

End-to-End Neural Entity Linking

Nikolaos Kolitsas *

ETH Zürich

nikos_kolitsas@hotmail.com

Octavian-Eugen Ganea *

ETH Zürich

octavian.ganea@inf.ethz.ch

Thomas Hofmann

ETH Zürich

thomas.hofmann@inf.ethz.ch

Abstract

Entity Linking (EL) is an essential task for semantic text understanding and information extraction. Popular methods separately address the Mention Detection (MD) and Entity Disambiguation (ED) stages of EL, without leveraging their mutual dependency. We here propose the first neural end-to-end EL system that jointly discovers and links entities in a text document. The main idea is to consider all possible spans as potential mentions and learn contextual similarity scores over their entity candidates that are useful for both MD and ED decisions. Key components are context-aware mention embeddings, entity embeddings and a probabilistic mention - entity map, without demanding other engineered features. Empirically, we show that our end-to-end method significantly outperforms popular systems on the Gerbil platform when enough training data is available. Conversely, if testing datasets follow different annotation conventions compared to the training set (e.g. queries/ tweets vs news documents), our ED model coupled with a traditional NER system offers the best or second best EL accuracy.

1 Introduction and Motivation

Towards the goal of automatic text understanding, machine learning models are expected to accurately extract potentially ambiguous mentions of entities from a textual document and link them to a knowledge base (KB), e.g. Wikipedia or Freebase. Known as entity linking, this problem is an essential building block for various Natural Language Processing tasks, e.g. automatic KB construction, question-answering, text summarization, or relation extraction.

An EL system typically performs two tasks: i) Mention Detection (MD) or Named Entity Recognition (NER) when restricted to named entities –

*Equal contribution.

1) MD may split a larger span into two mentions of less informative entities:

B. Obama's wife gave a speech [...]

Federer's coach [...]

2) MD may split a larger span into two mentions of incorrect entities:

Obama Castle was built in 1601 in Japan.

The Kennel Club is UK's official kennel club.

A bird dog is a type of gun dog or hunting dog.

Romeo and Juliet by Shakespeare [...]

Natural killer cells are a type of lymphocyte

Mary and Max, the 2009 movie [...]

3) MD may choose a shorter span, referring to an incorrect entity:

The Apple is played again in cinemas.

The New York Times is a popular newspaper.

4) MD may choose a longer span, referring to an incorrect entity:

Babies Romeo and Juliet were born hours apart.

Table 1: Examples where MD may benefit from ED and viceversa. Each wrong MD decision (underlined) can be avoided by proper context understanding. The correct spans are shown in blue.

– extracts entity references in a raw textual input, and ii) Entity Disambiguation (ED) – links these spans to their corresponding entities in a KB. Until recently, the common approach of popular systems Ceccarelli et al. (2013); van Erp et al. (2013); Piccinno and Ferragina (2014); Daiber et al. (2013); Hoffart et al. (2011); Steinmetz and Sack (2013) was to solve these two sub-problems independently. However, the important dependency between the two steps is ignored and errors caused by MD/NER will propagate to ED without possibility of recovery Sil and Yates (2013); Luo et al. (2015). We here advocate for models that address the end-to-

end EL task, informally arguing that humans understand and generate text in a similar joint manner, discussing about entities which are gradually introduced, referenced under multiple names and evolving during time Ji et al. (2017). Further, we emphasize the importance of the mutual dependency between MD and ED. First, numerous and more informative linkable spans found by MD obviously offer more contextual cues for ED. Second, finding the true entities appearing in a specific context encourages better mention boundaries, especially for multi-word mentions. For example, in the first sentence of Table 1, understanding the presence of the entity *Michelle Obama* helps detecting its true mention "*B. Obama's wife*", as opposed to separately linking *B. Obama* and *wife* to less informative concepts.

We propose a simple, yet competitive, model for end-to-end EL. Getting inspiration from the recent works of Lee et al. (2017) and Ganea and Hofmann (2017), our model first generates all possible spans (mentions) that have at least one possible entity candidate. Then, each mention - candidate pair receives a context-aware compatibility score based on word and entity embeddings coupled with a neural attention and a global voting mechanisms. During training, we enforce the scores of gold entity - mention pairs to be higher than all possible scores of incorrect candidates or invalid mentions, thus jointly taking the ED and MD decisions.

Our contributions are:

- We address the end-to-end EL task using a simple model that conditions the "linkable" quality of a mention to the strongest context support of its best entity candidate. We do not require expensive manually annotated negative examples of non-linkable mentions. Moreover, we are able to train competitive models using little and only partially annotated documents (with named entities only such as the CoNLL-AIDA dataset).
- We are among the first to show that, with one single exception, engineered features can be fully replaced by neural embeddings automatically learned for the joint MD & ED task.
- On the Gerbil¹ benchmarking platform, we empirically show significant gains for the end-to-end EL task when test and training data

¹<http://gerbil.aksw.org/gerbil/>

come from the same domain. Moreover, when testing datasets follow different annotation schemes or exhibit different statistics, our method is still effective in achieving state-of-the-art or close performance, but needs to be coupled with a popular NER system.

2 Related Work

With few exceptions, MD/NER and ED are treated separately in the vast EL literature.

Traditional NER models usually view the problem as a word sequence labeling that is modeled using conditional random fields on top of engineered features Finkel et al. (2005) or, more recently, using bi-LSTMs architectures Lample et al. (2016); Chiu and Nichols (2016); Liu et al. (2017) capable of learning complex lexical and syntactic features.

In the context of ED, recent neural methods He et al. (2013); Sun et al. (2015); Yamada et al. (2016); Ganea and Hofmann (2017); Le and Titov (2018); Yang et al. (2018); Radhakrishnan et al. (2018) have established state-of-the-art results, outperforming engineered features based models. Context aware word, span and entity embeddings, together with neural similarity functions, are essential in these frameworks.

End-to-end EL is the realistic task and ultimate goal, but challenges in joint NER/MD and ED modeling arise from their different nature. Few previous methods tackle the joint task, where errors in one stage can be recovered by the next stage. One of the first attempts, Sil and Yates (2013) use a popular NER model to over-generate mentions and let the linking step to take the final decisions. However, their method is limited by the dependence on a good mention spotter and by the usage of hand-engineered features. It is also unclear how linking can improve their MD phase. Later, Luo et al. (2015) presented one of the most competitive joint MD and ED models leveraging semi-Conditional Random Fields (semi-CRF). However, there are several weaknesses in this work. First, the mutual task dependency is weak, being captured only by type-category correlation features. The other engineered features used in their model are either NER or ED specific. Second, while their probabilistic graphical model allows for tractable learning and inference, it suffers from high computational complexity caused by the usage of the cartesian product of all possible document span segmentations, NER categories and entity assignments. Another

approach is J-NERD Nguyen et al. (2016) that addresses the end-to-end task using only engineered features and a probabilistic graphical model on top of sentence parse trees.

3 Neural Joint Mention Detection and Entity Disambiguation

We formally introduce the tasks of interest. For EL, the input is a text document (or a query or tweet) given as a sequence $D = \{w_1, \dots, w_n\}$ of words from a dictionary, $w_k \in \mathcal{W}$. The output of an EL model is a list of mention - entity pairs $\{(m_i, e_i)\}_{i \in \overline{1, T}}$, where each mention is a word subsequence of the input document, $m = w_q, \dots, w_r$, and each entity is an entry in a knowledge base KB (e.g. Wikipedia), $e \in \mathcal{E}$. For the ED task, the list of entity mentions $\{m_i\}_{i \in \overline{1, T}}$ that need to be disambiguated is additionally provided as input. The expected output is a list of corresponding annotations $\{e_i\}_{i \in \overline{1, T}} \in \mathcal{E}^T$.

Note that, in this work, we only link mentions that have a valid gold KB entity, setting referred in Röder et al. (2017) as *InKB* evaluation. Thus, we treat mentions referring to entities outside of the KB as "non-linkable". This is in line with few previous models, e.g. Luo et al. (2015); Ganea and Hofmann (2017); Yamada et al. (2016). We leave the interesting setting of discovering out-of-KB entities as future work.

We now describe the components of our neural end-to-end EL model, depicted in Figure 1. We aim for simplicity, but competitive accuracy.

Word and Char Embeddings. We use pre-trained Word2Vec vectors Mikolov et al. (2013). In addition, we train character embeddings that capture important word lexical information. Following Lample et al. (2016), for each word independently, we use bidirectional LSTMs Hochreiter and Schmidhuber (1997) on top of learnable char embeddings. These character LSTMs do not extend beyond single word boundaries, but they share the same parameters. Formally, let $\{z_1, \dots, z_L\}$ be the character vectors of word w . We use the forward and backward LSTMs formulations defined recursively as in Lample et al. (2016):

$$\begin{aligned} h_t^f &= FWD - LSTM(h_{t-1}^f, z_t) \\ h_t^b &= BKWD - LSTM(h_{t+1}^b, z_t) \end{aligned} \quad (1)$$

Then, we form the character embedding of w is $[h_L^f; h_1^b]$ from the hidden state of the forward LSTM

corresponding to the last character concatenated with the hidden state of the backward LSTM corresponding to the first character. This is then concatenated with the pre-trained word embedding, forming the context-independent word-character embedding of w . We denote the sequence of these vectors as $\{v_k\}_{k \in \overline{1, n}}$ and depict it as the first neural layer in Figure 1.

Mention Representation. We find it crucial to make word embeddings aware of their local context, thus being informative for both mention boundary detection and entity disambiguation (leveraging contextual cues, e.g. "newspaper"). We thus encode context information into words using a bi-LSTM layer on top of the word-character embeddings $\{v_k\}_{k \in \overline{1, n}}$. The hidden states of forward and backward LSTMs corresponding to each word are then concatenated into *context-aware word embeddings*, whose sequence is denoted as $\{x_k\}_{k \in \overline{1, n}}$.

Next, for each possible mention, we produce a fixed size representation inspired by Lee et al. (2017). Given a mention $m = w_q, \dots, w_r$, we first concatenate the embeddings of the first, last and the "soft head" words of the mention:

$$g^m = [x_q; x_r; \hat{x}^m] \quad (2)$$

The soft head embedding \hat{x}^m is built using an attention mechanism on top of the mention's word embeddings, similar with Lee et al. (2017):

$$\begin{aligned} \alpha_k &= \langle \mathbf{w}_\alpha, x_k \rangle \\ a_k^m &= \frac{\exp(\alpha_k)}{\sum_{t=q}^r \exp(\alpha_t)} \\ \hat{x}^m &= \sum_{k=q}^r a_k^m \cdot v_k \end{aligned} \quad (3)$$

However, we found the soft head embedding to only marginally improve results, probably due to the fact that most mentions are at most 2 words long. To learn non-linear interactions between the component word vectors, we project g^m to a final mention representation with the same size as entity embeddings (see below) using a shallow feed-forward neural network FFNN (a simple projection layer):

$$x^m = \text{FFNN}_1(g^m) \quad (4)$$

Entity Embeddings. We use fixed continuous entity representations, namely the pre-trained entity embeddings of Ganea and Hofmann (2017),

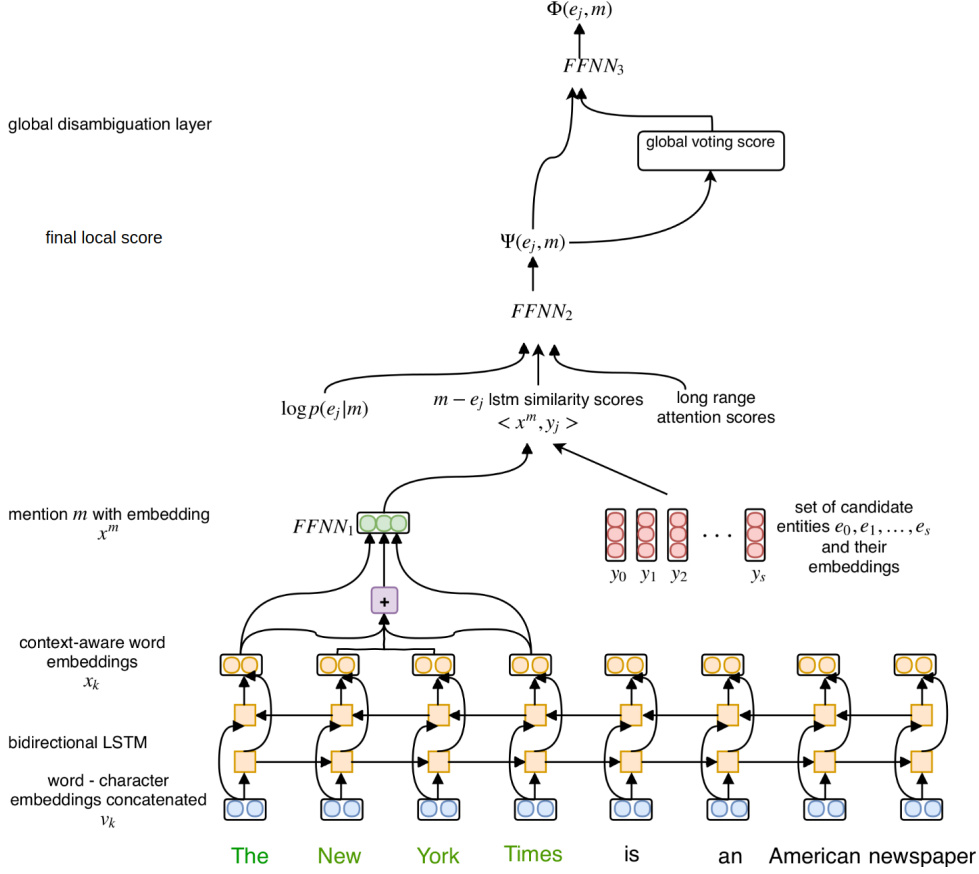


Figure 1: Our global model architecture shown for the mention *The New York Times*. The final score is used for both the mention linking and entity disambiguation decisions.

due to their simplicity and compatibility with the pre-trained word vectors of Mikolov et al. (2013). Briefly, these vectors are computed for each entity in isolation using the following exponential model that approximates the empirical conditional word-entity distribution $\hat{p}(w|e)$ obtained from co-occurrence counts.

$$\frac{\exp(\langle x_w, y_e \rangle)}{\sum_{w' \in \mathcal{W}} \exp(\langle x_{w'}, y_e \rangle)} \approx \hat{p}(w|e) \quad (5)$$

Here, x_w are fixed pre-trained word vectors and y_e is the entity embedding to be trained. In practice, Ganea and Hofmann (2017) re-write this as a max-margin objective loss.

Candidate Selection. For each span m we select up to s entity candidates that might be referred by this mention. These are top entities based on an empirical probabilistic entity - map $p(e|m)$ built by Ganea and Hofmann (2017) from Wikipedia hyperlinks, Crosswikis Spitzkovsky and Chang and YAGO Hoffart et al. (2011) dictionaries. We denote by $C(m)$ this candidate set and use it both at training and test time.

Final Local Score. For each span m that can possibly refer to an entity (i.e. $|C(m)| \geq 1$) and for each of its entity candidates $e_j \in C(m)$, we compute a similarity score using embedding dot-product that supposedly should capture useful information for both MD and ED decisions. We then combine it with the log-prior probability using a shallow FFNN, giving the context-aware entity - mention score:

$$\Psi(e_j, m) = \text{FFNN}_2([\log p(e_j|m); \langle x^m, y_j \rangle]) \quad (6)$$

Long Range Context Attention. In some cases, our model might be improved by explicitly capturing long context dependencies. To test this, we experimented with the attention model of Ganea and Hofmann (2017). This gives one context embedding per mention based on informative context words that are related to at least one of the candidate entities. We use this additional context embedding for computing dot-product similarity with any of the candidate entity embeddings. This value is fed as additional input of FFNN_2 in Eq. 6. We refer to this model as *long range context attention*.

Training. We assume a corpus with documents and gold entity - mention pairs $\mathcal{G} = \{(m_i, e_i^*)\}_{i=1, \dots, K}$ is available. At training time, for each input document we collect the set M of all (potentially overlapping) token spans m for which $|C(m)| \geq 1$. We then train the parameters of our model using the following minimization procedure:

$$\theta^* = \arg \min_{\theta} \sum_{m \in M} \sum_{e \in C(m)} V(\Psi_{\theta}(e, m)) \quad (7)$$

where the violation term V enforces the scores of gold pairs to be linearly separable from scores of negative pairs, i.e.

$$V(\Psi(e, m)) = \begin{cases} \max(0, \gamma - \Psi(e, m)), & \text{if } (e, m) \in \mathcal{G} \\ \max(0, \Psi(e, m)), & \text{otherwise} \end{cases} \quad (8)$$

Note that, in the absence of annotated negative examples of "non-linkable" mentions, we assume that all spans in M and their candidates that do not appear in \mathcal{G} should not be linked. The model will be enforced to only output negative scores for all entity candidates of such mentions.

We call *all spans training* the above setting. Our method can also be used to perform ED only, in which case we train only on gold mentions, i.e. $M = \{m | m \in \mathcal{G}\}$. This is referred as *gold spans training*.

Inference. At test time, our method can be applied for both *EL* and *ED only* as follows. First, for each document in our validation or test sets, we select all possibly linkable token spans, i.e. $M = \{m | |C(m)| \geq 1\}$ for EL, or the input set of mentions $M = \{m | m \in \mathcal{G}\}$ for ED, respectively. Second, the best linking threshold δ is found on the validation set such that the micro F1 metric is maximized when only linking mention - entity pairs with Ψ score greater than δ . At test time, only entity - mention pairs with a score higher than δ are kept and sorted according to their Ψ scores; the final annotations are greedily produced based on this set such that only spans not overlapping with previously selected spans (of higher scores) are chosen.

Global Disambiguation Our current model is "local", i.e. performs disambiguation of each candidate span independently. To enhance it, we add an extra layer to our neural network that will promote coherence among linked and disambiguated

entities inside the same document, i.e. the *global disambiguation layer*. Specifically, we compute a "global" mention-entity score based on which we produce the final annotations. We first define the set of mention-entity pairs that are allowed to participate in the global disambiguation voting, namely those that already have a high local score:

$$V_G = \{(m, e) | m \in M, e \in C(m), \Psi(e, m) \geq \gamma'\}$$

Since we initially consider all possible spans and for each span up to s candidate entities, this filtering step is important to avoid both undesired noise and exponential complexity for the EL task for which M is typically much bigger than for ED. The final "global" score $G(e_j, m)$ for entity candidate e_j of mention m is given by the cosine similarity between the entity embedding and the normalized average of all other voting entities' embeddings (of the other mentions m').

$$\begin{aligned} V_G^m &= \{e | (m', e) \in V_G \wedge m' \neq m\} \\ y_G^m &= \sum_{e \in V_G^m} y_e \\ G(e_j, m) &= \cos(y_{e_j}, y_G^m) \end{aligned}$$

This is combined with the local score, yielding

$$\Phi(e_j, m) = \text{FFNN}_3([\Psi(e_j, m); G(e_j, m)])$$

The final loss function is now slightly modified. Specifically, we enforce the linear separability in two places: in $\Psi(e, m)$ (exactly as before), but also in $\Phi(e, m)$, as follows

$$\theta^* = \arg \min_{\theta} \sum_{d \in D} \sum_{m \in M} \sum_{e \in C(m)} V(\Psi_{\theta}(e, m)) + V(\Phi_{\theta}(e, m)) \quad (9)$$

The inference procedure remains unchanged in this case, with the exception that it will only use the $\Phi(e, m)$ global score.

Coreference Resolution Heuristic. In a few cases we found important to be able to solve simple coreference resolution cases (e.g. "Alan" referring to "Alan Shearer"). These cases are difficult to handle by our candidate selection strategy. We thus adopt the simple heuristic described in [Ganea and Hofmann \(2017\)](#) and observed between 0.5% and 1% improvement on all datasets.

	AIDA A	AIDA B	MSNBC	OKE-2015	OKE-2016	N3-Reuters-128	N3-RSS-500	Derczynski	KORES0
F1@MA									
F1@MI									
FREME	23.6 37.6	23.8 36.3	15.8 19.9	26.1 31.6	22.7 28.5	26.8 30.9	32.5 27.8	31.4 18.9	12.3 14.5
FOX	54.7 58.0	58.1 57.0	11.2 8.3	53.9 56.8	49.5 50.5	52.4 53.3	35.1 33.8	42.0 38.0	28.3 30.8
Babelify	41.2 47.2	42.4 48.5	36.6 39.7	39.3 41.9	37.8 37.7	19.6 23.0	32.1 29.1	28.9 29.8	52.5 55.9
Entityclassifier.eu	43.0 44.7	42.9 45.0	41.4 42.2	29.2 29.5	33.8 32.5	24.7 27.9	23.1 22.7	16.3 16.9	25.2 28.0
Kea	36.8 40.4	39.0 42.3	30.6 30.9	44.6 46.2	46.3 46.4	17.5 18.1	22.7 20.5	31.3 26.5	41.0 46.8
DBpedia Spotlight	49.9 55.2	52.0 57.8	42.4 40.6	42.0 44.4	41.4 43.1	21.5 24.8	26.7 27.2	33.7 32.2	29.4 34.9
AIDA	68.8 72.4	71.9 72.8	62.7 65.1	58.7 63.1	0.0 0.0	42.6 46.4	42.6 42.4	40.6 32.6	49.6 55.4
WAT	69.2 72.8	70.8 73.0	62.6 64.5	53.2 56.4	51.8 53.9	45.0 49.2	45.3 42.3	44.4 38.0	37.3 49.6
Best baseline	69.2 72.8	71.9 73.0	62.7 65.1	58.7 63.1	51.8 53.9	52.4 53.3	45.3 42.4	44.4 38.0	52.5 55.9
base model	86.6 89.1	81.1 80.5	64.5 65.7	54.3 58.2	43.6 46.0	47.7 49.0	44.2 38.8	43.5 38.1	34.9 42.0
base model + att	86.5 88.9	81.9 82.3	69.4 69.5	56.6 60.7	49.2 51.6	48.3 51.1	46.0 40.5	47.9 42.3	36.0 42.2
base model + att + global	86.6 89.4	82.6 82.4	73.0 72.4	56.6 61.9	47.8 52.7	45.4 50.3	43.8 38.2	43.2 34.1	26.2 35.2
ED base model + att + global using Stanford NER mentions	75.7 80.3	73.3 74.6	71.1 71.0	62.9 66.9	57.1 58.4	54.2 54.6	45.9 42.2	48.8 42.3	40.3 46.0

Table 2: EL strong matching results on the Gerbil platform. Micro and Macro F1 scores are shown. We highlight the best and second best models, respectively. Training was done on AIDA-train set.

4 Experiments

Datasets and Metrics. We used Wikipedia 2014 as our KB. We conducted experiments on the most important public EL datasets using the Gerbil platform Röder et al. (2017). Datasets’ statistics are provided in Tables 5 and 6 (Appendix). This benchmarking framework offers reliable and trustable evaluation and comparison with state of the art EL/ED methods on most of the public datasets for this task. It also shows how well different systems generalize to datasets from very different domains and annotation schemes compared to their training sets. Moreover, it offers evaluation metrics for the end-to-end EL task, as opposed to some works that only evaluate NER and ED separately, e.g. Luo et al. (2015).

As previously explained, we do not use the NIL mentions (without a valid KB entity) and only compare against other systems using the InKB metrics.

For training we used the biggest publicly available EL dataset, AIDA/CoNLL Hoffart et al. (2011), consisting of a training set of 18,448 linked mentions in 946 documents, a validation set of 4,791 mentions in 216 documents, and a test set of 4,485 mentions in 231 documents.

We report micro and macro InKB F1 scores for

both EL and ED. For EL, these metrics are computed both in the *strong matching* and *weak matching* settings. The former requires exactly predicting the gold mention boundaries and their entity annotations, whereas the latter gives a perfect score to spans that just overlap with the gold mentions and are linked to the correct gold entities.

Baselines. We compare with popular and state-of-the-art EL and ED public systems, e.g. WAT Piccinno and Ferragina (2014), Dbpedia Spotlight Mendes et al. (2011), NERD-ML van Erp et al. (2013), KEA Steinmetz and Sack (2013), AIDA Hoffart et al. (2011), FREME, FOX Speck and Ngomo (2014), PBOH Ganea et al. (2016) or Babelify Moro et al. (2014). These are integrated within the Gerbil platform. However, we did not compare with Luo et al. (2015) and other models outside Gerbil that do not use end-to-end EL metrics.

Training Details and Model Hyper-Parameters. We use the same model hyper-parameters for both EL and ED training. The only difference between the two settings is the span set M : EL uses *all spans training*, whereas ED uses *gold spans training* settings as explained before.

Our pre-trained word and entity embeddings are

position of ground truth in $p(e m)$	Number of mentions	solved% matched gold entity%		ED Global model		ED Base + Att model		ED Base model		EL Global model		EL Base + Att model		EL Base model		EL Global -logp(e m)	
		upper	lower	upper	lower	upper	lower	upper	lower	upper	lower	upper	lower	upper	lower	upper	lower
1	3390	98.8	98.8	98.6	98.9	98.3	98.8	96.7	99.0	96.6	98.6	96.2	99.0	96.2	99.0	93.3	96.7
2	655	89.9	89.9	88.1	88.5	88.5	88.9	86.8	90.8	86.8	90.8	85.0	88.4	85.0	88.4	86.9	89.8
3	108	83.3	84.3	81.5	82.4	75.9	78.7	79.1	84.5	80.2	84.7	74.8	81.1	74.8	81.1	84.3	88.9
4-8	262	78.2	78.2	76.3	79.0	74.8	76.0	69.5	78.2	68.8	78.7	68.7	76.0	68.7	76.0	78.9	83.5
9+	247	59.9	63.6	53.4	60.7	53.0	58.7	47.8	58.2	46.2	59.4	50.4	54.8	50.4	54.8	62.7	67.5

Table 3: AIDA A dataset: Gold mentions are split by the position they appear in the $p(e|m)$ dictionary. In each cell, the upper value is the percentage of the gold mentions that were annotated with the correct entity (recall), whereas the lower value is the percentage of gold mentions for which our system’s highest scored entity is the ground truth entity, but that might not be annotated in the end because its score is below the threshold δ .

300 dimensional, while 50 dimensional trainable character vectors were used. The char LSTMs have also hidden dimension of 50. Thus, word-character embeddings are 400 dimensional. The contextual LSTMs have hidden size of 150, resulting in 300 dimensional context-aware word vectors. We apply dropout on the concatenated word-character embeddings, on the output of the bidirectional context LSTM and on the entity embeddings used in Eq. 6. The three FFNNs in our model are simple projections without hidden layers (no improvements were obtained with deeper layers). For the *long range context attention* we used a word window size of $K = 200$ and keep top $R = 10$ words after the hard attention layer (notations from [Ganea and Hofmann \(2017\)](#)). We use at most $s = 30$ entity candidate per mention both at train and test time. γ is set to 0.2 without further investigations. γ' is set to 0, but a value of 0.1 was giving similar results.

For the loss optimization we use Adam [Kingma and Ba \(2014\)](#) with a learning rate of 0.001. We perform early stopping by evaluating the model on the AIDA validation set each 10 minutes and stopping after 6 consecutive evaluations with no significant improvement in the macro F1 score.

Results and Discussion. The following models are used.

i) **Base model:** only uses the mention local score and the log-prior. It does not use long range attention, nor global disambiguation. It does not use the head attention mechanism.

ii) **Base model + att:** the Base Model plus Long Range Context Attention.

iii) **Base model + att + global:** our Global Model (depicted in figure 1)

iv) **ED base model + att + global Stanford NER:** our ED Global model that runs on top of the detected mentions of the Stanford NER system [Finkel et al. \(2005\)](#).

EL strong and weak matching results are presented in Tables 2 and 7 (Appendix).

We first note that our system outperforms all baselines on the end-to-end EL task on both AIDA-A (dev) and AIDA-B (test) datasets, which are the biggest EL datasets publicly available. Moreover, we surpass all competitors on both EL and ED by a large margin, at least 9%, showcasing the effectiveness of our method. We also outperform systems that optimize MD and ED separately, including our *ED base model + att + global Stanford NER*. This demonstrates the merit of joint MD + ED optimization.

In addition, one can observe that weak matching EL results are comparable with the strong matching results, showcasing that our method is very good at detecting mention boundaries.

At this point, our main goal was achieved: if enough training data is available with the same characteristics or annotation schemes as the test data, then our joint EL offers the best model. This is true not only when training on AIDA, but also for other types of datasets such as queries (Table 11) or tweets (Table 12). However, when testing data

1) Annotated document:

SOCCER - [SHEARER] NAMED AS [1 [ENGLAND] CAPTAIN]. [LONDON] 1996-08-30 The world 's costliest footballer [Alan Shearer] was named as the new [England] captain on Friday. The 26-year-old, who joined [Newcastle] for 15 million pounds sterling, takes over from [Tony Adams], who led the side during the [European] championship in June, and former captain [David Platt] . [2Adams] and [Platt] are both injured and will miss [England]'s opening [3[World Cup]] qualifier against [Moldova] on Sunday . [Shearer] takes the captaincy on a trial basis , but new coach [Glenn Hoddle] said he saw no reason why the former [Blackburn] and [Southampton] skipper should not make the post his own . "I 'm sure [4Alan] is the man for the job , " [Hoddle] said . [...] I spoke to [5Alan] he was up for it [...]. [Shearer] 's [Euro 96] striking partner [...].

Analysis:

[1 is considered a false negative and the ground truth is the entity `England_national_football_team`. Our annotation was for the span "ENGLAND CAPTAIN" and wrongly linked to the `England_cricket_team`. [3 The ground truth here is `1998_FIFA_World_Cup` whereas our model links it to `FIFA_World_Cup`. [2,4,5 are correctly solved due to our coreference resolution heuristic.

2) Annotated document:

[N. Korea] urges [S. Korea] to return war veteran . [SEOUL] 1996-08-31 [North Korea] demanded on Saturday that [South Korea] return a northern war veteran who has been in the [1 South] since the 1950-53 war , [Seoul] 's unification ministry said . " ...I request the immediate repatriation of Kim In-so to [North Korea] where his family is waiting , " [1 North Korean] Red Cross president Li Song-ho said in a telephone message to his southern counterpart , [2 Kang Young-hoon] . Li said Kim had been critically ill with a cerebral haemorrhage . The message was distributed to the press by the [South Korean] unification ministry . Kim , an unrepentant communist , was captured during the [2 [3 Korean] War] and released after spending more than 30 years in a southern jail . He submitted a petition to the [International Red Cross] in 1993 asking for his repatriation . The domestic [Yonhap] news agency said the [South Korean] government would consider the northern demand only if the [3 North] accepted [Seoul] 's requests , which include regular reunions of families split by the [4 Korean] War] . Government officials were not available to comment . [South Korea] in 1993 unconditionally repatriated Li In-mo , a northern partisan seized by the [5 South] during the war and jailed for more than three decades.

Analysis:

[1, [3, and [5 cases illustrate the main source of errors. These are false negatives in which our model has the correct ground truth entity pair as the highest scored one for that mention, but since it is not confident enough ($\text{score} < \gamma$) it decides not to annotate that mention. In this specific document these errors could probably be avoided easily with a better coreference resolution mechanism.

[3 and [4 cases illustrate that the gold standard can be problematic. Specifically, instead of annotating the whole span *Korean War* and linking it to the war of 1950, the gold annotation only include *Korean* and link it to the general entity of *Korea_(country)*.

[2 is correctly annotated by our system but it is not included in the gold standard.

Table 4: Error analysis on a sample document. Green corresponds to true positive (correctly discovered and annotated mention), red to false negative (ground truth mention or entity that was not annotated) and orange to false positive (incorrect mention or entity annotation).

has different statistics or follows different conventions than the training data, our method is shown to work best in conjunction with a state-of-the-art NER system as it can be seen from the results of our **ED base model + att + global Stanford NER** for different datasets in Table 2. It is expected that such a NER system designed with a broader generaliza-

tion scheme in mind would help in this case, which is confirmed by our results on different datasets.

While the main focus of this paper is the end-to-end EL and not the ED-only task, we do show ED results in Tables 8 and 9. We observe that our models are slightly behind recent top performing systems, but our unified EL - ED architecture has

to deal with other challenges, e.g. being able to exchange global information between many more mentions at the EL stage, and is thus not suitable for expensive global ED strategies. We leave bridging this gap for ED as future work.

Additional results and insights are shown in the Appendix.

Ablation study Table 3 shows an ablation study of our method. One can see that the $\log p(e|m)$ prior is very helpful for correctly linking unambiguous mentions, but is introducing noise when gold entities are not frequent. For this category of rare entities, removing this prior completely will result in a significant improvement, but this is not a practical choice since the gold entity is unknown at test time.

Error Analysis. We conducted a qualitative experiment shown in Table 4. We showcase correct annotations, as well as errors done by our system on the AIDA datasets. Inspecting the output annotations of our EL model, we discovered the remarkable property of not over-generating incorrect mentions, nor under-generating (missing) gold spans. We also observed that additional mentions generated by our model do correspond in the majority of time to actual KB entities, but are incorrectly forgotten from the gold annotations.

5 Conclusion

We presented the first neural end-to-end entity linking model and show the benefit of jointly optimizing entity recognition and linking. Leveraging key components, namely word, entity and mention embeddings, we prove that engineered features can be almost completely replaced by modern neural networks. Empirically, on the established Gerbil benchmarking platform, we exhibit state-of-the-art performance for EL on the biggest public dataset, AIDA/CoNLL, also showing good generalization ability on other datasets with very different characteristics when combining our model with the popular Stanford NER system.

Our code is publicly available².

References

- David Carmel, Ming-Wei Chang, Evgeniy Gabrilovich, Bo-June Paul Hsu, and Kuansan Wang. 2014.
- ²https://github.com/dalab/end2end_neural_el
- Erd'14: entity recognition and disambiguation challenge. In *ACM SIGIR Forum*, volume 48, pages 63–77. ACM.
- Diego Ceccarelli, Claudio Lucchese, Salvatore Orlando, Raffaele Perego, and Salvatore Trani. 2013. Dexter: an open source framework for entity linking. In *Proceedings of the sixth international workshop on Exploiting semantic annotations in information retrieval*, pages 17–20. ACM.
- Jason PC Chiu and Eric Nichols. 2016. Named entity recognition with bidirectional lstm-cnns. *Transactions of the Association for Computational Linguistics*, 4:357–370.
- Marco Cornolti, Paolo Ferragina, Massimiliano Ciaramita, Stefan Rüd, and Hinrich Schütze. 2016. A piggyback system for joint entity mention detection and linking in web queries. In *Proceedings of the 25th International Conference on World Wide Web*, pages 567–578. International World Wide Web Conferences Steering Committee.
- Joachim Daiber, Max Jakob, Chris Hokamp, and Pablo N Mendes. 2013. Improving efficiency and accuracy in multilingual entity extraction. In *Proceedings of the 9th International Conference on Semantic Systems*, pages 121–124. ACM.
- Leon Derczynski, Diana Maynard, Giuseppe Rizzo, Marieke van Erp, Genevieve Gorrell, Raphaël Troncy, Johann Petrak, and Kalina Bontcheva. 2015. Analysis of named entity recognition and linking for tweets. *Information Processing & Management*, 51(2):32–49.
- MGJ van Erp, G Rizzo, and R Troncy. 2013. Learning with the web: Spotting named entities on the intersection of nerd and machine learning. In *CEUR workshop proceedings*, pages 27–30.
- Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, pages 363–370. Association for Computational Linguistics.
- Octavian-Eugen Ganea, Marina Ganea, Aurelien Luchi, Carsten Eickhoff, and Thomas Hofmann. 2016. Probabilistic bag-of-hyperlinks model for entity linking. In *Proceedings of the 25th International Conference on World Wide Web*, pages 927–938. International World Wide Web Conferences Steering Committee.
- Octavian-Eugen Ganea and Thomas Hofmann. 2017. Deep joint entity disambiguation with local neural attention. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2619–2629.

- Zhengyan He, Shujie Liu, Mu Li, Ming Zhou, Longkai Zhang, and Houfeng Wang. 2013. Learning entity representation for entity disambiguation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 30–34.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Johannes Hoffart, Stephan Seufert, Dat Ba Nguyen, Martin Theobald, and Gerhard Weikum. 2012. Kore: keyphrase overlap relatedness for entity disambiguation. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, pages 545–554. ACM.
- Johannes Hoffart, Mohamed Amir Yosef, Ilaria Bordini, Hagen Fürstenau, Manfred Pinkal, Marc Spaniol, Bilyana Taneva, Stefan Thater, and Gerhard Weikum. 2011. Robust disambiguation of named entities in text. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 782–792. Association for Computational Linguistics.
- Yangfeng Ji, Chenhao Tan, Sebastian Martschat, Yejin Choi, and Noah A Smith. 2017. Dynamic entity representations in neural language models. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1830–1839.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of NAACL-HLT*, pages 260–270.
- Phong Le and Ivan Titov. 2018. Improving entity linking by modeling latent relations between mentions. *arXiv preprint arXiv:1804.10637*.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197.
- Liyuan Liu, Jingbo Shang, Frank Xu, Xiang Ren, Huan Gui, Jian Peng, and Jiawei Han. 2017. Empower sequence labeling with task-aware neural language model. *arXiv preprint arXiv:1709.04109*.
- Gang Luo, Xiaojiang Huang, Chin-Yew Lin, and Zaiqing Nie. 2015. Joint entity recognition and disambiguation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 879–888.
- Pablo N Mendes, Max Jakob, Andrés García-Silva, and Christian Bizer. 2011. Dbpedia spotlight: shedding light on the web of documents. In *Proceedings of the 7th international conference on semantic systems*, pages 1–8. ACM.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Andrea Moro, Alessandro Raganato, and Roberto Navigli. 2014. Entity linking meets word sense disambiguation: a unified approach. *Transactions of the Association for Computational Linguistics*, 2:231–244.
- Dat Ba Nguyen, Martin Theobald, and Gerhard Weikum. 2016. J-nerd: joint named entity recognition and disambiguation with rich linguistic features. *Transactions of the Association for Computational Linguistics*, 4:215–229.
- Andrea Giovanni Nuzzolese, Anna Lisa Gentile, Valentina Presutti, Aldo Gangemi, Darío Garigliotti, and Roberto Navigli. 2015. Open knowledge extraction challenge. In *Semantic Web Evaluation Challenge*, pages 3–15. Springer.
- Francesco Piccinno and Paolo Ferragina. 2014. From tagme to wat: a new entity annotator. In *Proceedings of the first international workshop on Entity recognition & disambiguation*, pages 55–62. ACM.
- Priya Radhakrishnan, Partha Talukdar, and Vasudeva Varma. 2018. Elden: Improved entity linking using densified knowledge graphs. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 1844–1853.
- Giuseppe Rizzo, Marieke van Erp, and Raphaël Troncy. Benchmarking the extraction and disambiguation of named entities on the semantic web.
- Michael Röder, Ricardo Usbeck, and Axel-Cyrille Ngonga Ngomo. 2017. Gerbil—benchmarking named entity recognition and linking consistently. *Semantic Web*, (Preprint):1–21.
- Avirup Sil and Alexander Yates. 2013. Re-ranking for joint named-entity recognition and linking. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, pages 2369–2374. ACM.
- René Speck and Axel-Cyrille Ngonga Ngomo. 2014. Ensemble learning for named entity recognition. In *International semantic web conference*, pages 519–534. Springer.
- Valentin I Spitkovsky and Angel X Chang. A cross-lingual dictionary for english wikipedia concepts.
- Nadine Steinmetz and Harald Sack. 2013. Semantic multimedia information retrieval based on contextual descriptions. In *Extended Semantic Web Conference*, pages 382–396. Springer.

- Yaming Sun, Lei Lin, Duyu Tang, Nan Yang, Zhenzhou Ji, and Xiaolong Wang. 2015. Modeling mention, context and entity with neural networks for entity disambiguation. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Ikuya Yamada, Hiroyuki Shindo, Hideaki Takeda, and Yoshiyasu Takefuji. 2016. Joint learning of the embedding of words and entities for named entity disambiguation. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 250–259.
- Yi Yang, Ozan Irsoy, and Kazi Shefaet Rahman. 2018. Collective entity disambiguation with structured gradient tree boosting. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 777–786.

³<https://github.com/dice-group/gerbil/issues/98>