).

BERT-based Entity Similarity

The local context model of Ganea and Hofmann (2017) mainly captures the topic level entity relatedness information based on a long range bag-of-words context. To capture latent entity type information, we design a BERT-based entity similarity score $\Psi_{\text{BERT}}(e, c)$. Specifically, given a short range (the immediate context where the mention m lies) context $c = (\text{lctx}, m, \text{rtctx})$, we firstly encode c using the same method defined by Equation 7.

$$\mathbf{c} = \text{BERT}(\{\text{lctx}, [\text{MASK}], \text{rtctx}\}) \quad (9)$$

Then we define the BERT-based entity similarity as the cosine similarity⁵ between the context representation \mathbf{c} and the entity representation \mathbf{B}_e .

$$\Psi_{\text{BERT}}(e, c) = \text{cosine}(\mathbf{B}_e, \mathbf{c}) \quad (10)$$

Finally, as for the local disambiguation model, we integrate the BERT-based entity similarity $\Psi_{\text{BERT}}(e, c)$ with the local context score $\Psi_{\text{long}}(e, c)$ (defined in Equation 2) and the prior $\hat{p}(e|m_i)$ with two fully connected layers of 100 hidden units and ReLU non-linearities following the same feature composition methods as Ganea and Hofmann (2017).

$$\Psi_{\text{local}}(e, m, c) = f(\Psi_{\text{long}}(e, c), \Psi_{\text{BERT}}(e, c), \log \hat{p}(e|m)) \quad (11)$$

As for the global disambiguation model, we firstly define the local context score $\Psi_{\text{localctx}}(e, c)$ by combining $\Psi_{\text{long}}(e, c)$ and $\Psi_{\text{BERT}}(e, c)$.⁶

$$\Psi_{\text{localctx}}(e, c) = f(\Psi_{\text{long}}(e, c), \Psi_{\text{BERT}}(e, c)) \quad (12)$$

Then we adopt exactly the same global model as Ganea and Hofmann (2017) which is already introduced in the Background section. Specifically, we adopt loopy belief propagation (LBP) to estimate the max-marginal probability $\hat{g}_i(e|D)$ and then combine it with the prior $\hat{p}(e|m_i)$ using a two-layer neural network to get the final score $\rho_i(e)$ for m_i .

$$\Phi(e, e') = \frac{2}{n-1} \mathbf{x}_e^\top \mathbf{C} \mathbf{x}_{e'} \quad (13)$$

$$\hat{g}_i(e|D) \approx \max_{\substack{e_1, \dots, e_{i-1} \\ e_{i+1}, \dots, e_n}} g(e_1, \dots, e_n|D) \quad (14)$$

$$\rho_i(e) = f'(\hat{g}_i(e|D), \hat{p}(e|m_i)) \quad (15)$$

⁵We also investigated calculating the similarity using a parameterized formula by adding a diagonal matrix between them, but found no significant improvements over Eq. 10.

⁶Following Ganea and Hofmann (2017), we did not integrate the prior score $\hat{p}(e|m)$ into the local scoring module of the global disambiguation model.

Model Training

We minimize the following max-margin ranking loss:

$$L(\theta) = \sum_{D \in \mathcal{D}} \sum_{m_i \in D} \sum_{e \in C_i} h(m_i, e) + \lambda \|\alpha\|_2^2 \quad (16)$$

$$h(m_i, e) = \max(0, \gamma - s(e_i^*) + s(e))$$

$$s(e) = \begin{cases} \Psi_{\text{local}}(e, m, c) & \text{if local model only} \\ \rho_i(e) & \text{local \& global model} \end{cases}$$

In order to discourage the model from biasing toward a particular feature, we add a L2 regularization term ($\lambda \|\alpha\|_2^2$) w.r.t parameters α in feature composition function f to the loss function in Equation 16, where λ is set 10^{-7} .

Experiments

Datasets

In order to verify the effectiveness of our model, we conduct experiments on standard benchmark datasets considering both in-domain and out-domain settings. For in-domain setting, we use AIDA-CoNLL dataset (Hoffart et al. 2011) for training, validation and testing. For out-domain setting, we evaluate the model trained with AIDA-CoNLL on five popular out-domain test sets: MSNBC, AQUAINT, ACE 2004 datasets cleaned and updated by Guo and Barbosa (2016) and WNED-CWEB (CWEB), WNED-WIKI (WIKI) automatically extracted from ClueWeb and Wikipedia (Guo and Barbosa 2016). Following previous work (Ganea and Hofmann 2017), we only consider in-KB mentions. Besides, our candidate generation strategy follows that of Ganea and Hofmann (2017) to make our results comparable.

Setup

The main goal of this work is to introduce a BERT-based entity similarity to capture latent entity type information which is supplementary to existing SOTA local context model (Ganea and Hofmann 2017). So we evaluate the performance when integrating the BERT-based entity similarity into the local context model of Ganea and Hofmann (2017). We also evaluate our model with or without global modeling method of Ganea and Hofmann (2017). In addition, we further compare our methods with other state-of-the-art models (Yamada et al. 2016; Le and Titov 2018). To verify the contribution of our proposed BERT-based entity embeddings, we also compare with a straightforward baseline which directly replaces the encoder of Ganea and Hofmann (2017) utilizing pre-trained BERT. To do so, we introduce a 768×300 dimensional matrix \mathbf{W} which projects BERT-based context representation \mathbf{c} into Ganea and Hofmann (2017)'s entity embeddings space when calculating the similarity score.

Hyper-parameter Setting

The resources (word and entity embeddings) used to train the local context model of Ganea and Hofmann (2017) are obtained from DeepED⁷. For each entity, we randomly sam-

⁷<https://github.com/dalab/deep-ed/>

ple at most 100 anchor contexts from Wikipedia⁸ to build the entity representation from BERT. We discard any articles appearing in WIKI dataset when building the entity representation from BERT. We take the anchor context as the surrounding sentence where the mention lies and replace the mention with a single [MASK] token. Each context is truncated to 128 tokens after WordPiece tokenization. We use the PyTorch implementation of pre-trained BERT models⁹ and choose the BERT-base-cased version. We adopt the Adam (Kingma and Ba 2014) implemented by BERT with $\beta_1 = 0.9$, $\beta_2 = 0.999$. Empirically, we found that it is helpful to set parameters in BERT a small initial learning rate and not BERT related parameters a larger initial learning rate to avoid the whole model biasing toward the BERT feature and disregarding other model components. In our experiments, pre-trained BERT model is fine-tuned with initial learning rate 10^{-5} whereas not BERT related parameters are trained with 10^{-3} . Similar learning rate usage can be found in the recent work by (Hwang et al. 2019). Similar to Ganea and Hofmann (2017), all the entity embeddings are fixed during fine-tuning. We randomly initialize the not BERT related parameters using Gaussian distribution $\mathcal{N}(0.0, 0.02)$ and the bias term is zeroed.

Note that all the hyper-parameters used in the local context and global model of Ganea and Hofmann (2017) were set to the same values as theirs for direct comparison purpose. Detailed hyper-parameters setting is described in the appendices. Our model is trained with 4 NVIDIA Tesla P100 GPUs. We run each of our model five times with different random seeds, and the performance is reported in the form of average \pm standard deviation.

Results

Table 1 shows the micro F1 scores on in-domain AIDA-B dataset of the SOTA methods and ours, which all use Wikipedia and YAGO mention-entity index. The models are divided into two groups: *local models* and *local & global models*. As we can see, our proposed model, BERT-Entity-Sim, outperforms all previous methods. Our local model achieves a 1.31 improvement in terms of F1 over its corresponding baseline (Ganea and Hofmann 2017), yielding a very competitive local model with an average 90.06 F1 score even surpassing the performance of four local & global models. Equipped with the global modeling method of Ganea and Hofmann (2017), the performance of our model further increase to 93.54 with an average 1.32 improvement in terms of F1 over Ganea and Hofmann (2017). In addition, our method outperforms Le and Titov (2018) model by 0.47 point. The model of Le and Titov (2018) is a multi-relational extension of Ganea and Hofmann (2017)’s global modeling method while keeps exactly the same local context model. Our better local context model should be orthogonal with them and has potential more applications on short texts (e.g. tweets) where global modeling has little benefits. Moreover, BERT+G&H’s embeddings performs signif-

⁸Using the same Wikipedia dump (Feb 2014) as the one which Ganea and Hofmann (2017) used to train their entity embeddings.

⁹<https://github.com/huggingface/pytorch-pretrained-BERT>

| Methods | AIDA-B |
|--|------------------------------------|
| <i>Local models</i> | |
| prior $\hat{p}(e m)$ | 71.9 |
| Lazic et al. (2015) | 86.4 |
| Globerson et al. (2016) | 87.9 |
| Yamada et al. (2016) | 87.2 |
| Ganea and Hofmann (2017) | 88.8 |
| Ganea and Hofmann (2017) (reproduce) | 88.75 \pm 0.30 |
| BERT-Entity-Sim (local) | 90.06 \pm 0.22 |
| <i>Local & Global models</i> | |
| Huang, Heck, and Ji (2015) | 86.6 |
| Ganea et al. (2016) | 87.6 |
| Chisholm and Hachey (2015) | 88.7 |
| Guo and Barbosa (2016) | 89.0 |
| Globerson et al. (2016) | 91.0 |
| Yamada et al. (2016) | 91.5 |
| Ganea and Hofmann (2017) | 92.22 \pm 0.14 |
| Le and Titov (2018) | 93.07 \pm 0.27 |
| <i>Explicit entity type injection models in Analysis section</i> | |
| Oracle type (Ultra-fine) | 95.38 \pm 0.07 |
| Oracle type (FIGER) | 96.35 \pm 0.14 |
| Predict type (Ultra-fine) | 91.35 \pm 0.18 |
| Predict type (ZOE) | 91.42 \pm 0.09 |
| BERT+G&H’s embeddings | 91.00 \pm 0.72 |
| BERT-Entity-Sim (local & global) | 93.54 \pm 0.12 |

Table 1: F1 scores on AIDA-B (test set).

icantly worse than the baseline (Ganea and Hofmann 2017) and our proposed BERT-Entity-Sim model. The reason is that BERT-based context representation space and Ganea and Hofmann’s entity embeddings space are heterogeneous. Ganea and Hofmann’s entity embeddings are bootstrapped from word embeddings which mainly capture topic level entity relatedness, while BERT-based context representation is derived from BERT which naturally captures type information. The non-parallel information in both context and entity sides makes it difficult to learn the alignment parameter \mathbf{W} and results in the poor generalization performance.

To evaluate the robustness of our model, Table 2 shows the performance of our method and SOTA methods on five out-domain test sets. On average, our proposed model (BERT-Entity-Sim) outperforms the local & global version of Ganea and Hofmann; Le and Titov (2017; 2018) by an average 0.80 and 0.51 on F1.

Analysis

We conduct experiment analysis to answer the following questions:

- Do the entity embeddings from BERT better capture latent entity type information than that of Ganea and Hofmann (2017)?
- Does the proposed model correct the type errors in the baseline (Ganea and Hofmann 2017)?
- Can straightforward integration of state-of-the-art fine grained entity typing systems improve entity linking performance?
- Can better global model further boost the performance of the proposed model?

| Methods | MSNBC | AQUAINT | ACE2004 | CWEB | WIKI | Avg |
|--|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|--------------|
| prior $\hat{p}(e m)$ | 89.3 | 83.2 | 84.4 | 69.8 | 64.2 | 78.18 |
| Milne and Witten (2008) | 78 | 85 | 81 | 64.1 | 81.7 | 77.96 |
| Hoffart et al. (2011) | 79 | 56 | 80 | 58.6 | 63 | 67.32 |
| Ratinov et al. (2011) | 75 | 83 | 82 | 56.2 | 67.2 | 72.68 |
| Cheng and Roth (2013) | 90 | 90 | 86 | 67.5 | 73.4 | 81.38 |
| Guo and Barbosa (2016) | 92 | 87 | 88 | 77 | 84.5 | 85.70 |
| Ganea and Hofmann (2017) | 93.7 \pm 0.1 | 88.5 \pm 0.4 | 88.5 \pm 0.3 | 77.9 \pm 0.1 | 77.5 \pm 0.1 | 85.22 |
| Le and Titov (2018) | 93.9 \pm 0.2 | 88.3 \pm 0.6 | 89.9 \pm 0.8 | 77.5 \pm 0.1 | 78.0 \pm 0.1 | 85.51 |
| <i>Explicit entity type injection models in Analysis section</i> | | | | | | |
| Oracle type (Ultra-fine) | 96.8 \pm 0.1 | 93.7 \pm 0.2 | 92.0 \pm 0.2 | 85.8 \pm 0.1 | 84.0 \pm 0.1 | 90.46 |
| Oracle type (FIGER) | 97.1 \pm 0.1 | 92.3 \pm 0.2 | 93.1 \pm 0.2 | 84.3 \pm 0.2 | 84.4 \pm 0.1 | 90.24 |
| Predict type (Ultra-fine) | 93.4 \pm 0.2 | 89.7 \pm 0.2 | 89.1 \pm 0.5 | 77.8 \pm 0.1 | 76.8 \pm 0.2 | 85.36 |
| Predict type (ZOE) | 93.2 \pm 0.2 | 89.5 \pm 0.2 | 89.2 \pm 0.4 | 77.6 \pm 0.1 | 77.0 \pm 0.1 | 85.30 |
| BERT+G&H’s embeddings | 93.3 \pm 0.4 | 89.1 \pm 0.6 | 88.1 \pm 0.6 | 75.7 \pm 0.5 | 76.3 \pm 0.5 | 84.50 |
| BERT-Entity-Sim (local & global) | 93.4 \pm 0.1 | <u>89.8 \pm 0.4</u> | 88.9 \pm 0.7 | 77.9 \pm 0.4 | <u>80.1 \pm 0.4</u> | 86.02 |

Table 2: F1 scores on five out-domain test sets. Underlined scores denote the corresponding model outperforms the baseline.

| Entity Embedding | FIGER | | | BBN | | | OntoNotes _{fine} | | |
|-----------------------------|-----------|-----------|-------|-----------|-----------|-------|---------------------------|-----------|-------|
| | $F1_{mi}$ | $F1_{ma}$ | Acc. | $F1_{mi}$ | $F1_{ma}$ | Acc. | $F1_{mi}$ | $F1_{ma}$ | Acc. |
| Ganea and Hofmann (2017) | 80.38 | 82.65 | 53.30 | 80.87 | 84.10 | 69.34 | 81.41 | 83.54 | 57.54 |
| BERT based Entity Embedding | 88.69 | 90.98 | 69.07 | 91.30 | 93.35 | 85.36 | 90.74 | 92.52 | 73.94 |

Table 3: Results of type classification task on three typing systems: FIGER, BBN, OntoNotes_{fine}

| Error Type | # Cases | Percentage (%) |
|-----------------------------|---------|----------------|
| <i>Due to prior</i> | 41 | 67.21 |
| <i>Due to global</i> | 14 | 22.95 |
| <i>Due to local context</i> | 6 | 9.84 |

Table 4: Remaining type error cases categorization

| Typing System | $F1_{mi}$ | $F1_{ma}$ | Acc. |
|--------------------|-----------|-----------|--------|
| Choi et al. (2018) | 26.52% | 26.60% | 0.36% |
| Zhou et al. (2018) | 66.12% | 67.98% | 46.08% |

Table 5: Performance of two state-of-the-art fine grained entity typing systems on AIDA-CoNLL development set

Effectiveness of BERT-based and Ganea & Hofmann (2017) Entity Embedding in Entity Type Prediction

In order to verify our claim that the entity embeddings from BERT better capture entity type information than those from Ganea and Hofmann (2017), we carry out an entity type prediction task based on its entity embedding. Specifically, we randomly sample 100K entities from Wikipedia, and randomly split them into training set (80K), development set (10K) and test set (10K). For each entity, we obtain its entity types from three typing systems: FIGER (Ling and Weld 2012), BBN (Weischedel and Brunstein 2005) and OntoNotes_{fine} (Gillick et al. 2014) via the entity type mapping provided by Zhou et al. (2018). The entity type prediction model is a simple linear classification model¹⁰ using the entity embedding of an entity as features; limiting its capacity enables us to focus on whether type information can be easily extracted from the entity embeddings. We evaluate the

¹⁰Due to space limitation, we put the detailed description of this model and training hyper-parameters in the Appendices.

model using standard entity typing metrics: Strict Accuracy (Acc.), Micro F1 ($F1_{mi}$) and Macro F1 ($F1_{ma}$).

As shown in Table 3, our proposed entity embedding from BERT significantly outperforms the entity embedding proposed by Ganea and Hofmann (2017) on three typing systems FIGER, BBN and OntoNotes_{fine}. Specifically, our method improves over the baseline with an absolute 8.31, 10.43 and 9.33 $F1_{mi}$ point than the baseline on three typing systems respectively. This demonstrates that our proposed entity embeddings from BERT indeed capture better latent entity type information than Ganea and Hofmann (2017).

Type Errors Correction As we have mentioned in the introduction section, more than half of the baseline model’s errors on the AIDA-A dataset are type errors. Type errors are error cases¹¹ where (1) the predicted entity’s type is different from the golden entity’s type; (2) contextual cue predictive of the type of the mention exists; (3) errors are not due to annotation errors. By doing so, we collect 185 type error cases which cover 57.45% of all (322) error cases. This indicates that Ganea and Hofmann (2017) produces many type errors due to its inability to consider the entity type information in mention context. By integrating the BERT-based entity similarity, our proposed model can correct 124 out of 185 (67.03%) type error cases of the baseline model which demonstrates that we correct more than two third of the type errors produced by the baseline. We have further examined and categorized the remaining 61 type error cases into three categories: (i) *Due to prior*: golden entities with very low $\hat{p}(e|m_i)$ prior, (ii) *Due to global*: both the local context score and prior score support predicting the golden

¹¹We discard the error cases due to candidate generation problem (i.e., gold entities that do not appear in mentions’ candidate list) which cover 2.98% mentions of AIDA-A dataset.

| Methods | AIDA-B | MSNBC | AQUAINT | ACE2004 | CWEB | WIKI | Avg |
|--------------------------------------|---------------------|-------------------|-------------------|-------------------|-------------------|-------------------|--------------|
| Ganea and Hofmann (2017) | 92.22 ± 0.14 | 93.7 ± 0.1 | 88.5 ± 0.4 | 88.5 ± 0.3 | 77.9 ± 0.1 | 77.5 ± 0.1 | 85.22 |
| Le and Titov (2018) | 93.07 ± 0.27 | 93.9 ± 0.2 | 88.3 ± 0.6 | 89.9 ± 0.8 | 77.5 ± 0.1 | 78.0 ± 0.1 | 85.51 |
| Yang et al. (2019) | 94.64 ± 0.20 | 94.6 ± 0.2 | 87.4 ± 0.5 | 89.4 ± 0.4 | 73.5 ± 0.1 | 78.2 ± 0.1 | 84.62 |
| BERT-Entity-Sim (local & global) | 93.54 ± 0.12 | 93.4 ± 0.1 | 89.8 ± 0.4 | 88.9 ± 0.7 | 77.9 ± 0.4 | 80.1 ± 0.4 | 86.02 |
| BERT-Entity-Sim (local & DCA global) | 93.66 ± 0.17 | 94.5 ± 0.3 | 89.1 ± 0.3 | 90.8 ± 0.4 | 78.2 ± 0.2 | 81.0 ± 0.3 | 86.72 |

Table 6: F1 scores of BERT-Entity-Sim equipped with the DCA global model (Yang et al. 2019) on six test sets.

entity, but the overall score supports predicting incorrect entity due to global modeling, (iii) *Due to local context*: the local context score misleads the model predicting the wrong entity, this is potentially due to the mention context can be misleading, e.g. a document discussing cricket will favor resolving the mention “Australian” in context “impressed by the positive influence of *Australian* coach Dave Gilbert” to the entity AUSTRALIA_NATIONAL_CRICKET_TEAM instead of the gold entity AUSTRALIA.

As shown in Table 4, 67.21% of the remaining type error cases are due to *prior* problem which are hard to solve in the current feature combination framework. We argue that prior should be considered as the final resort, only relying on it when the model can not make decision based on other features. Besides, there are 22.95% remaining type errors which are due to *global modeling* problem which shows the limitation of the global modeling method of Ganea and Hofmann (2017). Finally, 9.84% type error cases are due to *local context* problem that our BERT-based solution cannot address. We leave this to future work.

Incorporating Explicit Entity Types We have shown that our BERT-based local context model which implicitly captures entity type information and is effective in correcting two third of type error cases. It is nature to conjecture that we can also correct type errors by incorporating explicit type information into Ganea and Hofmann (2017). We investigate this approach in this section. Assuming that we have types for each mention and candidate entity, we calculate the Jaccard similarity between them and use it as a feature for local disambiguation model.

$$\text{JaccardSim}(e, m, c) = \frac{|T_m \cap T_e|}{|T_m \cup T_e|} \quad (17)$$

where T_m and T_e are the type sets of the mention m and candidate entity e respectively. The new local context score function considering explicit type information is defined as:

$$\Psi'_{\text{localctx}}(e, m, c) = f(\Psi_{\text{long}}(e, c), \text{JaccardSim}(e, m, c))$$

We consider both **Oracle** setting and **Predict** setting. In the oracle setting, the mention’s types are set as the golden entity’s types.¹² As for the entity’s types, we use two sources: one is the ultra-fine type sets from Choi et al. (2018) consisting of more than 10,000 ultra-fine grained types; the other one is the FIGER type sets (Ling and Weld 2012) consisting of 112 fine grained types. In the predict setting, we use two state-of-the-art fine grained entity typing systems: 1) Ultra-fine (Choi et al. 2018) which predicts types in ultra-fine type

¹²In practice, this setting is unachievable due to potentially insufficient context and the imperfect entity typing system.

sets; 2) ZOE (Zhou et al. 2018) which can predict types in FIGER type sets.

As we can see from both Table 1 and Table 2, in the oracle setting, the best model outperforms all the state-of-the-art entity linking models by a large margin, even surpass Le and Titov (2018) by 3.28 F1 points on AIDA-CoNLL test set. This result shows that a better type prediction system can further improve upon the state-of-the-state entity linking systems. However, in the predict setting, the type injection models have worse performance than the baseline. The degradation might be attributed to the poor performance of the two state-of-the-art fine grained entity typing systems. To verify this, we measure the performance of the two typing systems on AIDA-CoNLL development set.¹³ As shown in Table 5, the ultra-fine entity typing system (Choi et al. 2018) only achieves 26.52% $F1_{mi}$ score while the ZOE system (Zhou et al. 2018) achieves 66.12% $F1_{mi}$ score¹⁴ which are insufficient to improve state-of-the-art entity linking system with more than 92% F1 score.

Better Global Model In order to investigate whether better global model can further boost the performance of our model, we incorporate the recent proposed Dynamic Context Augmentation (DCA)¹⁵ (Yang et al. 2019). DCA is a global entity linking model featuring better efficiency and effectiveness than that of Ganea and Hofmann (2017) by breaking the “all-mention coherence” assumption. Compared to BERT-Entity-Sim equipped with Ganea and Hofmann (2017)’s global model in Table 8, BERT-Entity-Sim (local & DCA global) model gains a further improvement of 0.12 F1 on AIDA-B and 0.70 F1 on five out-domain test sets on average. This shows that better global model indeed can further boost the performance of our proposed model.

Notice that Yang et al. (2019) achieves a high performance on in-domain AIDA-CoNLL test set, but doesn’t perform well on five out-domain test sets (even inferior than Ganea and Hofmann (2017)). We found that Yang et al. (2019) includes an explicit type similarity which is based on a typing system¹⁶ trained with AIDA-train NER annotation. This explicit type similarity feature is tailored for AIDA-CoNLL data set and doesn’t achieve good generalization performance on out-domain test sets. In contrast, our BERT-Entity-Sim model capturing latent type information has potential better generalization performance with an average 2.10 F1 improvement over them.

¹³The mention’s golden types are set as its linked entity’s types.

¹⁴The size of type sets of ultra-fine typing system is much larger than that of ZOE.

¹⁵<https://github.com/YoungXiyuan/DCA/>

¹⁶It yields 95% accuracy on AIDA-A according to their paper.

| Model | Context | Context Sim | Golden Entity |
|----------|---|-------------|-----------------------|
| Our | In Milwaukee , Marc Newfield homered off Jose Parra (5-4) ... | - | MILWAUKEE |
| | (1) In Cleveland , Kevin Seitzer ’s two-out single ... | 0.968 | CLEVELAND |
| | (2) In Boston , Troy O’Leary homered off the right-field ... | 0.956 | BOSTON |
| | (3) In Houston , Jeff Bagwell homered and Donne Wall ... | 0.951 | HOUSTON |
| | (4) In Los Angeles , Greg Gagne had a run-scoring single ... | 0.949 | LOS_ANGELES |
| | (5) In Houston , Andy Benes allowed two runs over seven innings ... | 0.947 | HOUSTON |
| Baseline | ninth ninth seventh seventh Milwaukee Milwaukee inning games games victory ... | - | MILWAUKEE |
| | (1) eighth ninth Milwaukee inning league runs victory fourth rallied earned ... | 0.941 | MILWAUKEE_BREWERS |
| | (2) Denny runs fifth fifth allowed game fourth San win win ... | 0.940 | PHILADELPHIA_PHILLIES |
| | (3) Royals league games runs game won won win Minnesota straight ... | 0.938 | MINNESOTA_TWINS |
| | (4) Cleveland fifth games sixth inning innings Friday extra Sox month ... | 0.934 | CLEVELAND_INDIANS |
| | (5) games streak stay Reynoso runs runs second run won fourth ... | 0.926 | MIAMI_MARLINS |

Table 7: Nearest contexts for the example in Fig. 1 in BERT’s and baseline’s context representation space

| Model | STEVE_JOBS | NATIONAL_BASKETBALL_ASSOCIATION | BEIJING |
|-----------------------------|---------------|--|-----------|
| Ganea and Hofmann (2017) | APPLE_INC. | SACRAMENTO_KINGS | SEOUL |
| | STEVE_WOZNIAK | GOLDEN_STATE_WARRIORS | SHANGHAI |
| | BILL_GATES | LOS_ANGELES_CLIPPERS | CHINA |
| BERT based Entity Embedding | STEVE_WOZNIAK | AMERICAN_BASKETBALL_ASSOCIATION | GUANGZHOU |
| | BILL_GATES | WOMEN’S_NATIONAL_BASKETBALL_ASSOCIATION | SHANGHAI |
| | STEVE_BALLMER | NATIONAL_BASKETBALL_LEAGUE_(UNITED_STATES) | NANJING |

Table 8: Examples of nearest entities in Ganea and Hofmann (2017) and BERT based entity representation space

Case Study

We demonstrate the effectiveness of our proposed model by retrieving the nearest neighbours in both the context representation space and entity representation space.

Nearest Contexts We follow Papernot and McDaniel (2018) to retrieve training examples using nearest neighbour in the context representation space. For our model, we use the context representation c defined by Equation 9. For the baseline model, we use the attention-based context representation $h(c)$ defined in Equation 2. This can reveal which training instances support the prediction of a model. As shown in Table 7, for the example in Figure 1, the most similar contexts retrieved by our model’s context representation are all with preposition “In” ahead of the mention and the golden entities of them are all American cities. In contrast, the baseline’s local context is a bag-of-words representation which we denote using the top 10 attended contextual words sorted by attention weights. The most similar contexts retrieved by baseline’s context representation share common words like “games”, “victory” and the golden entities of them are all baseball teams. This explains why the baseline model incorrectly links the mention “Milwaukee” to MILWAUKEE_BREWERS while our model can link to the correct entity MILWAUKEE.

Nearest Entities We also retrieve nearest entities in the embedding space of Ganea and Hofmann (2017) and ours. As we can see, we query STEVE_JOBS, the nearest entity in Ganea and Hofmann (2017) is APPLE_INC. which is a different type. In contrast, all the entities retrieved by our approach share the same types like person, entrepreneur etc. Another example is when we query NATIONAL_BASKETBALL_ASSOCIATION, the most similar entities in Ganea and Hofmann (2017) are NBA teams which

are typically related, while the entities retrieved by our approach are all basketball leagues.

Conclusion

In this paper, we propose to improve entity linking by capturing latent entity type information with BERT. Firstly, we build entity embeddings from BERT by averaging all the context representation extracted from pre-trained BERT. Then we integrate a BERT-based entity similarity into the local model of the state-of-the-art method by (Ganea and Hofmann 2017). The experiment results show that our model significantly outperforms the baseline with an absolute improvement of 1.32% F1 on in-domain AIDA-CoNLL test set and average 0.80% F1 on five out-domain test datasets. The detailed experiment analysis shows that our method corrects most of the type errors produced by the baseline. In the future, we would like to design global modeling methods which can take advantage of the BERT architecture and investigate other ways to use the prior feature.

Acknowledgments

This work is partly funded by National Key Research and Development Program of China via grant 2018YFC0806800 and 2018YFC0832105.

Appendices

Classification Model of Entity Prediction Task

Given an entity e , we firstly retrieve its entity embedding e , then compute the probability for each type in the typeset T :

$$p_j^e = \sigma(\mathbf{w}_j^\top \mathbf{e} + b) \quad (18)$$

where σ is the sigmoid function, \mathbf{w}_j and b are respectively the weight and bias parameter. For each entity e , it is labeled

with t^e , a binary vector of all types where $t_j^e = 1$ if the j^{th} type is in the set of gold types of e and 0 otherwise. We optimize a multi-label binary cross entropy objective:

$$L_{\text{type}} = - \sum_j t_j^e \log p_j^e + (1 - t_j^e) \log(1 - p_j^e) \quad (19)$$

We optimize the model with Adam with an initial learning rate of 1e-3. Each model is trained for up to 200 epoches and training stops when the performance on the development set does not improve for 6 consecutive epoches.

Detailed Hyper-parameters Setting (Table 9)

| Hyper-parameters | Value |
|----------------------------------|-------------------------------------|
| BERT-based entity embedding dims | 768 |
| dumping factor | 0.5 |
| number of LBP loops | 10 |
| batch size | 1 document (≤ 64 mentions) |
| γ (margin) | 0.01 |
| epoch (local model) | 2 |
| epoch (local & global model) | 10 |

Table 9: Values of hyper-parameters.

References

- Cheng, X., and Roth, D. 2013. Relational inference for wikification. In *EMNLP*, 1787–1796.
- Chisholm, A., and Hachey, B. 2015. Entity disambiguation with web links. *TACL* 3.
- Choi, E.; Levy, O.; Choi, Y.; and Zettlemoyer, L. 2018. Ultra-fine entity typing. In *ACL*, 87–96.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- Durrett, G., and Klein, D. 2014. A joint model for entity analysis: Coreference, typing, and linking. *TACL*.
- Dyer, C.; Ballesteros, M.; Ling, W.; Matthews, A.; and Smith, N. A. 2015. Transition-based dependency parsing with stack long short-term memory. In *ACL*, 334–343.
- Ganea, O.-E., and Hofmann, T. 2017. Deep joint entity disambiguation with local neural attention. In *EMNLP*.
- Ganea, O.-E.; Ganea, M.; Lucchi, A.; Eickhoff, C.; and Hofmann, T. 2016. Probabilistic bag-of-hyperlinks model for entity linking. In *WWW*.
- Gillick, D.; Lazic, N.; Ganchev, K.; Kirchner, J.; and Huynh, D. 2014. Context-dependent fine-grained entity type tagging. *arXiv preprint arXiv:1412.1820*.
- Globerson, A.; Lazic, N.; Chakrabarti, S.; Subramanya, A.; Ringgaard, M.; and Pereira, F. 2016. Collective entity resolution with multi-focal attention. In *ACL*, volume 1.
- Guo, Z., and Barbosa, D. 2016. Robust named entity disambiguation with random walks. *Semantic Web, (Preprint)*.
- Harris, Z. S. 1954. Distributional structure. *Word* 10(2-3).
- Hoffart, J.; Yosef, M. A.; Bordino, I.; Fürstenaue, H.; Pinkal, M.; Spaniol, M.; Taneva, B.; Thater, S.; and Weikum, G. 2011. Robust disambiguation of named entities in text. In *EMNLP*.
- Hoffmann, R.; Zhang, C.; Ling, X.; Zettlemoyer, L.; and Weld, D. S. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *ACL*.
- Huang, H.; Heck, L.; and Ji, H. 2015. Leveraging deep neural networks and knowledge graphs for entity disambiguation. *arXiv preprint arXiv:1504.07678*.
- Hwang, W.; Yim, J.; Park, S.; and Seo, M. 2019. A comprehensive exploration on wikisql with table-aware word contextualization. *arXiv preprint arXiv:1902.01069*.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lazic, N.; Subramanya, A.; Ringgaard, M.; and Pereira, F. 2015. Plato: A selective context model for entity resolution. *TACL*.
- Le, P., and Titov, I. 2018. Improving entity linking by modeling latent relations between mentions. In *ACL*, 1595–1604.
- Le, P., and Titov, I. 2019a. Boosting entity linking performance by leveraging unlabeled documents. In *ACL*.
- Le, P., and Titov, I. 2019b. Distant learning for entity linking with automatic noise detection. In *ACL*.
- Ling, X., and Weld, D. S. 2012. Fine-grained entity recognition. In *AAAI*.
- Luo, G.; Huang, X.; Lin, C.-Y.; and Nie, Z. 2015. Joint entity recognition and disambiguation. In *EMNLP*, 879–888.
- Martins, P. H.; Marinho, Z.; and Martins, A. F. 2019. Joint learning of named entity recognition and entity linking. In *ACL: Student Research Workshop*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, 3111–3119.
- Milne, D., and Witten, I. H. 2008. Learning to link with wikipedia. In *CIKM*. ACM.
- Nguyen, D. B.; Theobald, M.; and Weikum, G. 2016. J-nerd: joint named entity recognition and disambiguation with rich linguistic features. *TACL* 4:215–229.
- Papernot, N., and McDaniel, P. 2018. Deep k-nearest neighbors: Towards confident, interpretable and robust deep learning. *arXiv preprint arXiv:1803.04765*.
- Ratinov, L.; Roth, D.; Downey, D.; and Anderson, M. 2011. Local and global algorithms for disambiguation to Wikipedia. In *ACL*.
- Wainwright, M. J.; Jordan, M. I.; et al. 2008. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning* 1(1–2):1–305.
- Weischedel, R., and Brunstein, A. 2005. Bbn pronoun coreference and entity type corpus. *Linguistic Data Consortium*.
- Yamada, I.; Shindo, H.; Takeda, H.; and Takefuji, Y. 2016. Joint learning of the embedding of words and entities for named entity disambiguation. *CoNLL*.
- Yang, X.; Gu, X.; Lin, S.; Tang, S.; Zhuang, Y.; Wu, F.; Chen, Z.; Hu, G.; and Ren, X. 2019. Learning dynamic context augmentation for global entity linking. In *EMNLP*.
- Yih, W.-t.; Chang, M.-W.; He, X.; and Gao, J. 2015. Semantic parsing via staged query graph generation: Question answering with knowledge base. In *ACL*.
- Zhou, B.; Khashabi, D.; Tsai, C.-T.; and Roth, D. 2018. Zero-shot open entity typing as type-compatible grounding. In *EMNLP*.