

Multi-level Representations for Fine-Grained Typing of Knowledge Base Entities

Yadollah Yaghoobzadeh and Hinrich Schütze
Center for Information and Language Processing
LMU Munich, Germany
yadollah@cis.lmu.de

Abstract

Entities are essential elements of natural language. In this paper, we present methods for learning multi-level representations of entities on three complementary levels: **character** (character patterns in entity names extracted, e.g., by neural networks), **word** (embeddings of words in entity names) and **entity** (entity embeddings). We investigate state-of-the-art learning methods on each level and find large differences, e.g., for deep learning models, traditional ngram features and the subword model of *fasttext* (Bojanowski et al., 2016) on the **character** level; for *word2vec* (Mikolov et al., 2013) on the **word** level; and for the order-aware model *wang2vec* (Ling et al., 2015a) on the **entity** level.

We confirm experimentally that each level of representation contributes complementary information and a joint representation of all three levels improves the existing embedding based baseline for fine-grained entity typing by a large margin. Additionally, we show that adding information from entity descriptions further improves multi-level representations of entities.

1 Introduction

Knowledge about entities is essential for understanding human language. This knowledge can be attributional (e.g., *canFly*, *isEdible*), type-based (e.g., *isFood*, *isPolitician*, *isDisease*) or relational (e.g., *marriedTo*, *bornIn*). Knowledge bases (KBs) are designed to store this information in a structured way, so that it can be queried easily. Examples of such KBs are Freebase (Bollacker et al., 2008), Wikipedia, Google knowledge graph and

YAGO (Suchanek et al., 2007). For automatic updating and completing the entity knowledge, text resources such as news, user forums, textbooks or any other data in the form of text are important sources. Therefore, information extraction methods have been introduced to extract knowledge about entities from text. In this paper, we focus on the extraction of entity types, i.e., assigning types to – or *typing* – entities. Type information can help extraction of relations by applying constraints on relation arguments.

We address a problem setting in which the following are given: a KB with a set of entities E , a set of types T and a membership function $m : E \times T \mapsto \{0, 1\}$ such that $m(e, t) = 1$ iff entity e has type t ; and a large corpus C in which mentions of E are annotated. In this setting, we address the task of *fine-grained entity typing*: we want to learn a probability function $S(e, t)$ for a pair of entity e and type t and based on $S(e, t)$ infer whether $m(e, t) = 1$ holds, i.e., whether entity e is a member of type t .

We address this problem by learning a multi-level representation for an entity that contains the information necessary for typing it. One important source is the *contexts in which the entity is used*. We can take the standard method of learning embeddings for words and extend it to learning embeddings for entities. This requires the use of an entity linker and can be implemented by replacing all occurrences of the entity by a unique token. We refer to entity embeddings as *entity-level representations*. Previously, entity embeddings have been learned mostly using bag-of-word models like *word2vec* (e.g., by Wang et al. (2014) and Yaghoobzadeh and Schütze (2015)). We show below that order information is critical for high-quality entity embeddings.

Entity-level representations are often uninformative for rare entities, so that using only entity

embeddings is likely to produce poor results. In this paper, we use *entity names* as a source of information that is complementary to entity embeddings. We define an entity name as a noun phrase that is used to refer to an entity. We learn character and word level representations of entity names.

For the *character-level representation*, we adopt different character-level neural network architectures. Our intuition is that there is sub/cross word information, e.g., orthographic patterns, that is helpful to get better entity representations, especially for rare entities. A simple example is that a three-token sequence containing an initial like “P.” surrounded by two capitalized words (“Rolph P. Kugl”) is likely to refer to a person.

We compute the *word-level representation* as the sum of the embeddings of the words that make up the entity name. The sum of the embeddings accumulates evidence for a type/property over all constituents, e.g., a name containing “stadium”, “lake” or “cemetery” is likely to refer to a location. In this paper, we compute our word level representation with two types of word embeddings: (i) using only contextual information of words in the corpus, e.g., by `word2vec` (Mikolov et al., 2013) and (ii) using subword as well as contextual information of words, e.g., by Facebook’s recently released `fasttext` (Bojanowski et al., 2016).

In this paper, we integrate character-level and word-level with entity-level representations to improve the results of previous work on fine-grained typing of KB entities. We also show how descriptions of entities in a KB can be a complementary source of information to our multi-level representation to improve the results of entity typing, especially for rare entities.

Our main contributions in this paper are:

- We propose new methods for learning entity representations on three levels: character-level, word-level and entity-level.
- We show that these levels are complementary and a joint model that uses all three levels improves the state of the art on the task of fine-grained entity typing by a large margin.
- We experimentally show that an order dependent embedding is more informative than its bag-of-word counterpart for entity representation.

We release our dataset and source codes: cistern.cis.lmu.de/figment2/.

2 Related Work

Entity representation. Two main sources of information used for learning entity representation are: (i) links and descriptions in KB, (ii) name and contexts in corpora. We focus on name and contexts in corpora, but we also include (Wikipedia) descriptions. We represent entities on three levels: entity, word and character. Our entity-level representation is similar to work on relation extraction (Wang et al., 2014; Wang and Li, 2016), entity linking (Yamada et al., 2016; Fang et al., 2016), and entity typing (Yaghoobzadeh and Schütze, 2015). Our word-level representation with distributional word embeddings is similarly used to represent entities for entity linking (Sun et al., 2015) and relation extraction (Socher et al., 2013; Wang et al., 2014). Novel entity representation methods we introduce in this paper are representation based on `fasttext` (Bojanowski et al., 2016) subword embeddings, several character-level representations, “order-aware” entity-level embeddings and the combination of several different representations into one multi-level representation.

Character-subword level neural networks. Character-level convolutional neural networks (CNNs) are applied by dos Santos and Zadrozny (2014) to part of speech (POS) tagging, by dos Santos and Guimarães (2015), Ma and Hovy (2016), and Chiu and Nichols (2016) to named entity recognition (NER), by Zhang et al. (2015) and Zhang and LeCun (2015) to sentiment analysis and text categorization, and by Kim et al. (2016) to language modeling (LM). Character-level LSTM is applied by Ling et al. (2015b) to LM and POS tagging, by Lample et al. (2016) to NER, by Ballesteros et al. (2015) to parsing morphologically rich languages, and by Cao and Rei (2016) to learning word embeddings. Bojanowski et al. (2016) learn word embeddings by representing words with the average of their character ngrams (subwords) embeddings. Similarly, Chen et al. (2015) extends `word2vec` for Chinese with joint modeling with characters.

Fine-grained entity typing. Our task is to infer fine-grained types of KB entities. KB completion is an application of this task. Yaghoobzadeh and Schütze (2015)’s FIGMENT system addresses this task with only contextual information; they do not use character-level and word-level features of entity names. Neelakantan and Chang (2015) and Xie et al. (2016) also address a similar task,

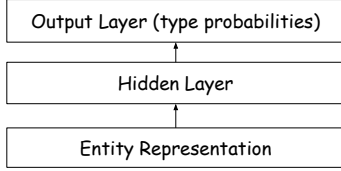


Figure 1: Schematic diagram of our architecture for entity classification. “Entity Representation” ($\vec{v}(e)$) is the (one-level or multi-level) vector representation of entity. Size of output layer is $|T|$.

but they rely on entity descriptions in KBs, which in many settings are not available. The problem of Fine-grained mention typing (FGMT) (Yosef et al., 2012; Ling and Weld, 2012; Yogatama et al., 2015; Del Corro et al., 2015; Shimaoka et al., 2016; Ren et al., 2016) is related to our task. FGMT classifies single *mentions* of named entities to their context dependent types whereas we attempt to identify all types of a KB *entity* from the aggregation of all its mentions. FGMT can still be evaluated in our task by aggregating the mention level decisions but as we will show in our experiments for one system, i.e., FIGER (Ling and Weld, 2012), our entity embedding based models are better in entity typing.

3 Fine-grained entity typing

Given (i) a KB with a set of entities E , (ii) a set of types T , and (iii) a large corpus C in which mentions of E are linked, we address the task of *fine-grained entity typing* (Yaghoobzadeh and Schütze, 2015): predict whether entity e is a member of type t or not. To do so, we use a set of training examples to learn $P(t|e)$: the probability that entity e has type t . These probabilities can be used to assign *new types* to entities covered in the KB as well as typing *unknown entities*.

We learn $P(t|e)$ with a general architecture; see Figure 1. The output layer has size $|T|$. Unit t of this layer outputs the probability for type t . “Entity Representation” ($\vec{v}(e)$) is the vector representation of entity e – we will describe in detail in the rest of this section what forms $\vec{v}(e)$ takes. We model $P(t|e)$ as a multi-label classification, and train a multilayer perceptron (MLP) with one hidden layer:

$$[P(t_1|e) \dots P(t_T|e)] = \sigma(\mathbf{W}_{\text{out}} f(\mathbf{W}_{\text{in}} \vec{v}(e))) \quad (1)$$

where $\mathbf{W}_{\text{in}} \in \mathbb{R}^{h \times d}$ is the weight matrix from

$\vec{v}(e) \in \mathbb{R}^d$ to the hidden layer with size h . f is the rectifier function. $\mathbf{W}_{\text{out}} \in \mathbb{R}^{|T| \times h}$ is the weight matrix from hidden layer to output layer of size $|T|$. σ is the sigmoid function. Our objective is binary cross entropy summed over types:

$$\sum_t - (m_t \log p_t + (1 - m_t) \log (1 - p_t))$$

where m_t is the truth and p_t the prediction.

The key difficulty when trying to compute $P(t|e)$ is in learning a good representation for entity e . We make use of contexts and name of e to represent its feature vector on the three levels of entity, word and character.

3.1 Entity-level representation

Distributional representations or embeddings are commonly used for words. The underlying hypothesis is that words with similar meanings tend to occur in similar contexts (Harris, 1954) and therefore cooccur with similar context words. We can extend the distributional hypothesis to entities (cf. Wang et al. (2014), Yaghoobzadeh and Schütze (2015)): entities with similar meanings tend to have similar contexts. Thus, we can learn a d dimensional embedding $\vec{v}(e)$ of entity e from a corpus in which all mentions of the entity have been replaced by a special identifier. We refer to these entity vectors as the *entity level representation* (ELR).

In previous work, order information of context words (relative position of words in the contexts) was generally ignored and objectives similar to the SkipGram (henceforth: *SKIP*) model were used to learn $\vec{v}(e)$. However, the bag-of-words context is difficult to distinguish for pairs of types like (restaurant,food) and (author,book). This suggests that *using order aware embedding models is important for entities*. Therefore, we apply Ling et al. (2015a)’s extended version of SKIP, Structured SKIP (SSKIP). It incorporates the order of context words into the objective. We compare it with SKIP embeddings in our experiments.

3.2 Word-level representation

Words inside entity names are important sources of information for typing entities. We define the word-level representation (WLR) as the *average of the embeddings of the words* that the entity name contains $\vec{v}(e) = 1/n \sum_{i=1}^n \vec{v}(w_i)$ where $\vec{v}(w_i)$ is the embedding of the i^{th} word of an entity name

of length n . We opt for simple averaging since entity names often consist of a small number of words with clear semantics. Thus, averaging is a promising way of combining the information that each word contributes.

The word embedding, \vec{w} , itself can be learned from models with different granularity levels. Embedding models that consider words as atomic units in the corpus, e.g., SKIP and SSKIP, are word-level. On the other hand, embedding models that represent words with their character ngrams, e.g., *fasttext* (Bojanowski et al., 2016), are subword-level. Based on this, we consider and evaluate **word-level WLR (WWLR)** and **subword-level WLR (SWLR)** in this paper.¹

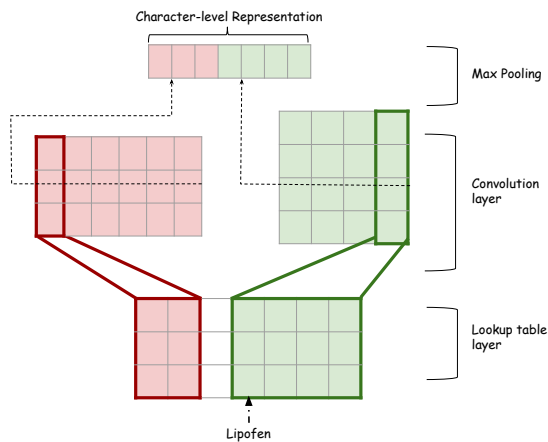


Figure 2: Example architecture for the character-level CNN with max pooling. The input is “Lipofen”. Character embedding size is three. There are three filters of width 2 and four filters of width 4.

3.3 Character-level representation

For computing the *character level representation* (CLR), we design models that try to type an entity based on the sequence of characters of its name. Our hypothesis is that names of entities of a specific type often have similar character patterns. Entities of type ETHNICITY often end in “ish” and “ian”, e.g., “Spanish” and “Russian”. Entities of type MEDICINE often end in “en”: “Lipofen”, “acetaminophen”. Also, some types tend to have specific cross-word shapes in their entities, e.g.,

¹Subword models have properties of both character-level models (subwords are character ngrams) and of word-level models (they do not cross boundaries between words). They probably could be put in either category, but in our context fit the word-level category better because we see the granularity level with respect to the entities and not words.

PERSON names usually consist of two words, or MUSIC names are usually long, containing several words.

The first layer of the character-level models is a *lookup table* that maps each character to an embedding of size d_c . These embeddings capture similarities between characters, e.g., similarity in type of phoneme encoded (consonant/vowel) or similarity in case (lower/upper). The output of the lookup layer for an entity name is a matrix $C \in \mathbb{R}^{l \times d_c}$ where l is the maximum length of a name and all names are padded to length l . This length l includes special start/end characters that bracket the entity name.

We experiment with four architectures to produce character-level representations in this paper: FORWARD (direct forwarding of character embeddings), CNNs, LSTMs and BiLSTMs. The output of each architecture then takes the place of the entity representation $\vec{v}(e)$ in Figure 1.

FORWARD simply concatenates all rows of matrix C ; thus, $\vec{v}(e) \in \mathbb{R}^{d_c * l}$.

The **CNN** uses k filters of different window widths w to narrowly convolve C . For each filter $H \in \mathbb{R}^{d_c \times w}$, the result of the convolution of H over matrix C is feature map $f \in \mathbb{R}^{l-w+1}$:

$$f[i] = \text{rectifier}(C_{[:,i:i+w-1]} \odot H + b)$$

where *rectifier* is the activation function, b is the bias, $C_{[:,i:i+w-1]}$ are the columns i to $i + w - 1$ of C , $1 \leq w \leq 10$ are the window widths we consider and \odot is the sum of element-wise multiplication. Max pooling then gives us one feature for each filter. The concatenation of all these features is our representation: $\vec{v}(e) \in \mathbb{R}^k$. An example CNN architecture is show in Figure 2.

The input to the **LSTM** is the character sequence in matrix C , i.e., $x_1, \dots, x_l \in \mathbb{R}^{d_c}$. It generates the state sequence h_1, \dots, h_{l+1} and the output is the last state $\vec{v}(e) \in \mathbb{R}^{d_h}$.²

The **BiLSTM** consists of two LSTMs, one going forward, one going backward. The first state of the backward LSTM is initialized as h_{l+1} , the last state of the forward LSTM. The BiLSTM entity representation is the concatenation of last states of forward and backward LSTMs, i.e., $\vec{v}(e) \in \mathbb{R}^{2*d_h}$.

3.4 Multi-level representations

Our different levels of representations can give complementary information about entities.

²We use Blocks (van Merriënboer et al., 2015).

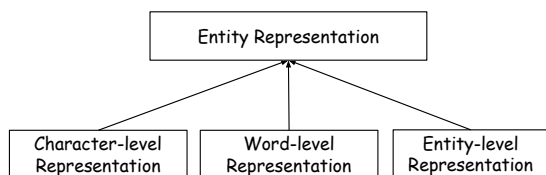


Figure 3: Multi-level representation

WLR and CLR. Both WLR models, SWLR and WWLR, do not have access to the cross-word character ngrams of entity names while CLR models do. Also, CLR is task specific by training on the entity typing dataset while WLR is generic. On the other hand, WWLR and SWLR models have access to information that CLR ignores: the tokenization of entity names into words and embeddings of these words. It is clear that words are particularly important character sequences since they often correspond to linguistic units with clearly identifiable semantics – which is not true for most character sequences. For many entities, the words they contain are a better basis for typing than the character sequence. For example, even if “nectarine” and “compote” did not occur in any names in the training corpus, we can still learn good word embeddings from their non-entity occurrences. This then allows us to correctly type the entity “Aunt Mary’s Nectarine Compote” as FOOD based on the sum of the word embeddings.

WLR/CLR and ELR. Representations from entity names, i.e., WLR and CLR, by themselves are limited because many classes of names can be used for different types of entities; e.g., person names do not contain hints as to whether they are referring to a politician or athlete. In contrast, the ELR embedding is based on an entity’s contexts, which are often informative for each entity and can distinguish politicians from athletes. On the other hand, not all entities have sufficiently many informative contexts in the corpus. For these entities, their name can be a complementary source of information and character/word level representations can increase typing accuracy.

Thus, we introduce joint models that use combinations of the three levels. Each multi-level model concatenates several levels. We train the constituent embeddings as follows. WLR and ELR are computed as described above and are not changed during training. CLR – produced by one of the character-level networks described above – is initialized randomly and then tuned during

training. Thus, it can focus on complementary information related to the task that is not already present in other levels. The schematic diagram of our multi-level representation is shown in Figure 3.

4 Experimental setup and results

4.1 Setup

Entity datasets and corpus. We address the task of fine-grained entity typing and use Yaghoobzadeh and Schütze (2015)’s FIGMENT dataset³ for evaluation. The FIGMENT corpus is part of a version of ClueWeb in which Freebase entities are annotated using FACC1 (URL, 2016b; Gabrilovich et al., 2013). The FIGMENT entity datasets contain 200,000 Freebase entities that were mapped to 102 FIGER types (Ling and Weld, 2012). We use the same train (50%), dev (20%) and test (30%) partitions as Yaghoobzadeh and Schütze (2015) and extract the names from mentions of dataset entities in the corpus. We take the most frequent name for dev and test entities and three most frequent names for train (each one tagged with entity types).

Adding parent types to refine entity dataset. FIGMENT ignores that FIGER is a proper hierarchy of types; e.g., while HOSPITAL is a subtype of BUILDING according to FIGER, there are entities in FIGMENT that are hospitals, but not buildings.⁴ Therefore, we modified the FIGMENT dataset by adding for each assigned type (e.g., HOSPITAL) its parents (e.g., BUILDING). This makes FIGMENT more consistent and eliminates spurious false negatives (BUILDING in the example).

We now describe our **baselines**: (i) BOW & NSL: hand-crafted features, (ii) FIGMENT (Yaghoobzadeh and Schütze, 2015) and (iii) adapted version of FIGER (Ling and Weld, 2012).

We implement the following two feature sets from the literature as a *hand-crafted baseline* for our character and word level models. (i) *BOW*: individual words of entity name (both as-is and lowercased); (ii) *NSL* (ngram-shape-length): shape and length of the entity name (cf. Ling and Weld (2012)), character n -grams, $1 \leq n \leq n_{\max}$, $n_{\max} = 5$ (we also tried $n_{\max} = 7$, but results were worse on dev) and normalized character n -grams: lowercased, digits replaced by “7”, punctuation replaced by “.”. These features are represented as a sparse

³cistern.cis.lmu.de/figment/

⁴See github.com/xiaoling/figer for FIGER

binary vector $\vec{v}(e)$ that is input to the architecture in Figure 1.

FIGMENT is the model for entity typing presented by Yaghoobzadeh and Schütze (2015). The authors only use entity-level representations for entities trained by SkipGram, so the *FIGMENT* baseline corresponds to the entity-level result shown as ELR(SKIP) in the tables.

The third baseline is using an existing mention-level entity typing system, *FIGER* (Ling and Weld, 2012). *FIGER* uses a wide variety of features on different levels (including parsing-based features) from contexts of entity mentions as well as the mentions themselves and returns a score for each mention-type instance in the corpus. We provide the ClueWeb/FACC1 segmentation of entities, so *FIGER* does not need to recognize entities.⁵ We use the trained model provided by the authors and normalize *FIGER* scores using softmax to make them comparable for aggregation. We experimented with different aggregation functions (including maximum and k-largest-scores for a type), but we use the average of scores since it gave us the best result on dev. We call this baseline AGG-*FIGER*.

Distributional embeddings. For WWLR and ELR, we use SkipGram model in `word2vec` and SSkip model in `wang2vec` (Ling et al., 2015a) to learn embeddings for words, entities and types. To obtain embeddings for all three in the same space, we process ClueWeb/FACC1 as follows. For each sentence s , we add three copies: s itself, a copy of s in which each entity is replaced with its Freebase identifier (MID) and a copy in which each entity (not test entities though) is replaced with an ID indicating its notable type. The resulting corpus contains around 4 billion tokens and 1.5 billion types.

We run SKIP and SSkip with the same setup (200 dimensions, 10 negative samples, window size 5, word frequency threshold of 100)⁶ on this corpus to learn embeddings for words, entities and *FIGER* types. Having entities and types in the same vector space, we can add another feature vector $\vec{v}(e) \in \mathbb{R}^{|T|}$ (referred to as TC below): for each entity, we compute cosine similarity of its entity vector with all type vectors.

For SWLR, we use `fasttext`⁷ to learn word

⁵Mention typing is separated from recognition in *FIGER* model. So it can use our segmentation of entities.

⁶The threshold does not apply for MIDs.

⁷github.com/facebookresearch/fastText

embeddings from the ClueWeb/FACC1 corpus. We use similar settings as our WWLR SKIP and SSkip embeddings and keep the defaults of other hyperparameters. Since the trained model of `fasttext` is applicable for new words, we apply the model to get embeddings for the filtered rare words as well.

model	hyperparameters
CLR(FF)	$d_c = 15, h_{mlp} = 600$
CLR(LSTM)	$d_c = 70, d_h = 70, h_{mlp} = 300$
CLR(BiLSTM)	$d_c = 50, d_h = 50, h_{mlp} = 200$
CLR(CNN)	$d_c = 10, w = [1, \dots, 8]$ $n = 100, h_{mlp} = 800$
CLR(NSL)	$h_{mlp} = 800$
BOW	$h_{mlp} = 200$
BOW+CLR(NSL)	$h_{mlp} = 300$
WWLR	$h_{mlp} = 400$
SWLR	$h_{mlp} = 400$
WWLR+CLR(CNN)	$w = [1, \dots, 7]$ $d_c = 10, n = 50, h_{mlp} = 700$
SWLR+CLR(CNN)	$w = [1, \dots, 7]$ $d_c = 10, n = 50, h_{mlp} = 700$
ELR(SKIP)	$h_{mlp} = 400$
ELR(SSkip)	$h_{mlp} = 400$
ELR+CLR	$d_c = 10, w = [1, \dots, 7]$ $n = 100, h_{mlp} = 700$
ELR+WWLR	$h_{mlp} = 600$
ELR+SWLR	$h_{mlp} = 600$
ELR+WWLR+CLR	$d_c = 10, w = [1, \dots, 7]$ $n = 50, h_{mlp} = 700$
ELR+SWLR+CLR	$d_c = 10, w = [1, \dots, 7]$ $n = 50, h_{mlp} = 700$
ELR+WWLR+CNN+TC	$d_c = 10, w = [1, \dots, 7]$ $n = 50, h_{mlp} = 900$
ELR+SWLR+CNN+TC(MuLR)	$d_c = 10, w = [1, \dots, 7]$ $n = 50, h_{mlp} = 900$
AVG-DES	$h_{mlp} = 400$
MuLR+AVG-DES	$d_c = 10, w = [1, \dots, 7]$ $n = 50, h_{mlp} = 1000$

Table 1: Hyperparameters of different models. w is the filter size. n is the number of CNN feature maps for each filter size. d_c is the character embedding size. d_h is the LSTM hidden state size. h_{mlp} is the number of hidden units in the MLP.

Our **hyperparameter values** are given in Table 1. The values are optimized on dev. We use AdaGrad and minibatch training. For each experiment, we select the best model on dev.

We use these **evaluation measures**: (i) accuracy: an entity is correct if all its types and no incorrect types are assigned to it; (ii) micro average F_1 : F_1 of all type-entity assignment decisions; (iii) entity macro average F_1 : F_1 of types assigned to an entity, averaged over entities; (iv) type macro average F_1 : F_1 of entities assigned to a type, averaged over types.

The assignment decision is based on thresholding the probability function $P(t|e)$. For each model and type, we select the threshold that maximizes F_1 of entities assigned to the type on dev.

	all entities			head entities			tail entities		
	acc	mic	mac	acc	mic	mac	acc	mic	mac
1 MFT	.000	.041	.041	.000	.044	.044	.000	.038	.038
2 CLR(FORWARD)	.066	.379	.352	.067	.342	.369	.061	.374	.350
3 CLR(LSTM)	.121	.425	.396	.122	.433	.390	.116	.408	.391
4 CLR(BiLSTM)	.133	.440	.404	.129	.443	.394	.135	.428	.404
5 CLR(NSL)	.164	.484	.464	.157	.470	.443	.173	.483	.472
6 CLR(CNN)	.177	.494	.468	.171	.484	.450	.187	.489	.474
7 BOW	.113	.346	.379	.109	.323	.353	.120	.356	.396
8 WWLR(SKIP)	.214	.581	.531	.293	.660	.634	.173	.528	.478
9 WWLR(SSKIP)	.223	.584	.543	.306	.667	.642	.183	.533	.494
10 SWLR	.236	.590	.554	.301	.665	.632	.209	.551	.522
11 BOW+CLR(NSL)	.156	.487	.464	.157	.480	.452	.159	.485	.469
12 WWLR+CLR(CNN)	.257	.603	.568	.317	.668	.637	.235	.567	.538
13 SWLR+CLR(CNN)	.241	.594	.561	.295	.659	.628	.227	.560	.536
14 ELR(SKIP)	.488	.774	.741	.551	.834	.815	.337	.621	.560
15 ELR(SSKIP)	.515	.796	.763	.560	.839	.819	.394	.677	.619
16 AGG-FIGER	.320	.694	.660	.396	.762	.724	.220	.593	.568
17 ELR+CLR	.554	.816	.788	.580	.844	.825	.467	.733	.690
18 ELR+WWLR	.557	.819	.793	.582	.846	.827	.480	.749	.708
19 ELR+SWLR	.558	.820	.796	.584	.846	.829	.480	.751	.714
20 ELR+WWLR+CLR	.568	.823	.798	.590	.847	.829	.491	.755	.716
21 ELR+SWLR+CLR	.569	.824	.801	.590	.849	.831	.497	.760	.724
22 ELR+WWLR+CLR+TC	.572	.824	.801	.594	.849	.831	.499	.759	.722
23 ELR+SWLR+CLR+TC	.575	.826	.802	.597	.851	.831	.508	.762	.727

Table 2: Accuracy (acc), micro (mic) and macro (mac) F_1 on test for all, head and tail entities.

4.2 Results

Table 2 gives results on the test entities for all (about 60,000 entities), head (frequency > 100; about 12,200) and tail (frequency < 5; about 10,000). *MFT* (line 1) is the most frequent type baseline that ranks types according to their frequency in the train entities. Each level of representation is separated with dashed lines, and – unless noted otherwise – the best of each level is joined in multi level representations.⁸

Character-level models are on lines 2-6. The order of systems is: CNN > NSL > BiLSTM > LSTM > FORWARD. The results show that complex neural networks are more effective than simple forwarding. BiLSTM works better than LSTM, confirming other related work. CNNs probably work better than LSTMs because there are few complex non-local dependencies in the sequence, but many important local features. CNNs with maxpooling can more straightforwardly capture local and position-independent features. CNN also beats NSL baseline; a possible reason is that CNN – an automatic method of feature learning

⁸For accuracy measure: in the following ordered lists of sets, $A < B$ means that all members (row numbers in Table 2) of A are significantly worse than all members of B : $\{1\} < \{2\} < \{3, \dots, 11\} < \{12, 13\} < \{14, 15, 16\} < \{17, \dots, 23\}$. Test of equal proportions, $\alpha < 0.05$. See Table 6 in the appendix for more details.

types:	all	head	tail
AGG-FIGER	.566	.702	.438
ELR	.621	.784	.480
MuLR	.669	.811	.541

Table 3: Type macro average F_1 on test for all, head and tail types. MuLR = ELR+SWLR+CLR+TC

	all known?	
	yes	no
CLR(NSL)	.484	.521
CLR(CNN)	.494	.524
BOW	.346	.435
SWLR	.590	.612
BOW+NSL	.497	.535
SWLR+CLR(CNN)	.594	.616

Table 4: Micro F_1 on test of character, word level models for all, known (“known? yes”) and unknown (“known? no”) entities.

– is more robust than hand engineered feature based NSL. We show more detailed results in Section 4.3.

Word-level models are on lines 7-10. BOW performs worse than WWLR because it cannot deal well with sparseness. SSKIP uses word order information in WWLR and performs better than SKIP. SWLR uses subword information and performs better than WWLR, especially for tail entities. Integrating subword information improves the quality of embeddings for rare words and mitigates the problem of unknown words.

Joint word-character level models are on lines 11-13. WWLR+CLR(CNN) and SWLR+CLR(CNN) beat the component models. This confirms our underlying assumption in designing the complementary multi-level models. BOW problem with rare words does not allow its joint model with NSL to work better than NSL. WWLR+CLR(CNN) works better than BOW+CLR(NSL) by 10% micro F_1 , again due to the limits of BOW compared to WWLR. Interestingly WWLR+CLR works better than SWLR+CLR and this suggests that WWLR is indeed richer than SWLR when CLR mitigates its problem with rare/unknown words

Entity-level models are on lines 14–15 and they are better than all previous models on lines

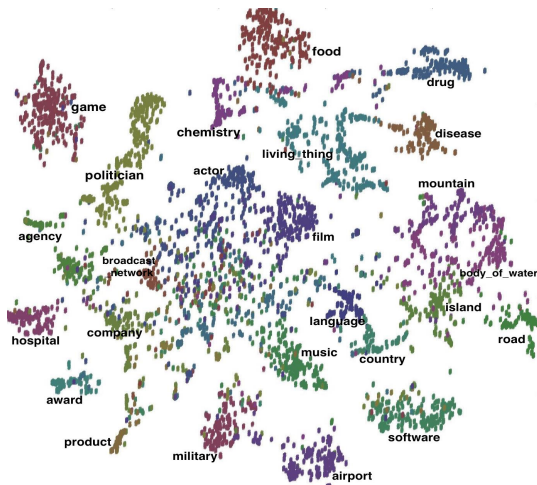


Figure 4: t-SNE result of entity-level representations

1–13. This shows the power of entity-level embeddings. In Figure 4, a t-SNE (Van der Maaten and Hinton, 2008) visualization of ELR(SKIP) embeddings using different colors for entity types shows that entities of the same type are clustered together. SSKIP works marginally better than SKIP for ELR, especially for tail entities, confirming our hypothesis that order information is important for a good distributional entity representation. This is also confirming the results of Yaghoobzadeh and Schütze (2016), where they also get better entity typing results with SSKIP compared to SKIP. They propose to use entity typing as an extrinsic evaluation for embedding models.

Joint entity, word, and character level models are on lines 16-23. The AGG-FIGER baseline works better than the systems on lines 1-13, but worse than ELRs. This is probably due to the fact that AGG-FIGER is optimized for mention typing and it is trained using distant supervision assumption. Parallel to our work, Yaghoobzadeh et al. (2017) optimize a mention typing model for our entity typing task by introducing multi instance learning algorithms, resulting comparable performance to ELR(SKIP). We will investigate their method in future.

Joining CLR with ELR (line 17) results in large improvements, especially for tail entities (5% micro F_1). This demonstrates that for rare entities, contextual information is often not sufficient for an informative representation, hence name features are important. This is also true for the joint models of WWLR/SWLR and ELR (lines 18-19). Joining WWLR works better than

CLR, and SWLR is slightly better than WWLR. Joint models of WWLR/SWLR with ELR+CLR gives more improvements, and SWLR is again slightly better than WWLR. ELR+WWLR+CLR and ELR+SWLR+CLR, are better than their two-level counterparts, again confirming that these levels are complementary.

We get a further boost, especially for tail entities, by also including TC (type cosine) in the combinations (lines 22-23). This demonstrates the potential advantage of having a common representation space for entities and types. Our best model, ELR+SWLR+CLR+TC (line 22), which we refer to as MuLR in the other tables, beats our initial baselines (ELR and AGG-FIGER) by large margins, e.g., in tail entities improvements are more than 8% in micro F1.

Table 3 shows **type macro F_1** for MuLR (ELR+SWLR+CLR+TC) and two baselines. There are 11 head types (those with ≥ 3000 train entities) and 36 tail types (those with < 200 train entities). These results again confirm the superiority of our multi-level models over the baselines: AGG-FIGER and ELR, the best single-level model baseline.

4.3 Analysis

Unknown vs. known entities. To analyze the complementarity of character and word level representations, as well as more fine-grained comparison of our models and the baselines, we divide test entities into *known entities* – at least one word of the entity’s name appears in a train entity – and *unknown entities* (the complement). There are 45,000 (resp. 15,000) known (resp. unknown) test entities.

Table 4 shows that the CNN works only slightly better (by 0.3%) than NSL on known entities, but works much better on unknown entities (by 3.3%), justifying our preference for deep learning CLR models. As expected, BOW works relatively well for known entities and really poorly for unknown entities. SWLR beats CLR models as well as BOW. The reason is that in our setup, word embeddings are induced on the entire corpus using an unsupervised algorithm. Thus, even for many words that did not occur in train, SWLR has access to informative representations of words. The joint model, SWLR+CLR(CNN), is significantly better than BOW+CLR(NSL) again due to limits of BOW. SWLR+CLR(CNN) is better than SWLR

in unknown entities.

Case study of LIVING-THING. To understand the interplay of different levels better, we perform a case study of the type LIVING-THING. Living beings that are not humans belong to this type.

WLRs incorrectly assign “Walter Leaf” (PERSON) and “Along Came A Spider” (MUSIC) to LIVING-THING because these names contain a word referring to a LIVING-THING (“leaf”, “spider”), but the entity itself is not a LIVING-THING. In these cases, the averaging of embeddings that WLR performs is misleading. The CLR(CNN) types these two entities correctly because their names contain character ngram/shape patterns that are indicative of PERSON and MUSIC.

ELR incorrectly assigns “Zumpango” (CITY) and “Lake Kasumigaura” (LOCATION) to LIVING-THING because these entities are rare and words associated with living things (e.g., “wildlife”) dominate in their contexts. However, CLR(CNN) and WLR enable the joint model to type the two entities correctly: “Zumpango” because of the informative suffix “-go” and “Lake Kasumigaura” because of the informative word “Lake”.

While some of the **remaining errors** of our best system MuLR are due to the inherent difficulty of entity typing (e.g., it is difficult to correctly type a one-word entity that occurs once and whose name is not informative), many other errors are due to artifacts of our setup. First, ClueWeb/FACCI is the result of an automatic entity linking system and any entity linking errors propagate to our models. Second, due to the incompleteness of Freebase (Yaghoobzadeh and Schütze, 2015), many entities in the FIGMENT dataset are incompletely annotated, resulting in correctly typed entities being evaluated as incorrect.

Adding another source: description-based embeddings. While in this paper, we focus on the contexts and names of entities, there is a textual source of information about entities in KBs which we can also make use of: descriptions of entities. We extract Wikipedia descriptions of FIGMENT entities filtering out the entities ($\sim 40,000$ out of $\sim 200,000$) without description.

We then build a simple entity representation by averaging the embeddings of the top k words (wrt tf-idf) of the description (henceforth, AVG-DES).⁹ This representation is used as input in Figure 1 to train the MLP. We also train our best multi-

⁹ $k = 20$ gives the best results on dev.

entities:	all	head	tail
AVG-DES	.773	.791	.745
MuLR	.825	.846	.757
MuLR+AVG-DES	.873	.877	.852

Table 5: Micro average F_1 results of MuLR and description based model and their joint.

level model as well as the joint of the two on this smaller dataset. Since the descriptions are coming from Wikipedia, we use 300-dimensional Glove (URL, 2016a) embeddings pretrained on Wikipedia+Gigaword to get more coverage of words. For MuLR, we still use the embeddings we trained before.

Results are shown in Table 5. While for head entities, MuLR works marginally better, the difference is very small in tail entities. The joint model of the two (by concatenation of vectors) improves the micro F1, with clear boost for tail entities. This suggests that for tail entities, the contextual and name information is not enough by itself and some keywords from descriptions can be really helpful. Integrating more complex description-based embeddings, e.g., by using CNN (Xie et al., 2016), may improve the results further. We leave it for future work.

5 Conclusion

In this paper, we have introduced representations of entities on different levels: character, word and entity. The character level representation is learned from the entity name. The word level representation is computed from the embeddings of the words w_i in the entity name where the embedding of w_i is derived from the corpus contexts of w_i . The entity level representation of entity e_i is derived from the corpus contexts of e_i . Our experiments show that each of these levels contributes complementary information for the task of fine-grained typing of entities. The joint model of all three levels beats the state-of-the-art baseline by large margins. We further showed that extracting some keywords from Wikipedia descriptions of entities, when available, can considerably improve entity representations, especially for rare entities. We believe that our findings can be transferred to other tasks where entity representation matters.

Acknowledgments. This work was supported by DFG (SCHU 2246/8-2).

References

- Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2015. Improved transition-based parsing by modeling characters instead of words with lstms. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 349–359, Lisbon, Portugal, September. Association for Computational Linguistics.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. Enriching word vectors with subword information. *CoRR*, abs/1607.04606.
- Kurt D. Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2008, Vancouver, BC, Canada, June 10-12, 2008*, pages 1247–1250.
- Kris Cao and Marek Rei. 2016. A joint model for word embedding and word morphology. In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pages 18–26, Berlin, Germany, August. Association for Computational Linguistics.
- Xinxiong Chen, Lei Xu, Zhiyuan Liu, Maosong Sun, and Huan-Bo Luan. 2015. Joint learning of character and word embeddings. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 1236–1242.
- Jason Chiu and Eric Nichols. 2016. Named entity recognition with bidirectional lstm-cnns. *Transactions of the Association for Computational Linguistics*, 4:357–370.
- Luciano Del Corro, Abdalghani Abujabal, Rainer Gemulla, and Gerhard Weikum. 2015. Finet: Context-aware fine-grained named entity typing. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 868–878, Lisbon, Portugal, September. Association for Computational Linguistics.
- Cícero Nogueira dos Santos and Victor Guimarães. 2015. Boosting named entity recognition with neural character embeddings. *CoRR*, abs/1505.05008.
- Cícero Nogueira dos Santos and Bianca Zadrozny. 2014. Learning character-level representations for part-of-speech tagging. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 1818–1826.
- Wei Fang, Jianwen Zhang, Dilin Wang, Zheng Chen, and Ming Li. 2016. Entity disambiguation by knowledge and text jointly embedding. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 260–269, Berlin, Germany, August. Association for Computational Linguistics.
- Evgeniy Gabrilovich, Michael Ringgaard, and Amarnag Subramanya. 2013. Facc1: Freebase annotation of cluweb corpora.
- Zellig S. Harris. 1954. Distributional structure. *Word*, 10:146–162.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander M. Rush. 2016. Character-aware neural language models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 2741–2749.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California, June. Association for Computational Linguistics.
- Xiao Ling and Daniel S. Weld. 2012. Fine-grained entity recognition. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada*.
- Wang Ling, Chris Dyer, Alan W Black, and Isabel Trancoso. 2015a. Two/too simple adaptations of word2vec for syntax problems. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1299–1304, Denver, Colorado, May–June. Association for Computational Linguistics.
- Wang Ling, Chris Dyer, Alan W Black, Isabel Trancoso, Ramon Fernandez, Silvio Amir, Luis Marujo, and Tiago Luis. 2015b. Finding function in form: Compositional character models for open vocabulary word representation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1520–1530, Lisbon, Portugal, September. Association for Computational Linguistics.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1064–1074, Berlin, Germany, August. Association for Computational Linguistics.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781.
- Arvind Neelakantan and Ming-Wei Chang. 2015. Inferring missing entity type instances for knowledge base completion: New dataset and methods. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 515–525, Denver, Colorado, May–June. Association for Computational Linguistics.

- Xiang Ren, Wenqi He, Meng Qu, Clare R. Voss, Heng Ji, and Jiawei Han. 2016. Label noise reduction in entity typing by heterogeneous partial-label embedding. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, pages 1825–1834.
- Sonse Shimaoka, Pontus Stenetorp, Kentaro Inui, and Sebastian Riedel. 2016. An attentive neural architecture for fine-grained entity type classification. pages 69–74, June.
- Richard Socher, Danqi Chen, Christopher D. Manning, and Andrew Y. Ng. 2013. Reasoning with neural tensor networks for knowledge base completion. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States.*, pages 926–934.
- Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007*, pages 697–706.
- 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 *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 1333–1339.
- URL. 2016a. Glove project. <http://nlp.stanford.edu/projects/glove>.
- URL. 2016b. Lemur project. <http://lemurproject.org/clueweb12/FACC1>.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(2579-2605):85.
- Bart van Merriënboer, Dzmitry Bahdanau, Vincent Dumoulin, Dmitriy Serdyuk, David Warde-Farley, Jan Chorowski, and Yoshua Bengio. 2015. Blocks and fuel: Frameworks for deep learning. *CoRR*, abs/1506.00619.
- Zhigang Wang and Juan-Zi Li. 2016. Text-enhanced representation learning for knowledge graph. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 1293–1299.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph and text jointly embedding. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1591–1601, Doha, Qatar, October. Association for Computational Linguistics.
- Ruobing Xie, Zhiyuan Liu, Jia Jia, Huanbo Luan, and Maosong Sun. 2016. Representation learning of knowledge graphs with entity descriptions. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 2659–2665.
- Yadollah Yaghoobzadeh and Hinrich Schütze. 2015. Corpus-level fine-grained entity typing using contextual information. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 715–725, Lisbon, Portugal, September. Association for Computational Linguistics.
- Yadollah Yaghoobzadeh and Hinrich Schütze. 2016. Intrinsic subspace evaluation of word embedding representations. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 236–246, Berlin, Germany, August. Association for Computational Linguistics.
- Yadollah Yaghoobzadeh, Heike Adel, and Hinrich Schütze. 2017. Noise mitigation for neural entity typing and relation extraction. In *EACL*, Valencia, Spain.
- Ikuya Yamada, Hiroyuki Shindo, Hideaki Takeda, and Yoshiyasu Takefuji. 2016. Joint learning of the embedding of words and entities for named entity disambiguation. pages 250–259, August.
- Dani Yogatama, Daniel Gillick, and Nevena Lazic. 2015. Embedding methods for fine grained entity type classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 291–296, Beijing, China, July. Association for Computational Linguistics.
- Mohamed Amir Yosef, Sandro Bauer, Johannes Hoffart, Marc Spaniol, and Gerhard Weikum. 2012. HYENA: hierarchical type classification for entity names. In *COLING 2012, 24th International Conference on Computational Linguistics, Proceedings of the Conference: Posters, 8-15 December 2012, Mumbai, India*, pages 1361–1370.
- Xiang Zhang and Yann LeCun. 2015. Text understanding from scratch. *CoRR*, abs/1502.01710.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. pages 649–657.

A Supplementary Material

All entities		01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23
01	MFT	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
02	CLR(FORWARD)	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
03	CLR(LSTM)	*	*	0	0	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
04	CLR(BiLSTM)	*	*	*	0	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
05	CLR(CNN)	*	*	*	*	0	*	*	0	0	0	*	*	0	0	0	0	0	0	0	0	0	0	0
06	CLR(NSL)	*	*	*	*	0	0	*	0	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0
07	BOW	*	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
08	WWLR(SkipG)	*	*	*	*	*	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
09	WWLR(SSkipG)	*	*	*	*	*	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
10	SWLR	*	*	*	*	*	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
11	BOW+CLR(NSL)	*	*	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
12	WWLR+CLR(CNN)	*	*	*	*	*	*	*	*	*	*	0	*	0	0	0	0	0	0	0	0	0	0	0
13	SWLR+CLR(CNN)	*	*	*	*	*	*	*	*	*	*	0	0	0	0	0	0	0	0	0	0	0	0	0
14	ELR(SkipG)	*	*	*	*	*	*	*	*	*	*	*	*	0	0	*	*	*	*	*	*	*	*	*
15	ELR(SSkipG)	*	*	*	*	*	*	*	*	*	*	*	*	*	0	*	*	*	*	*	*	*	*	*
16	AGG-FIGER	*	*	*	*	*	*	*	*	*	*	*	*	*	0	0	0	0	0	0	0	0	0	0
17	ELR+CLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
18	ELR+WWLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
19	ELR+SWLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
20	ELR+WWLR+CLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
21	ELR+SWLR+CLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
22	ELR+WWLR+CLR+TC	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
23	ELR+SWLR+CLR+TC	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*

Head entities		01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23
01	MFT	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
02	CLR(FORWARD)	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
03	CLR(LSTM)	*	*	0	0	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
04	CLR(BiLSTM)	*	*	0	0	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
05	CLR(CNN)	*	*	*	*	0	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
06	CLR(NSL)	*	*	*	*	0	0	*	0	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0
07	BOW	*	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
08	WWLR(SkipG)	*	*	*	*	*	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
09	WWLR(SSkipG)	*	*	*	*	*	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
10	SWLR	*	*	*	*	*	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
11	BOW+CLR(NSL)	*	*	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
12	WWLR+CLR(CNN)	*	*	*	*	*	*	*	*	*	*	0	*	0	0	0	0	0	0	0	0	0	0	0
13	SWLR+CLR(CNN)	*	*	*	*	*	*	*	*	*	*	0	0	0	0	0	0	0	0	0	0	0	0	0
14	ELR(SkipG)	*	*	*	*	*	*	*	*	*	*	*	*	0	0	*	*	*	*	*	*	*	*	*
15	ELR(SSkipG)	*	*	*	*	*	*	*	*	*	*	*	*	*	0	*	*	*	*	*	*	*	*	*
16	AGG-FIGER	*	*	*	*	*	*	*	*	*	*	*	*	*	0	0	0	0	0	0	0	0	0	0
17	ELR+CLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
18	ELR+WWLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
19	ELR+SWLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
20	ELR+WWLR+CLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
21	ELR+SWLR+CLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
22	ELR+WWLR+CLR+TC	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
23	ELR+SWLR+CLR+TC	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*

Tail entities		01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23
01	MFT	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
02	CLR(FORWARD)	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
03	CLR(LSTM)	*	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
04	CLR(BiLSTM)	*	*	*	0	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
05	CLR(CNN)	*	*	*	*	0	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
06	CLR(NSL)	*	*	*	*	0	0	*	0	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0
07	BOW	*	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
08	WWLR(SkipG)	*	*	*	*	0	0	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
09	WWLR(SSkipG)	*	*	*	*	0	0	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
10	SWLR	*	*	*	*	*	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0
11	BOW+CLR(NSL)	*	*	*	*	0	0	*	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
12	WWLR+CLR(CNN)	*	*	*	*	*	*	*	*	*	*	0	0	0	0	0	0	0	0	0	0	0	0	0
13	SWLR+CLR(CNN)	*	*	*	*	*	*	*	*	*	*	0	0	0	0	0	0	0	0	0	0	0	0	0
14	ELR(SkipG)	*	*	*	*	*	*	*	*	*	*	*	*	*	0	0	*	*	*	*	*	*	*	*
15	ELR(SSkipG)	*	*	*	*	*	*	*	*	*	*	*	*	*	*	0	*	*	*	*	*	*	*	*
16	AGG-FIGER	*	*	*	*	*	*	*	*	*	*	*	*	*	*	0	0	0	0	0	0	0	0	0
17	ELR+CLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
18	ELR+WWLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
19	ELR+SWLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
20	ELR+WWLR+CLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
21	ELR+SWLR+CLR	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
22	ELR+WWLR+CLR+TC	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*
23	ELR+SWLR+CLR+TC	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*	*

Table 6: Significance-test results for accuracy measure for all, head and tail entities. If the result for the model in a row is significantly larger than the result for the model in a column, then the value in the corresponding (row,column) is * and otherwise is 0.