



HAL
open science

Multimodal entity linking for tweets

Omar Adjali, Romaric Besançon, Olivier Ferret, Hervé Le Borgne, Brigitte Grau

► **To cite this version:**

Omar Adjali, Romaric Besançon, Olivier Ferret, Hervé Le Borgne, Brigitte Grau. Multimodal entity linking for tweets. ECIR 2020 - 42nd European Conference on Information Retrieval Research, Apr 2020, Lisbonne (Online event), Portugal. pp.463-478, <10.1007/978-3-030-45439-5_31>. <hal-04315181>

HAL Id: hal-04315181

<https://hal.science/hal-04315181v1>

Submitted on 30 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Multimodal Entity Linking for Tweets*

Adjali, Omar¹, Besançon, Romaric¹, Ferret, Olivier¹, Le Borgne, Hervé¹, and Grau, Brigitte²

¹Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

²Université Paris-Saclay, CNRS, LIMSI, 91400, Orsay, France

Abstract

In many information extraction applications, entity linking (EL) has emerged as a crucial task that allows leveraging information about named entities from a knowledge base. In this paper, we address the task of multimodal entity linking (MEL), an emerging research field in which textual and visual information is used to map an ambiguous mention to an entity in a knowledge base (KB). First, we propose a method for building a fully annotated Twitter dataset for MEL, where entities are defined in a Twitter KB. Then, we propose a model for jointly learning a representation of both mentions and entities from their textual and visual contexts. We demonstrate the effectiveness of the proposed model by evaluating it on the proposed dataset and highlight the importance of leveraging visual information when it is available.

1 Introduction

Entity linking (EL) is a crucial task for natural language processing applications that require disambiguating named mentions within textual documents. It consists in mapping ambiguous named mentions to entities defined in a knowledge base (KB). To address the EL problem, most of the state-of-the-art approaches use some form of textual representations associated with the target mention and its corresponding entity. [4] first proposed to link a named mention to an entity in a knowledge base, followed by several other local approaches [10, 36, 15, 16, 11] where mentions are individually disambiguated using lexical mention-entity measures and context features extracted from their surrounding text. In contrast, global approaches [43, 22, 40, 20] use global features at the document-level to disambiguate all the mentions within the document and can also exploit semantic relationships between entities in the KB. These approaches have shown their effectiveness on standard EL datasets such as TAC KBP [28], CoNLL-YAGO [24], and ACE [3] which contain structured documents that provide a rich context for disambiguation. However, EL becomes more challenging when a limited, informal, textual context is available, such as in social media posts. On the other hand, social media posts often include images to illustrate the text that could be exploited to improve the disambiguation by adding a visual context.

In this paper, we address the problem of multimodal entity linking (MEL), by leveraging both semantic textual and visual information extracted from mention and entity contexts. We propose to apply MEL on Twitter posts as they form a prototypical framework where the textual context is generally poor and visual information is available. Indeed, an important mechanism in Twitter communication is the usage of a user's screen name (@UserScreenName) in a tweet text. This helps to explicitly mention that the corresponding user is somehow related to the posted tweet. One observation we made is that some Twitter users tend to mention other users without resorting to screen names but rather use their first name, last name, or acronyms (when the Twitter user is an organization). Consider the following example tweet (see Figure 1):

Andrew explains how embracing AI doesn't have to be daunting

*Preprint of Adjali O., Besançon R., Ferret O., Le Borgne H., Grau B. (2020) Multimodal Entity Linking for Tweets. In: Jose J. et al. (eds) Advances in Information Retrieval. ECIR 2020. Lecture Notes in Computer Science, vol 12035. Springer, Cham. The final authenticated version is available online at https://doi.org/10.1007/978-3-030-45439-5_31




Figure 1: An illustrative example of multimodal Entity Linking on our Twitter dataset. The example depicts a subset of (text, image) pairs representing the entity @AndrewYNg (left), and some example of (text, image) pairs related to the mention *Andrew* (right).

In this example, the mention *Andrew* refers to the user’s screen name @AndrewYNg. Obviously, the mention *Andrew* could have referred to any Twitter user whose name includes Andrew. Such a practice, when it occurs, may lead to ambiguities that standard EL systems may not be able to resolve, as most of Twitter users do not have entries in standard knowledge bases. We thus propose to link ambiguous mentions to a specific Twitter KB, composed of Twitter users.

Our contributions are: (i) we investigate the multimodal entity linking task. To the best of our knowledge, this is the first EL work that combines textual and visual contexts on a Twitter dataset; (ii) we construct a new large-scale dataset for multimodal EL and above all, define a method for building such datasets at convenience; (iii) we present a new EL model-based for learning a multimodal joint representation of tweets and show, on the proposed dataset, the interest of adding visual features for the EL task. Code and data are available at https://github.com/OA256864/MEL_Tweets¹.

2 Related Work

Entity Linking on Twitter Recently, several research efforts proposed to meet the challenges posed by the EL task on Twitter media posts. Collective approaches are preferred and leverage global information about tweets in relation with the target mention. [34] collectively resolves a set of mentions by aggregating all their related tweets to compute mention-mention and mention-entity similarities while [26] takes advantage of both semantic relatedness and mention co-referencing. In a more original way, [45] determines the user’s topics of interest from all its posted tweets in order to collectively link all its named entity mentions. Similarly [25] considers social (user’s interest + popularity) and temporal contexts. Other collective approaches include in their EL model additional non-textual features. For example, [17] and [7] use global information of tweets that are close in space and time to the tweet of the target mention. Finally, [14] proposes a joint cross-document co-reference resolution and disambiguation approach including temporal information associated with their corpus to improve EL performance. While these works yield interesting results using non-textual features, they often depend on the availability of social data and do not exploit visual information.

Multimodal Representation Learning Joint multimodal representations are used in several multimodal applications such as visual question answering [27, 2], text-image retrieval [49, 5, 8] and image captioning [31, 29] and exploits different approaches e.g., Canonical Correlation Analysis [47], linear ranking-based models and non-linear deep learning models to learn projecting image features and text features into a joint space [18]. In our work, we take inspiration from [48] for our joint representation learning model as it showed satisfying results for multimodal learning.

To our knowledge, [37] is the only other work that leverages multimodal information for an entity disambiguation task in a social media context. However, they use a dataset of 12K annotated image-caption pairs from Snapchat (which is not made available), whereas our work relies on a much larger

¹forked at https://github.com/cea-list-lasti/MEL_Tweets

base (85K samples for the benchmark and 2M to build the KB) and is more in line with the state-of-the-art which mostly uses Twitter as case study for EL on social media. Furthermore, they use Freebase as KB, which does not contain image information and does not allow a multimodal representation at the KB level, leading to a very different disambiguation model.

3 Problem Formulation

In Twitter terms, each user account has a unique screen name (@AndrewYNg), a user name (Andrew Ng), a user description (Stanford CS adjunct faculty) and a timeline containing all the tweets (text+image) posted by the user (see Figure 1). On this basis, we consider the following definitions:

- an entity e corresponds to a Twitter user account u (generally associated with a person or an organization);
- a mention m corresponds to an ambiguous textual occurrence of an entity e_j mentioned in a tweet t that does not belong to the timeline of e_j ;
- the knowledge base is the set of entities (i.e. Twitter accounts).

Formally, we denote the knowledge base $KB = \{e_j\}$ as a set of entities, each entity being defined as a tuple $e_j = (s_j, u_j, TL_j)$ including their screen name s_j , user name u_j and timeline TL_j (we did not use the user description in the representation of the entity). The timeline contains both texts and images. A mention m_j is defined as a pair (w_i, t_i) composed of the word (or set of words) w_i characterizing the mention and the tweet t_i in which it occurs: t_i contains both text and images.

The objective of the task consists in finding the most similar entity $e^*(m_i)$ to the mention m_i in the KB according to a given similarity. From a practical point of view, we do not compute the similarities between the mention and all entities in the KB. We first select a subset of the entities that are good candidates to disambiguate m_i . It is defined as $Cand(m_i) = \{e_j \in KB | w_i \sim u_j\}$ where $w_i \sim u_j$ indicates that the mention words w_i are close to the entity name u_j . In our case, due to the nature of the dataset, we only use a simple inclusion (i.e. the mention words are present in the entity name) but a more complex lexical distance that takes into account more variations could be considered [37]. The disambiguation of m_i is then formalized as finding the best multimodal similarity measure between the tweet containing the mention and the timeline of the correct entity, both containing text and images:

$$e^*(m_i) = \underset{e_j \in Cand(m_i)}{\operatorname{argmax}} \operatorname{sim}(t_i, TL_j) \quad (1)$$

4 Twitter-MEL Dataset

One important contribution of our work is a novel dataset for multimodal entity linking on short texts accompanied by images derived from Twitter posts. The process to build this dataset is mostly automatic and can, therefore, be applied to generate a new dataset at convenience. More details and further analysis on the dataset can be found in [1].

All the tweets and user’s metadata were collected using the Twitter official API². The dataset creation process comprises two phases: the construction of the knowledge base and the generation of ambiguous mentions. As mentioned in Section 1, our multimodal EL task aims at mapping mentions to entities (i.e. twitter users) from our Twitter knowledge base, which are characterized by a timeline, namely the collection of the most recent tweets posted by a given user.

As a first step, we established a non-exhaustive initial list of Twitter user’s screen names using Twitter lists in order to have users that are likely to produce a sufficient number of tweets and be referred by a sufficient number of other users. A Twitter list is a curated set of Twitter accounts³ generally grouped by topic. From this initial list of users, we started building the KB by collecting the tweets of each user’s timeline along with its meta-information ensuring that both re-tweets and tweets without images were discarded. Moreover, as explained previously, users tend to create ambiguous mentions in tweets when they employ any expression (for example first or last name) other than Twitter screen names to mention other users in their post (see Section 1). Consequently, we have drawn

²<https://dev.twitter.com>

³<https://help.twitter.com/en/using-twitter/twitter-lists>

Table 1: Statistics on timeline and entity distributions in MEL dataset.

	Mean	Median	Max	Min	StdDev
nb Tweets / timeline (text+image)	127.9	52	3,117	1	222.2
nb ambiguous entities/mention	16.5	16	67	2	12

inspiration from this usage to elaborate a simple process for both candidate entity and ambiguous mention generation.

4.1 Selection of Possibly Ambiguous Entities

To ensure that the KB contains sufficiently ambiguous entities, i.e. entities with possible ambiguous mentions, to make the EL task challenging, we expanded the KB from the initial lists with ambiguous entities. More precisely, we first extracted the last name from each Twitter account name of the initial list of users. Then, we used these names as search queries in the Twitter API user search engine to collect data about similar users⁴. Users that have been inactive for a long period, non-English users and non-verified user accounts were filtered. Furthermore, to ensure more diversity in the entities and not only use person names, we manually collected data about organization accounts. We relied on Wikipedia acronym disambiguation pages to form groups of ambiguous (organization) entities that share the same acronym.

4.2 Generation of Ambiguous Mentions

After building the KB, we used the collected entities to search for tweets that mention them. The Twitter Search API⁵ returns a collection of relevant tweets matching the specified query. Thus, for each entity in the KB, **(i)** we set its screen name (@user) as the query search; **(ii)** we collect all the retrieved tweets; **(iii)** we filtered out tweets without images. Given the resulting collections of tweets mentioning the different entities of the KB, we systematically replaced the screen name mentioned in the tweet with its corresponding ambiguous mention: last names for *person* entities and acronyms for *organization* entities. Finally, we track of the ground truths of each tweet in the dataset reducing the cost of a manual annotation task and resulting in a dataset composed of annotated pairs of text and image. Although ambiguous mentions are synthetically generated, they are comparable to some extent with real-world ambiguous mentions in tweets. We applied a named entity recognition (NER) system [13] on the tweets of the dataset which achieved a 77% accuracy score on the generated mentions. This suggests that these latter are somehow close to real-world named mentions.

4.3 Dataset Statistics and Analysis

Altogether, we collected and processed 14M tweets, 10M timeline tweets for entity characterization and 4M tweets with ambiguous mentions (mention tweets) covering 20k entities in the KB. Filtering these tweets drastically reduced the size of our data set. Regarding mention tweets, a key factor in the reduction is the elimination of noisy tweets where mentions have no syntactic role in the tweet text (e.g. where the mention is included in a list of recipients of the message). Discarding these irrelevant tweets as well as tweets without image left a dataset of 2M timeline tweets and 85k mention tweets. After 3 months of data collection, we found that only 10% of tweet posts are accompanied by images. In the end, we randomly split the set of mention tweets into training (40%), validation (20%) and test (40%) while ensuring that 50% of mention tweets in the test set correspond to entities unseen during the training.

Table 1 shows the timeline tweet distribution of all entities in our KB. As noted by [25], this distribution reveals that most Twitter users are information seekers, i.e. they rarely tweet, in contrast to users that are content generators who tweet frequently. Along with user’s popularity, this has an influence on the number of mention tweets we can collect. We necessarily gathered more mention tweets from content generator entities, as they are more likely to be mentioned by others than information seeker entities.

⁴Only the first 1,000 matching results are available with the Twitter API.

⁵Twitter API searches within a sampling of tweets published in the past 7 days.

5 Proposed MEL Approach

Visual and textual representations of mentions and entities are extracted with pre-trained networks. We then learn a two-branch feed-forward neural network to minimize a triplet-loss defining an implicit joint feature space. This network provides a similarity score between a given mention and an entity that is combined with other external features (such as popularity or other similarity measures) as the input of a multi-layer perceptron (MLP) which performs a binary classification on a (mention, entity) pair.

5.1 Features

Textual context representation We used the unsupervised Sent2Vec [38] model to learn tweet representations, pre-trained on large Twitter corpus. We adopted this model as training on the same type of data (short noisy texts) turns out to be essential for performing well on the MEL task (see Sec. 6.2). The Sent2Vec model extends the CBOW model proposed by [35] to learn a vector representation of words. More precisely, it learns a vector representation of a sentence S by calculating the average of the embeddings of the words (unigrams) making up the sentence S and the n-grams present in S :

$$V_S = \frac{1}{|R(S)|} \sum_{w \in R(S)} v_w \quad (2)$$

where $R(S)$ is the set of n-grams (including unigrams) in the sentence S , and v_w is the embedding vector of the word w . Therefore, the textual context of a mention m_i within the tweet t_i is represented by the sentence embedding vector of t_i . We produce then for each mention two continuous vector representations ($D=700$), a sentence embedding $U_m^{(i)}$ inferred using only tweet unigrams and a sentence embedding $B_m^{(i)}$ inferred using tweet unigrams and bigrams. Combining their vectors is generally beneficial [38]. An entity context being represented by a set of tweets (see Section 3), given an entity e_i , we average the unigram and bigram embeddings of all e_i 's timeline tweets yielding two average embedding vectors $U_e^{(i)}$, $B_e^{(i)}$ representing the entity textual context used as features.

BM25 features

Given that the disambiguation task aims at finding the correct entity for a tweet, it can be viewed as an IR problem, where we try to associate a given tweet with the most relevant timeline in the KB. We, therefore, consider, as a baseline, a *tf-idf*-like model to match the mention with the entity: in our case, both the tweet and the timeline are represented as bag-of-words vectors and we used the standard BM25 weighting scheme to perform the comparison.

Popularity features Given an entity e representing a Twitter user u , we consider 3 popularity features represented by: N_{fo} the number of followers, N_{fr} the number of friends and N_t the number of tweets posted by u .

Visual context features The visual features are extracted with the Inception_v3 model [46], pre-learned on the 1.2M images of the ILSVRC challenge [44]. We use its last layer ($D = 1,000$), which encodes high-level information that may help discriminating between entities. For an entity e_i , we retain a unique feature vector that is the average of the feature vectors of all the images within its timeline, similarly to the process for the textual context. The visual feature vector of a mention is extracted from the image of the tweet that contains the mention.

5.2 Joint Multimodal Representation Learning

The proposed model measures the multimodal context similarity between a mention and its candidate entities. Figure 2 shows the architecture of the proposed joint representation learning model. It follows a triplet loss structure with two sub-models as in [23, 42], one processing the mention contexts and the other processing the entity contexts. Each sub-model has a structure resembling the similarity model proposed by [48], i.e., it comprises three branches (unigram embedding, bigram embedding, image feature), each in turn including 2 fully connected (FC) layers with a Rectified Linear Unit (ReLU) activation followed by a normalization layer [33]. Then, the output vectors of the three branches are merged using concatenation, followed by a final FC layer. Other merging approaches exist such as element-wise product and compact bilinear pooling [18, 48]. However, in our work, simple concatenation showed satisfying results as a first step. Moreover, the mention and entity inputs differ




Figure 2: Triplet ranking model for learning entity and mention joint representations.

in our task, a mention being characterized by one (text, image) pair and an entity by a large set of (text, image) pairs. Thus, we investigated the performance of our model when the parameters of the two sub-models are partially or fully shared. We found that shared parameters yielded better accuracy results on the validation set, thus the weight of the FC layers of both branches are shared.

In summary, using the extracted visual and textual features $\{U_m, B_m, I_m\}$, $\{U_e, B_e, I_e\}$ respectively of a mention m_i and an entity e_i , our model is trained to project each modality through the branches into an implicit joint space. The resulting representations are concatenated into C_m and C_e and passed through the final FC layer yielding two multimodal vectors J_m and J_e . The objective is to minimize the following triplet loss function:

$$\min_W \sum_{e^- \neq e^+} \max(0, 1 - \|f_m(m), f_e(e^+)\| - \|f_m(m), f_m(e^-)\|)$$

where m, e^+, e^- , are the mention, positive and negative entities respectively. $f_m(\cdot)$ is the mention sub-network output function, $f_e(\cdot)$ the entity sub-network function and $\|\cdot\|$ is the L_2 -norm. The objective aims at minimizing the L_2 distance between the multimodal representation of m and the representation of a positive entity e^+ , and maximizing the distance to the negative entity e^- multimodal representation. For MEL, we calculate the cosine similarity between the two vectors of a pair (m_i, e_i) given by $f_m(\cdot)$ and $f_e(\cdot)$ to represent their multimodal context similarity: $sim(m_i, e_i) = \text{cosine}(J_m^{(i)}, J_e^{(i)})$.

Finally, we propose to integrate the popularity of the entity to the estimation of the similarity by combining the multimodal context similarity score and popularity features through a MLP. Other external features may also be combined at the MLP level, as in our case the BM25 similarity. This MLP is trained to minimize a binary cross-entropy, with label 0 (resp. 1) for negative (resp. positive) entities w.r.t the mention.

Features	Description (see Section 5.1)
Popularity (Pop)	Baseline feature where the most popular entity is selected;
BM25	Standard textual context similarity with BM25 weighting;
S2V-uni	Similarity measured between the unigram embeddings extracted using the Sent2Vec language model;
S2V-bi	Similarity measured between the bigram embeddings extracted using the Sent2Vec language model;
S2V	For easy readability, we use the S2V notation to represent the combination of S2V-uni and S2V-bi;
Img	Similarity measured between the image features extracted using the pretrained Inception-V3 model;
ET(X)	Combination of features X using an Extra-Trees classifier;
JMEL(X)	Combination of features X with our joint multimodal representation model (see Section 5.2).

6 Experiments Analysis and Evaluation Setup

6.1 Parameter Settings

We initialized the weights of each FC layer using the Xavier initialization [21]. Training is performed over 100 epochs, each covering 600k (mention, entity) pairs of which 35k are positive samples. The triplet loss is minimized by stochastic gradient descent, with a momentum of 0.9, an initial learning rate of 0.1 and a batch size of 256. The learning rate is divided by 10 when the loss does not change by more than $1e^{-4}$ during 6 epochs. After 50 epochs, if the accuracy on the validation set does not increase during 5 successive evaluations, the training is early stopped. Regarding the BCE branch, the MLP has two hidden layers with one more neuron than the number of inputs and *tanh* non-linearities. The BCE loss is minimized using L-BGFS with a learning rate of 10^{-5} . The network is implemented in PyTorch [39].

Since our approach is not directly comparable with previous works [37], we compare the results with different configurations and baselines. As 50% of mentions are unique in each split of our dataset, we do not take the standard most frequent entity prediction as a baseline, which links the entity that was seen most in training to its corresponding mention. We rather consider the popularity features and standard textual similarity measures (BM25), presented in section 5.1. Moreover, for comparison sake, we also report in Table 3 the performance of a baseline combination of the features using an Extra-Trees classifier [19], compared with our Joint Multimodal Entity Linking approach (JMEL), which combines the textual and visual features at the representation level. Table 2 summarizes the features and the combination models used in these experiments.

In order to more thoroughly assess the contribution of textual and visual information in the MEL task, we compare the performance of the proposed model when the visual features are combined with sentence embeddings that are derived from different models the following notable models: SkipThought [32] (D=4,800) trained on a large corpus of novels, Sent2Vec [38] unigram and bigram vectors (D=700) trained respectively on English Wikipedia and on a Twitter corpus, InferSent [9] (D=2,048) Bi-LSTM+max pooling model trained on the Stanford SNLI corpus, BERT [13] trained on BooksCorpus [50] and on English Wikipedia and ELMo [41] trained on the One Billion Word Benchmark [6]. For BERT and ELMo, sentences are represented by the average word embedding.

6.2 Results

Table 3 reports the accuracy results on the validation and test sets for the binary classification task about the correctness of the first entity selected from the KB for a given mention. First, we note that the baseline points out an imbalance in our dataset, as 59% of the mentions in the test set correspond to the most popular entity among the candidate entities, compared to 36.9% in the validation set. Note that for some Entity linking datasets, popularity can achieve up to 82% accuracy score [12]. We also observe that our popularity baseline outperforms the combination of textual and visual features with the Extra-Trees classifier: this indicates that the features extracted from the textual and visual

Table 3: Multimedia Entity Linking results (accuracy).

	Valid.	Test
Single features		
Popularity	0.369	0.590
BM25	0.415	0.433
S2V-uni	0.482	0.513
S2V-bi	0.487	0.523
Img	0.290	0.299
Combination of features with an ExtraTrees Classifier		
ET(S2V)	0.495	0.529
ET(S2V + Img)	0.507	0.542
ET(S2V + Img + Pop)	0.585	0.627
ET(S2V + Img + Pop + BM25)	0.654	0.671
Combination of features with our JMEL model		
JMEL(S2V)	0.628	0.724
JMEL(S2V + Img)	0.639	0.731
JMEL(S2V + Img + Pop)	0.767	0.776
JMEL(S2V + Img + Pop + BM25)	0.795	0.803

contexts, when naively used and combined, produce poor results. In contrast, our model achieves significant improvements on both the validation and test sets compared with the popularity baseline and the Extra-Trees combination. We see that combining additional features in the JMEL model (popularity and BM25) also provides significant performance gain. Regarding the visual modality, although considering it alone leads to poor results, its integration in a global model always improve the performance, compared to text-only features.

Image and sentence representation impact analysis Table 4 reports the MEL accuracy on the validation and test sets with various sentence representation models. It shows that the integration of visual features always improves the performance of our approach, whatever the sentence embedding model we use, even though the level of improvement varies depending on the sentence model. For example, while the results of averaging BERT word embeddings and InferSent are comparable using images on the test set, InferSent performs significantly better than BERT using text only. This emphasizes the role of the visual context representation of mentions and entities to help in the EL task. If we look at the performance of the various sentence embeddings models, we can see that the sent2vec model trained on a Twitter corpus outperforms all other embeddings: this reveals the importance of training data in a transfer learning setting. Indeed, we observe that the sent2vec model trained on English Wikipedia produces worse results than the model trained on the target task data (Twitter). Hence, we can assume that other models that achieve good results when trained on a general corpus (such as InferSent or BERT) would get better results if trained on a Twitter collection.

Error analysis We identified several potential sources of errors that may be addressed in future work. First, in our approach we characterize entity contexts with a collection of text/image pairs and by taking their mean, we consider all these pairs equally important to represent an entity. However, by manually checking some entities, we note that each timeline may contain a subset of outlier pairs and more specifically, images that are not representative of an entity. Sampling strategies may be employed to select the most relevant images and to discard the misleading outliers. Moreover, our model fails on some difficult cases where the visual and textual contexts of entity candidates are indistinguishable. For example, it fails on the mention "post", by linking it to the entity *@nationalpost* instead of *@nypost*, two entities representing news organizations whose posts cover various topics. One additional bias is that we restricted our dataset and KB to have only tweets with images. Meanwhile, we observed that tweets without images tend to have more textual context. Thus, it would be interesting to include tweets without images for further experiments.

Table 4: Impact of sentence embeddings on EL results.

Sent. Embedding	Valid		Test	
	Txt	Txt+Img	Txt	Txt+Img
S2V-uni(Twitter)	0.592	0.611	0.698	0.708
S2V-uni(Wiki)	0.499	0.538	0.625	0.654
S2V-bi(Twitter)	0.616	0.637	0.709	0.716
S2V-bi(Wiki)	0.511	0.547	0.639	0.663
InferSent(GloVe)	0.559	0.579	0.666	0.683
InferSent(fastText [30])	0.551	0.570	0.671	0.689
Avg. BERT	0.580	0.594	0.641	0.687
Avg. ELMo	0.524	0.563	0.605	0.655
Skip-Thought	0.464	0.511	0.575	0.605
S2V(Twitter)	0.628	0.639	0.724	0.731
S2V(Wiki)	0.524	0.551	0.652	0.666

7 Conclusion

We explore a novel approach that makes use of text and image information for the entity linking task applied to tweets. Specifically, we emphasize the benefit of leveraging visual features to help in the disambiguation. We propose a model that first extracts textual and visual contexts for both mentions and entities and learns a joint representation combining textual and visual features. This representation is used to compute a multimodal context similarity between a mention and an entity. Preliminary experiments on a dedicated dataset demonstrated the effectiveness of the proposed model and revealed the importance of leveraging visual contexts in the EL task. Furthermore, our work is all the more relevant with the emergence of social media, which offer abundant textual and visual information. In that perspective, we propose a new multimodal EL dataset based on Twitter posts and a process for collecting and constructing a fully annotated multimodal EL dataset, where entities are defined in a Twitter KB.

Further exploration is still needed concerning certain points in our model: in particular, our future work includes exploring attention mechanisms for text and images and experimenting different sampling strategies for the triplet loss.

Acknowledgments. This research was partially supported by Labex DigiCosme (project ANR11LABEX-0045DIGICOSME) operated by ANR as part of the program “Investissements d’Avenir” Idex Paris Saclay (ANR11IDEX000302).

References

- [1] Omar Adjali, romaric Besancon, olivier Ferret, Hervé Le Borgne, and Brigitte Grau. Building a multimodal entity linking dataset from tweets. In *International Conference on Language Resources and Evaluation (LREC)*, Marseille, France, may 2020.
- [2] Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C Lawrence Zitnick, Devi Parikh, and Dhruv Batra. VQA: Visual question answering. *International Journal of Computer Vision*, 123(1):4–31, 2017.
- [3] Luisa Bentivogli, Pamela Forner, Claudio Giuliano, Alessandro Marchetti, Emanuele Pianta, and Kateryna Tymoshenko. Extending English ACE 2005 corpus annotation with ground-truth links to Wikipedia. In *Proceedings of the 2nd Workshop on The People’s Web Meets NLP: Collaboratively Constructed Semantic Resources*, pages 19–27, 2010.
- [4] Razvan Bunescu and Marius Paşca. Using encyclopedic knowledge for named entity disambiguation. In *11th Conference of the European Chapter of the Association for Computational Linguistics*, 2006.

- [5] Ines Chami, Youssef Tamaazousti, and Hervé Le Borgne. Amecon: Abstract meta-concept features for text-illustration. In *International Conference on Multimedia Retrieval (ICMR)*, Bucharest, Romania, 2017.
- [6] Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, Phillipp Koehn, and Tony Robinson. One Billion Word Benchmark for Measuring Progress in Statistical Language Modeling. In *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [7] Wen-Haw Chong, Ee-Peng Lim, and William Cohen. Collective entity linking in tweets over space and time. In *European Conference on Information Retrieval*, pages 82–94. Springer, 2017.
- [8] Mithun Chowdhury, Panda Rameswar, Evangelos Papalexakis, and Amit Roy-Chowdhury. Webly Supervised Joint Embedding for Cross-Modal Image-Text Retrieval. In *ACM International Conference on Multimedia*, 2018.
- [9] Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. Supervised Learning of Universal Sentence Representations from Natural Language Inference Data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark, 2017. Association for Computational Linguistics.
- [10] Silviu Cucerzan. Large-scale named entity disambiguation based on Wikipedia data. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 708–716, 2007.
- [11] Hani Daher, Romaric Besançon, Olivier Ferret, Hervé Le Borgne, Anne-Laure Daquo, and Youssef Tamaazousti. Supervised learning of entity disambiguation models by negative sample selection. In *International Conference on Computational Linguistics and Intelligent Text Processing (CI-CLing)*, Budapest, Hungary, 2017. 17 – 23 avril.
- [12] Hongliang Dai, Yangqiu Song, Liwei Qiu, and Rijia Liu. Entity Linking within a Social Media Platform: A Case Study on Yelp. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2023–2032, 2018.
- [13] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186, 2019.
- [14] Mark Dredze, Nicholas Andrews, and Jay DeYoung. Twitter at the grammys: A social media corpus for entity linking and disambiguation. In *Proceedings of The Fourth International Workshop on Natural Language Processing for Social Media*, pages 20–25, 2016.
- [15] Mark Dredze, Paul McNamee, Delip Rao, Adam Gerber, and Tim Finin. Entity disambiguation for knowledge base population. In *Proceedings of the 23rd International Conference on Computational Linguistics*, pages 277–285. Association for Computational Linguistics, 2010.
- [16] Yotam Eshel, Noam Cohen, Kira Radinsky, Shaul Markovitch, Ikuya Yamada, and Omer Levy. Named Entity Disambiguation for Noisy Text. In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 58–68, Vancouver, Canada, 2017. Association for Computational Linguistics.
- [17] Yuan Fang and Ming-Wei Chang. Entity linking on microblogs with spatial and temporal signals. *Transactions of the Association for Computational Linguistics*, 2:259–272, 2014.
- [18] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal Compact Bilinear Pooling for Visual Question Answering and Visual Grounding. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 457–468, 2016.
- [19] Pierre Geurts, Damien Ernst, and Louis Wehenkel. Extremely randomized trees. *Machine learning*, 63(1):3–42, 2006.

- [20] Amir Globerson, Nevena Lazic, Soumen Chakrabarti, Amarnag Subramanya, Michael Ringgaard, and Fernando Pereira. Collective entity resolution with multi-focal attention. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 621–631, 2016.
- [21] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 249–256, 2010.
- [22] Zhaochen Guo and Denilson Barbosa. Entity linking with a unified semantic representation. In *Proceedings of the 23rd International Conference on World Wide Web*, pages 1305–1310. ACM, 2014.
- [23] Hua He, Kevin Gimpel, and Jimmy Lin. Multi-perspective sentence similarity modeling with convolutional neural networks. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1576–1586, 2015.
- [24] Johannes Hoffart, Mohamed Amir Yosef, Iliaria Bordino, Hagen Fürstenu, Manfred Pinkal, Marc Spaniol, Bilyana Taneva, Stefan Thater, and Gerhard Weikum. 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, 2011.
- [25] Wen Hua, Kai Zheng, and Xiaofang Zhou. Microblog entity linking with social temporal context. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, pages 1761–1775. ACM, 2015.
- [26] Hongzhao Huang, Yunbo Cao, Xiaojiang Huang, Heng Ji, and Chin-Yew Lin. Collective tweet wikification based on semi-supervised graph regularization. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 380–390, 2014.
- [27] Allan Jabri, Armand Joulin, and Laurens Van Der Maaten. Revisiting visual question answering baselines. In *European conference on computer vision*, pages 727–739. Springer, 2016.
- [28] Heng Ji, Ralph Grishman, Hoa Trang Dang, Kira Griffitt, and Joe Ellis. Overview of the TAC 2010 knowledge base population track. In *Third Text Analysis Conference (TAC 2010)*, 2010.
- [29] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4565–4574, 2016.
- [30] Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Herve Jégou, and Tomas Mikolov. FastText.zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*, 2016.
- [31] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference On Computer Vision And Pattern Recognition*, pages 3128–3137, 2015.
- [32] Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Skip-thought vectors. In *Advances in Neural Information Processing Systems*, pages 3294–3302, 2015.
- [33] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [34] Xiaohua Liu, Yitong Li, Haocheng Wu, Ming Zhou, Furu Wei, and Yi Lu. Entity linking for tweets. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 1304–1311, 2013.
- [35] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [36] David Milne and Ian H Witten. Learning to link with wikipedia. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management*, pages 509–518. ACM, 2008.

- [37] Seungwhan Moon, Leonardo Neves, and Vitor Carvalho. Multimodal Named Entity Disambiguation for Noisy Social Media Posts. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 2000–2008, 2018.
- [38] Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. Unsupervised Learning of Sentence Embeddings using Compositional n-Gram Features. In *2018 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2018)*, 2018.
- [39] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS 2017 Autodiff Workshop*, 2017.
- [40] Maria Pershina, Yifan He, and Ralph Grishman. Personalized page rank for named entity disambiguation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 238–243, 2015.
- [41] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *Proceedings of NAACL-HLT*, pages 2227–2237, 2018.
- [42] Jinfeng Rao, Hua He, and Jimmy Lin. Noise-contrastive estimation for answer selection with deep neural networks. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 1913–1916. ACM, 2016.
- [43] Lev Ratinov, Dan Roth, Doug Downey, and Mike Anderson. Local and global algorithms for disambiguation to wikipedia. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1375–1384. Association for Computational Linguistics, 2011.
- [44] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [45] Wei Shen, Jianyong Wang, Ping Luo, and Min Wang. Linking named entities in tweets with knowledge base via user interest modeling. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 68–76. ACM, 2013.
- [46] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [47] Thi Quynh Nhi Tran, Hervé Le Borgne, and Michel Crucianu. Aggregating Image and Text Quantized Correlated Components. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, USA, 2016.
- [48] Liwei Wang, Yin Li, Jing Huang, and Svetlana Lazebnik. Learning two-branch neural networks for image-text matching tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):394–407, 2019.
- [49] Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5005–5013, 2016.
- [50] Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 19–27, 2015.