

Neural Fine-Grained Entity Type Classification with Hierarchy-Aware Loss

Peng Xu

Department of Computing Science
University of Alberta
Edmonton, Canada
pxu4@ualberta.ca

Denilson Barbosa

Department of Computing Science
University of Alberta
Edmonton, Canada
denilson@ualberta.ca

Abstract

The task of Fine-grained Entity Type Classification (FETC) consists of assigning types from a hierarchy to entity mentions in text. Existing methods rely on distant supervision and are thus susceptible to noisy labels that can be *out-of-context* or *overly-specific* for the training sentence. Previous methods that attempt to address these issues do so with heuristics or with the help of hand-crafted features. Instead, we propose an end-to-end solution with a neural network model that uses a variant of cross-entropy loss function to handle *out-of-context* labels, and hierarchical loss normalization to cope with *overly-specific* ones. Also, previous work solve FETC a multi-label classification followed by ad-hoc post-processing. In contrast, our solution is more elegant: we use public word embeddings to train a single-label that jointly learns representations for entity mentions and their context. We show experimentally that our approach is robust against noise and consistently outperforms the state-of-the-art on established benchmarks for the task.

1 Introduction

Fine-grained Entity Type Classification (FETC) aims at labeling entity mentions in context with one or more specific types organized in a hierarchy (e.g., **actor** as a subtype of **artist**, which in turn is a subtype of **person**). Fine-grained types help in many applications, including relation extraction (Mintz et al., 2009), question answering (Li and Roth, 2002), entity linking (Lin et al., 2012), knowledge base completion (Dong et al., 2014) and entity recommendation (Yu et al., 2014). Because of the high cost in labeling large training corpora with fine-grained types, current FETC systems resort to distant supervision (Mintz et al., 2009) and annotate mentions in the training corpus with *all types* associated with the entity in a knowledge graph. This is illustrated in

Figure 1, with three training sentences about entity *Steve Kerr*. Note that while the entity belongs to three fine-grained types (**person**, **athlete**, and **coach**), some sentences provide evidence of only some of the types: **person** and **coach** from **S1**, **person** and **athlete** from **S2**, and just **person** for **S3**. Clearly, direct distant supervision leads to noisy training data which can hurt the accuracy of the FETC model.

One kind of noise introduced by distant supervision is assigning labels that are *out-of-context* (**athlete** in **S1** and **coach** in **S2**) for the sentence. Current FETC systems sidestep the issue by either ignoring *out-of-context* labels or using simple pruning heuristics like discarding training examples with entities assigned to multiple types in the knowledge graph. However, both strategies are inelegant and hurt accuracy. Another source of noise introduced by distant supervision is when the type is *overly-specific* for the context. For instance, example **S3** does not support the inference that Mr. Kerr is either an **athlete** or a **coach**. Since existing knowledge graphs give more attention to notable entities with more specific types, *overly-specific* labels bias the model towards popular subtypes instead of generic ones, *i.e.*, preferring **athlete** over **person**. Instead of correcting for this bias, most existing FETC systems ignore the problem and treat each type equally and independently, ignoring that many types are semantically related.

Besides failing to handle noisy training data there are two other limitations of previous FETC approaches we seek to address. First, they rely on hand-crafted features derived from various NLP tools; therefore, the inevitable errors introduced by these tools propagate to the FETC systems via the training data. Second, previous systems treat FETC as a multi-label classification problem: during type inference they predict a plausibility score for each type, and, then, either classify types

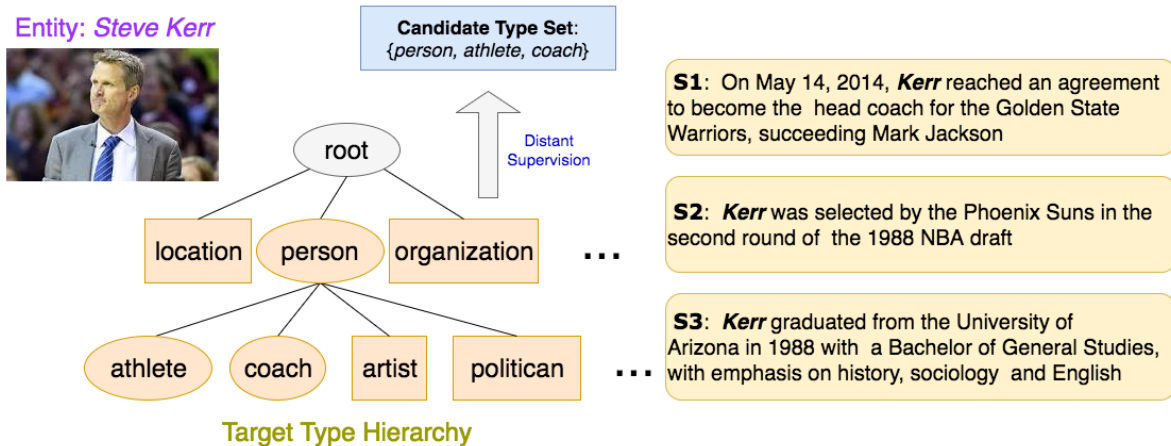


Figure 1: With distant supervision, all the three mentions of *Steve Kerr* shown are labeled with the same types in oval boxes in the target type hierarchy. While only part of the types are correct: **person** and **coach** for S1, **person** and **athlete** for S2, and just **person** for S3.

with scores above a threshold (Mintz et al., 2009; Gillick et al., 2014; Shimaoka et al., 2017) or perform a top-down search in the given type hierarchy (Ren et al., 2016a; Abhishek et al., 2017).

Contributions: We propose a neural network based model to overcome the drawbacks of existing FETC systems mentioned above. With publicly available word embeddings as input, we learn two different entity representations and use bidirectional long-short term memory (LSTM) with attention to learn the context representation. We propose a variant of cross entropy loss function to handle *out-of-context* labels automatically during the training phase. Also, we introduce hierarchical loss normalization to adjust the penalties for correlated types, allowing our model to understand the type hierarchy and alleviate the negative effect of *overly-specific* labels.

Moreover, in order to simplify the problem and take advantage of previous research on hierarchical classification, we transform the multi-label classification problem to a single-label classification problem. Based on the assumption that each mention can only have one *type-path* depending on the context, we leverage the fact that type hierarchies are forests, and represent each *type-path* uniquely by the terminal type (which might not be a leaf node). For Example, *type-path root-person-coach* can be represented as just **coach**, while *root-person* can be unambiguously represented as the non-leaf **person**.

Finally, we report on an experimental validation against the state-of-the-art on established bench-

marks that shows that our model can adapt to noise in training data and consistently outperform previous methods. In summary, we describe a single, much simpler and more elegant neural network model that attempts FETC “end-to-end” without post-processing or ad-hoc features and improves on the state-of-the-art for the task.

2 Related Work

Fine-Grained Entity Type Classification: The first work to use distant supervision (Mintz et al., 2009) to induce a large - but noisy - training set and manually label a significantly smaller dataset to evaluate their FETC system, was Ling and Weld (2012) who introduced both a training and evaluation dataset FIGER (GOLD). They used a linear classifier perceptron for multi-label classification. While initial work largely assumed that mention assignments could be done independently of the mention context, Gillick et al. (2014) introduced the concept of context-dependent FETC where the types of a mention are constrained to what can be deduced from its context and introduced a new OntoNotes-derived (Weischedel et al., 2011) manually annotated evaluation dataset. In addition, they addressed the problem of label noise induced by distant supervision and proposed three label cleaning heuristics. Yogatama et al. (2015) proposed an embedding-based model where user-defined features and labels were embedded into a low dimensional feature space to facilitate information sharing among labels. Ma et al. (2016) presented a label embedding method that incor-

	Attentive	AFET	LNR	AAA	NFETC
no hand-crafted features	—	—	—	✓	✓
uses attentive neural network	✓	—	—	—	✓
adopts single label setting	—	—	—	—	✓
handles <i>out-of-context</i> noise	—	✓	✓	✓	✓
handles <i>overly-specific</i> noise	—	✓	✓	—	✓

Table 1: Summary comparison to related FETC work. FETC systems listed in the table: (1) **Attentive** (Shimaoka et al., 2017); (2) **AFET** (Ren et al., 2016a); (3) **LNR** (Ren et al., 2016b); (4) **AAA** (Abhishek et al., 2017).

porates prototypical and hierarchical information to learn pre-trained label embeddings and adapted a zero-shot framework that can predict both seen and previously unseen entity types.

Shimaoka et al. (2016) proposed an attentive neural network model that used LSTMs to encode the context of an entity mention and used an attention mechanism to allow the model to focus on relevant expressions in such context. Shimaoka et al. (2017) summarizes many neural architectures for FETC task. These models ignore the *out-of-context noise*, that is, they assume that all labels obtained via distant supervision are “correct” and appropriate for every context in the training corpus. In our paper, a simple yet effective variant of cross entropy loss function is proposed to handle the problem of *out-of-context noise*.

Ren et al. (2016a) have proposed AFET, an FETC system, that separates the loss function for *clean* and *noisy* entity mentions and uses label-label correlation information obtained by given data in its parametric loss function. Considering the noise reduction aspects for FETC systems, Ren et al. (2016b) introduced a method called LNR to reduce label noise without data loss, leading to significant performance gains on both the evaluation dataset of FIGER(GOLD) and OntoNotes. Although these works consider both *out-of-context noise* and *overly-specific noise*, they rely on hand-crafted features which become an impediment to further improvement of the model performance. For LNR, because the noise reduction step is separated from the FETC model, the inevitable errors introduced by the noise reduction will be propagated into the FETC model which is undesirable. In our FETC system, we handle the problem induced from *irrelevant noise* and *overly-specific noise* seamlessly inside the model and avoid the usage of hand-crafted features.

Most recently, following the idea from AFET, Abhishek et al. (2017) proposed a simple neural network model which incorporates noisy label information using a variant of non-parametric

hinge loss function and gain great performance improvement on FIGER(GOLD). However, their work overlooks the effect of *overly-specific noise*, treating each type label equally and independently when learning the classifiers and ignores possible correlations among types.

Hierarchical Loss Function: Due to the intrinsic type hierarchy existing in the task of FETC, it is natural to adopt the idea of hierarchical loss function to adjust the penalties for FETC mistakes depending on how far they are in the hierarchy. The penalty for predicting **person** instead of **athlete** should be less than the penalty for predicting **organization**. To the best of our knowledge, the first use of a hierarchical loss function was originally introduced in the context of document categorization with support vector machines (Cai and Hofmann, 2004). However, that work assumed that weights to control the hierarchical loss would be solicited from domain experts, which is inapplicable for FETC. Instead, we propose a method called hierarchical loss normalization which can overcome the above limitations and be incorporated with cross entropy loss used in our neural architecture.

Table 1 provides a summary comparison of our work against the previous state-of-the-art in fine grained entity typing.

3 Background and Problem

Our task is to automatically reveal the type information for entity mentions in context. The input is a knowledge graph Ψ with schema \mathcal{Y}_Ψ , whose types are organized into a type hierarchy \mathcal{Y} , and an automatically labeled training corpus \mathcal{D} obtained by distant supervision with \mathcal{Y} . The output is a *type-path* in \mathcal{Y} for each named entity mentioned in a test sentence from a corpus \mathcal{D}_t .

More precisely, a labeled corpus for entity type classification consists of a set of extracted *entity mentions* $\{m_i\}_{i=1}^N$ (i.e., token spans representing entities in text), the *context* (e.g., sentence, paragraph) of each mention $\{c_i\}_{i=1}^N$, and the *candidate*

type sets $\{\mathcal{Y}_i\}_{i=1}^N$ automatically generated for each mention.

We represent the training corpus using a set of mention-based triples $\mathcal{D} = \{(m_i, c_i, \mathcal{Y}_i)\}_{i=1}^N$.

If \mathcal{Y}_i is free of *out-of-context noise*, the type labels for each m_i should form a single *type-path* in \mathcal{Y}_i . However, \mathcal{Y}_i may contain *type-paths* that are irrelevant to m_i in c_i if there exists *out-of-context noise*.

We denote the type set including all terminal types for each *type-path* as the target type set \mathcal{Y}_i^t . In the example type hierarchy shown in Figure 1, if \mathcal{Y}_i contains types **person**, **athlete**, **coach**, \mathcal{Y}_i^t should contain **athlete**, **coach**, but not **person**. In order to understand the trade-off between the effect of *out-of-context noise* and the size of the training set, we report on experiments with two different training sets: $\mathcal{D}_{filtered}$ only with triples whose \mathcal{Y}_i form a single *type-path* in \mathcal{D} , and \mathcal{D}_{raw} with all triples.

We formulate fine-grained entity classification problem as follows:

Definition 1 Given an entity mention $m_i = (w_p, \dots, w_t)$ ($p, t \in [1, T], p \leq t$) and its context $c_i = (w_1, \dots, w_T)$ where T is the context length, our task is to predict its most specific type \hat{y}_i depending on the context.

In practice, c_i is generated by truncating the original context with words beyond the context window size C both to the left and to the right of m_i . Specifically, we compute a probability distribution over all the $K = |\mathcal{Y}|$ types in the target type hierarchy \mathcal{Y} . The type with the highest probability is classified as the predicted type \hat{y}_i which is the terminal type of the predicted *type-path*.

4 Methodology

This section details our Neural Fine-Grained Entity Type Classification (NFETC) model.

4.1 Input Representation

As stated in Section 3, the input is an entity mention m_i with its context c_i . First, we transform each word in the context c_i into a real-valued vector to provide lexical-semantic features. Given a word embedding matrix W^{word} of size $d_w \times |V|$, where V is the input vocabulary and d_w is the size of word embedding, we map every w_i to a column vector $\mathbf{w}_i^d \in \mathbb{R}^{d_w}$.

To additionally capture information about the relationship to the target entities, we incorporate

word position embeddings (Zeng et al., 2014) to reflect relative distances between the i -th word to the entity mention. Every relative distance is mapped to a randomly initialized position vector in \mathbb{R}^{d_p} , where d_p is the size of position embedding. For a given word, we obtain the position vector \mathbf{w}_i^p . The overall embedding for the i -th word is $\mathbf{w}_i^E = [(\mathbf{w}_i^d)^\top, (\mathbf{w}_i^p)^\top]^\top$.

4.2 Context Representation

For the context c_i , we want to apply a non-linear transformation to the vector representation of c_i to derive a context feature vector $h_i = f(c_i; \theta)$ given a set of parameters θ . In this paper, we adopt bidirectional LSTM with d_s hidden units as $f(c_i; \theta)$. The network contains two sub-networks for the forward pass and the backward pass respectively. Here, we use element-wise sum to combine the forward and backward pass outputs. The output of the i -th word is shown in the following equation:

$$h_i = [\vec{h}_i \oplus \overleftarrow{h}_i] \quad (1)$$

Following Zhou et al. (2016), we employ word-level attention mechanism, which makes our model able to softly select the most informative words during training. Let H be a matrix consisting of output vectors $[h_1, h_2, \dots, h_T]$ that the LSTM produced. The context representation r is formed by a weighted sum of these output vectors:

$$G = \tanh(H) \quad (2)$$

$$\alpha = \text{softmax}(w^\top G) \quad (3)$$

$$r_c = H\alpha^\top \quad (4)$$

where $H \in \mathbb{R}^{d_s \times T}$, w is a trained parameter vector. The dimension of w, α, r_c are d_s, T, d_s respectively.

4.3 Mention Representation

Averaging encoder: Given the entity mention $m_i = (w_p, \dots, w_t)$ and its length $L = t - p + 1$, the averaging encoder computes the average word embedding of the words in m_i . Formally, the averaging representation r_a of the mention is computed as follows:

$$r_a = \frac{1}{L} \sum_{i=p}^t \mathbf{w}_i^d \quad (5)$$

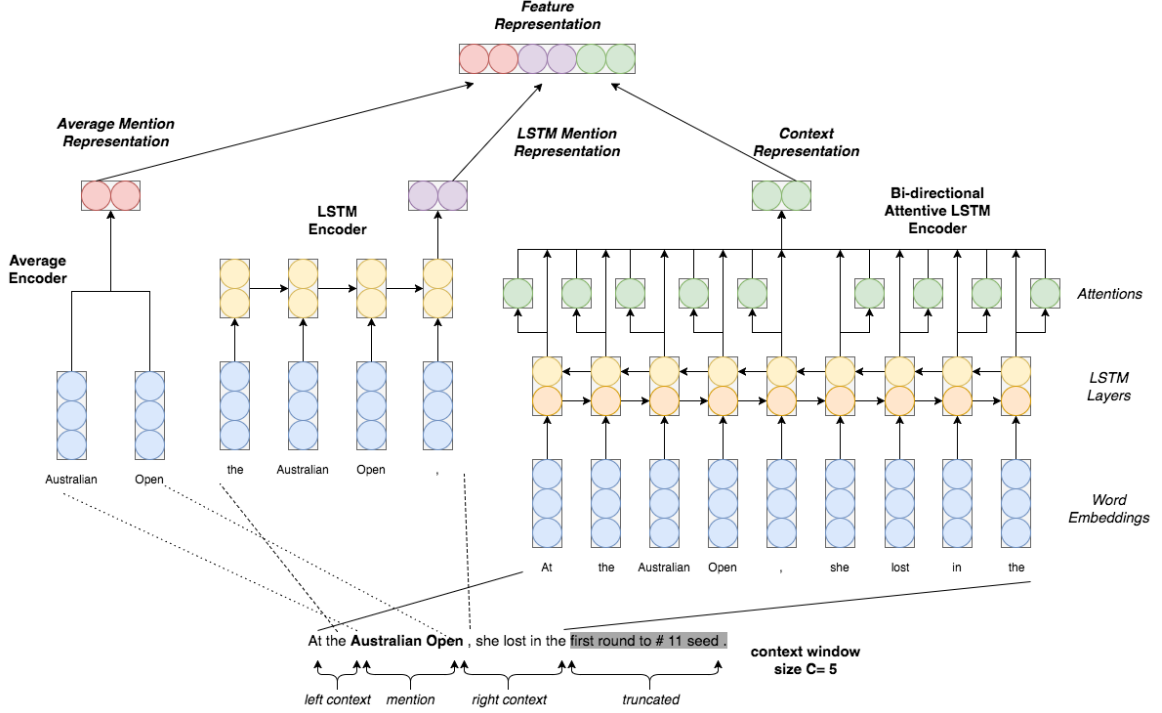


Figure 2: The architecture of the NFETC model.

This relatively simple method for composing the mention representation is motivated by it being less prone to overfitting (Shimaoka et al., 2017).

LSTM encoder: In order to capture more semantic information from the mentions, we add one token before and another after the target entity to the mention. The extended mention can be represented as $m_i^* = (w_{p-1}, w_p, \dots, w_t, w_{t+1})$. The standard LSTM is applied to the mention sequence from left to right and produces the outputs h_{p-1}, \dots, h_{t+1} . The last output h_{t+1} then serves as the LSTM representation r_l of the mention.

4.4 Optimization

We concatenate context representation and two mention representations together to form the overall feature representation of the input $R = [r_c, r_a, r_l]$. Then we use a softmax classifier to predict \hat{y}_i from a discrete set of classes for a entity mention m and its context c with R as input:

$$\hat{p}(y|m, c) = \text{softmax}(WR + b) \quad (6)$$

$$\hat{y} = \arg \max_y \hat{p}(y|m, c) \quad (7)$$

where W can be treated as the learned type embeddings and b is the bias.

The traditional cross-entropy loss function is represented as follows:

$$J(\theta) = -\frac{1}{N} \sum_{i=1}^N \log(\hat{p}(y_i|m_i, c_i)) + \lambda \|\Theta\|^2 \quad (8)$$

where y_i is the only element in \mathcal{Y}_i^t and $(m_i, c_i, \mathcal{Y}_i) \in \mathcal{D}_{filtered}$. λ is an L2 regularization hyperparameter and Θ denotes all parameters of the considered model.

In order to handle data with *out-of-context noise* (in other words, with multiple labeled types) and take full advantage of them, we introduce a simple yet effective variant of the cross-entropy loss:

$$J(\theta) = -\frac{1}{N} \sum_{i=1}^N \log(\hat{p}(y_i^*|m_i, c_i)) + \lambda \|\Theta\|^2 \quad (9)$$

where $y_i^* = \arg \max_{y \in \mathcal{Y}_i^t} \hat{p}(y|m_i, c_i)$ and $(m_i, c_i, \mathcal{Y}_i) \in \mathcal{D}_{raw}$. With this loss function, we assume that the type with the highest probability among \mathcal{Y}_i^t during training as the correct type. If there is only one element in \mathcal{Y}_i^t , this loss function is equivalent to the cross-entropy loss function. Wherever there are multiple elements, it can filter the less probable types based on the local context automatically.

4.5 Hierarchical Loss Normalization

Since the fine-grained types tend to form a forest of type hierarchies, it is unreasonable to treat every type equally. Intuitively, it is better to predict an ancestor type of the true type than some other unrelated type. For instance, if one example is labeled as **athlete**, it is reasonable to predict its type as **person**. However, predicting other high level types like **location** or **organization** would be inappropriate. In other words, we want the loss function to penalize less the cases where types are related. Based on the above idea, we adjust the estimated probability as follows:

$$p^*(\hat{y}|m, c) = p(\hat{y}|m, c) + \beta * \sum_{t \in \Gamma} p(t|m, c) \quad (10)$$

where Γ is the set of ancestor types along the *type-path* of \hat{y} , β is a hyperparameter to tune the penalty. Afterwards, we re-normalize it back to a probability distribution, a process which we denote as *hierarchical loss normalization*.

As discussed in Section 1, there exists *overly-specific noise* in the automatically labeled training sets which hurt the model performance severely. With *hierarchical loss normalization*, the model will get less penalty when it predicts the actual type for one example with *overly-specific noise*. Hence, it can alleviate the negative effect of *overly-specific noise* effectively. Generally, *hierarchical loss normalization* can make the model somewhat understand the given type hierarchy and learn to detect those overly-specific cases. During classification, it will make the models prefer generic types unless there is a strong indicator for a more specific type in the context.

4.6 Regularization

Dropout, proposed by Hinton et al. (2012), prevents co-adaptation of hidden units by randomly omitting feature detectors from the network during forward propagation. We employ both input and output dropout on LSTM layers. In addition, we constrain L2-norms for the weight vectors as shown in Equations 8, 9 and use early stopping to decide when to stop training.

5 Experiments

This section reports an experimental evaluation of our NFETC approach using the previous state-of-the-art as baselines.

	FIGER(GOLD)	OntoNotes
# types	113	89
# raw training mentions	2009898	253241
# raw testing mentions	563	8963
% filtered training mentions	64.46	73.13
% filtered testing mentions	88.28	94.00
Max hierarchy depth	2	3

Table 2: Statistics of the datasets

5.1 Datasets

We evaluate the proposed model on two standard and publicly available datasets, provided in a pre-processed tokenized format by Shimaoka et al. (2017). Table 2 shows statistics about the benchmarks. The details are as follows:

- **FIGER(GOLD):** The training data consists of Wikipedia sentences and was automatically generated with distant supervision, by mapping Wikipedia identifiers to Freebase ones. The test data, mainly consisting of sentences from news reports, was manually annotated as described by Ling and Weld (2012).
- **OntoNotes:** The OntoNotes dataset consists of sentences from newswire documents present in the OntoNotes text corpus (Weischedel et al., 2013). DBpedia spotlight (Daiber et al., 2013) was used to automatically link entity mention in sentences to Freebase. Manually annotated test data was shared by Gillick et al. (2014).

Because the type hierarchy can be somewhat understood by our proposed model, the quality of the type hierarchy can also be a key factor to the performance of our model. We find that the type hierarchy for FIGER(GOLD) dataset following Freebase has some flaws. For example, **software** is not a subtype of **product** and **government** is not a subtype of **organization**. Following the proposed type hierarchy of Ling and Weld (2012), we refine the Freebase-based type hierarchy. The process is a one-to-one mapping for types in the original dataset and we didn't add or drop any type or sentence in the original dataset. As a result, we can directly compare the results of our proposed model with or without this refinement.

Aside from the advantages brought by adopting the single label classification setting, we can see one disadvantage of this setting based on Table 2. That is, the performance upper bounds of

our proposed model are no longer 100%: for example, the best strict accuracy we can get in this setting is 88.28% for **FIGER(GOLD)**. However, as the strict accuracy of state-of-the-art methods are still nowhere near 80% (Table 3), the evaluation we perform is still informative.

5.2 Baselines

We compared the proposed model with state-of-the-art FETC systems ¹: (1) **Attentive** (Shimaoka et al., 2017); (2) **AFET** (Ren et al., 2016a); (3) **LNR+FIGER** (Ren et al., 2016b); (4) **AAA** (Abhishek et al., 2017).

We compare these baselines with variants of our proposed model: (1) **NFETC(f)**: basic neural model trained on $\mathcal{D}_{filtered}$ (recall Section 4.4); (2) **NFETC-hier(f)**: neural model with hierarchical loss normalization trained on $\mathcal{D}_{filtered}$. (3) **NFETC(r)**: neural model with proposed variant of cross-entropy loss trained on \mathcal{D}_{raw} ; (4) **NFETC-hier(r)**: neural model with proposed variant of cross-entropy loss and hierarchical loss normalization trained on \mathcal{D}_{raw} .

5.3 Experimental Setup

For evaluation metrics, we adopt the same criteria as Ling and Weld (2012), that is, we evaluate the model performance by strict accuracy, loose macro, and loose micro F-scores. These measures are widely used in existing FETC systems (Shimaoka et al., 2017; Ren et al., 2016b,a; Abhishek et al., 2017).

We use pre-trained word embeddings that were not updated during training to help the model generalize to words not appearing in the training set. For this purpose, we used the freely available 300-dimensional cased word embedding trained on 840 billion tokens from the Common Crawl supplied by Pennington et al. (2014). For both datasets, we randomly sampled 10% of the test set as a development set, on which we do the hyperparameters tuning. The remaining 90% is used for final evaluation. We run each model with the well-tuned hyperparameter setting five times and report their average strict accuracy, macro F1 and micro F1 on the test set. The proposed model was implemented using the TensorFlow framework. ²

¹The results of the baselines are all as reported in their corresponding papers.

²The code to replicate the work is available at: <https://github.com/billy-inn/NFETC>

Parameter	FIGER(GOLD)	OntoNotes
lr	0.0002	0.0002
d_p	85	20
d_s	180	440
p_i	0.7	0.5
p_o	0.9	0.5
λ	0.0	0.0001
β	0.4	0.3

Table 4: Hyperparameter Settings

5.4 Hyperparameter Setting

In this paper, we search different hyperparameter settings for FIGER(GOLD) and OntoNotes separately, considering the differences between the two datasets. The hyperparameters include the learning rate lr for Adam Optimizer, size of word position embeddings (WPE) d_p , state size for LSTM layers d_s , input dropout keep probability p_i and output dropout keep probability p_o for LSTM layers ³, L2 regularization parameter λ and parameter to tune hierarchical loss normalization β . The values of these hyperparameters, obtained by evaluating the model performance on the development set, for each dataset can be found in Table 4.

5.5 Performance comparison and analysis

Table 3 compares our models with other state-of-the-art FETC systems on FIGER(GOLD) and OntoNotes. The proposed model performs better than the existing FETC systems, consistently on both datasets. This indicates benefits of the proposed representation scheme, loss function and hierarchical loss normalization.

Discussion about *Out-of-context Noise*: For dataset FIGER(GOLD), the performance of our model with the proposed variant of cross-entropy loss trained on \mathcal{D}_{raw} is significantly better than the basic neural model trained on $\mathcal{D}_{filtered}$, suggesting that the proposed variant of the cross-entropy loss function can make use of the data with *out-of-context noise* effectively. On the other hand, the improvement introduced by our proposed variant of cross-entropy loss is not as significant for the OntoNotes benchmark. This may be caused by the fact that OntoNotes is much smaller than FIGER(GOLD) and proportion of examples without *out-of-context noise* are also higher, as shown in Table 2.

³Following TensorFlow terminology.

Model	FIGER(GOLD)			OntoNotes		
	Strict Acc.	Macro F1	Micro F1	Strict Acc.	Macro F1	Micro F1
Attentive	59.68	78.97	75.36	51.74	70.98	64.91
AFET	53.3	69.3	66.4	55.1	71.1	64.7
LNR+FIGER	59.9	76.3	74.9	57.2	71.5	66.1
AAA	65.8	81.2	77.4	52.2	68.5	63.3
NFETC(f)	57.9 ± 1.3	78.4 ± 0.8	75.0 ± 0.7	54.4 ± 0.3	71.5 ± 0.4	64.9 ± 0.3
NFETC-hier(f)	68.0 ± 0.8	81.4 ± 0.8	77.9 ± 0.7	59.6 ± 0.2	76.1 ± 0.2	69.7 ± 0.2
NFETC(r)	56.2 ± 1.0	77.2 ± 0.9	74.3 ± 1.1	54.8 ± 0.4	71.8 ± 0.4	65.0 ± 0.4
NFETC-hier(r)	68.9 ± 0.6	81.9 ± 0.7	79.0 ± 0.7	60.2 ± 0.2	76.4 ± 0.1	70.2 ± 0.2

Table 3: Strict Accuracy, Macro F1 and Micro F1 for the models tested on the FIGER(GOLD) and OntoNotes datasets.

Test Sentence	Ground Truth
S1: <i>Hopkins</i> said four fellow elections is curious , considering the ...	Person
S2: ... for WiFi communications across all <i>the SD cards</i> .	Product
S3: A handful of professors in the <i>UW</i> Department of Chemistry ...	Educational Institution
S4: Work needs to be done and, in <i>Washington state</i> , ...	Province
S5: <i>ASC</i> Director Melvin Taing said that because the commission is ...	Organization

Table 5: Examples of test sentences in FIGER(GOLD) where the entity mentions are marked as bold italics.

Investigations on *Overly-Specific Noise*: With hierarchical loss normalization, the performance of our models are consistently better no matter whether trained on \mathcal{D}_{raw} or $\mathcal{D}_{filtered}$ on both datasets, demonstrating the effectiveness of this hierarchical loss normalization and showing that *overly-specific noise* has a potentially significant influence on the performance of FETC systems.

5.6 T-SNE Visualization of Type Embeddings

By visualizing the learned type embeddings (Figure 3), we can observe that the parent types are mixed with their subtypes and forms clear distinct clusters without hierarchical loss normalization, making it hard for the model to distinguish subtypes like **actor** or **athlete** from their parent types **person**. This also biases the model towards the most popular subtype. While the parent types tend to cluster together and the general pattern is more complicated with hierarchical loss normalization. Although it’s not as easy to interpret, it hints that our model can learn rather subtle intricacies and correlations among types latent in the data with the help of hierarchical loss normalization, instead of sticking to a pre-defined hierarchy.

5.7 Error Analysis on FIGER(GOLD)

Since there are only 563 sentences for testing in FIGER(GOLD), we look into the predictions for

all the test examples of all variants of our model. Table 5 shows 5 examples of test sentence. Without hierarchical loss normalization, our model will make too aggressive predictions for S1 with **Politician** and for S2 with **Software**. This kind of mistakes are very common and can be effectively reduced by introducing hierarchical loss normalization leading to significant improvements on the model performance. Using the changed loss function to handle multi-label (noisy) training data can help the model distinguish ambiguous cases. For example, our model trained on $\mathcal{D}_{filtered}$ will misclassify S5 as **Title**, while the model trained on \mathcal{D}_{raw} can make the correct prediction.

However, there are still some errors that can’t be fixed with our model. For example, our model cannot make correct predictions for S3 and S4 due to the fact that our model doesn’t know that *UW* is an abbreviation of *University of Washington* and *Washington state* is the name of a province. In addition, the influence of *overly-specific noise* can only be alleviated but not eliminated. Sometimes, our model will still make too aggressive or conservative predictions. Also, mixing up very ambiguous entity names is inevitable in this task.




Figure 3: T-SNE visualization of the type embeddings learned from FIGER(GOLD) dataset where subtypes share the same color as their parent type. The seven parent types are shown in the black boxes. The below sub-figure uses the hierarchical loss normalization, while the above not.

6 Conclusion and Further Work

In this paper, we studied two kinds of noise, namely *out-of-context noise* and *overly-specific noise*, for noisy type labels and investigate their effects on FETC systems. We proposed a neural network based model which jointly learns representations for entity mentions and their context. A variant of cross-entropy loss function was used to handle *out-of-context noise*. Hierarchical loss normalization was introduced into our model to alleviate the effect of *overly-specific noise*. Experimental results on two publicly available datasets demonstrate that the proposed model is robust to these two kind of noise and outperforms previous state-of-the-art methods significantly.

More work can be done to further develop hierarchical loss normalization since currently it's very simple. Considering type information is valuable in various NLP tasks, we can incorporate results produced by our FETC system to other tasks, such as relation extraction, to check our model's effectiveness and help improve other tasks' per-

formance. In addition, tasks like relation extraction are complementary to the task of FETC and therefore may have potentials to be digged to help improve the performance of our system in return.

Acknowledgments

This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

- Abhishek Abhishek, Ashish Anand, and Amit Awekar. 2017. Fine-grained entity type classification by jointly learning representations and label embeddings. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL)* pages 797–807.
- Lijuan Cai and Thomas Hofmann. 2004. Hierarchical document categorization with support vector machines. In *Proceedings of the thirteenth ACM international conference on Information and knowledge management*. ACM, pages 78–87.
- Joachim Daiber, Max Jakob, Chris Hokamp, and Pablo N Mendes. 2013. Improving efficiency and accuracy in multilingual entity extraction. *Proceedings of the 9th International Conference on Semantic Systems* pages 121–124.
- Xin Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmann, Shaohua Sun, and Wei Zhang. 2014. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining* pages 601–610.
- Dan Gillick, Nevena Lazic, Kuzman Ganchev, Jesse Kirchner, and David Huynh. 2014. Context-dependent fine-grained entity type tagging. *arXiv preprint arXiv:1412.1820*.
- Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. 2012. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*.
- Xin Li and Dan Roth. 2002. Learning question classifiers. In *Proceedings of the 19th international conference on Computational linguistics-Volume 1*. Association for Computational Linguistics, pages 1–7.
- Thomas Lin, Oren Etzioni, et al. 2012. No noun phrase left behind: detecting and typing unlinkable entities. *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning* pages 893–903.

- Xiao Ling and Daniel S Weld. 2012. Fine-grained entity recognition. *AAAI* .
- Yukun Ma, Erik Cambria, and Sa Gao. 2016. Label embedding for zero-shot fine-grained named entity typing. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*. pages 171–180.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP* pages 1003–1011.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. *EMNLP* 14(1532–1543).
- Xiang Ren, Wenqi He, Meng Qu, Lifu Huang, Heng Ji, and Jiawei Han. 2016a. Afet: Automatic fine-grained entity typing by hierarchical partial-label embedding. *EMNLP* 16(17).
- Xiang Ren, Wenqi He, Meng Qu, Clare R Voss, Heng Ji, and Jiawei Han. 2016b. Label noise reduction in entity typing by heterogeneous partial-label embedding. *KDD* .
- Sonse Shimaoka, Pontus Stenetorp, Kentaro Inui, and Sebastian Riedel. 2016. An attentive neural architecture for fine-grained entity type classification. *arXiv preprint arXiv:1604.05525* .
- Sonse Shimaoka, Pontus Stenetorp, Kentaro Inui, and Sebastian Riedel. 2017. Neural architectures for fine-grained entity type classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL)* .
- Ralph Weischedel, Eduard Hovy, Mitchell Marcus, Martha Palmer, Robert Belvin, Sameer Pradhan, Lance Ramshaw, and Nianwen Xue. 2011. Ontonotes: A large training corpus for enhanced processing. *Handbook of Natural Language Processing and Machine Translation*. Springer .
- Ralph Weischedel, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Nianwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, et al. 2013. Ontonotes release 5.0 ldc2013t19. *Linguistic Data Consortium, Philadelphia, PA* .
- Dani Yogatama, Daniel Gillick, and Nevena Lazic. 2015. Embedding methods for fine grained entity type classification. *ACL (2)* pages 291–296.
- Xiao Yu, Xiang Ren, Yizhou Sun, Quanquan Gu, Bradley Sturt, Urvashi Khandelwal, Brandon Norick, and Jiawei Han. 2014. Personalized entity recommendation: A heterogeneous information network approach. *Proceedings of the 7th ACM international conference on Web search and data mining* pages 283–292.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, Jun Zhao, et al. 2014. Relation classification via convolutional deep neural network. *COLING* pages 2335–2344.
- Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. Attention-based bidirectional long short-term memory networks for relation classification. *The 54th Annual Meeting of the Association for Computational Linguistics* .