

Joint Embedding of Hierarchical Categories and Entities for Concept Categorization and Dataless Classification

Yuezhang Li, Ronghuo Zheng, Tian Tian, Zhiting Hu, Rahul Iyer, Katia Sycara

Language Technology Institute, School of Computer Science

Carnegie Mellon University

yuezhanl@andrew.cmu.edu

Abstract

Existing work learning distributed representations of knowledge base entities has largely failed to incorporate rich categorical structure, and is unable to induce category representations. We propose a new framework that embeds entities and categories jointly into a semantic space, by integrating structured knowledge and taxonomy hierarchy from large knowledge bases. Our framework enables to compute meaningful semantic relatedness between entities and categories in a principled way, and can handle both single-word and multiple-word concepts. Our method shows significant improvement on the tasks of concept categorization and dataless hierarchical classification.

1 Introduction

Hierarchies, most commonly represented as tree or directed acyclic graph (DAG) structures, provide a natural way to categorize and locate knowledge in large knowledge bases (KBs). For example, WordNet (Fellbaum, 1998), Freebase (Bollacker et al., 2008), and Wikipedia¹ use hierarchical taxonomy to organize entities into category hierarchies. These hierarchical categories could benefit applications such as object and concept categorization (Verma et al., 2012; Rothenhäusler and Schütze, 2009), document classification (Gopal and Yang, 2013; Hu et al., 2016), and link prediction on knowledge graphs (Lin et al., 2015; Zhang et al., 2016). In all of these applications, it is essential to have a good representation of categories and entities as well as a good semantic relatedness measure.

In this paper, we propose two models to learn distributed representations of both entities and categories from large-scale knowledge bases (KBs). The **Category Embedding (CE) model** extends the entity embedding method of (Hu et al., 2015) and induces category embeddings in addition to entity vectors. The **Hierarchical Category Embedding (HCE) model** further enhances the CE model by integrating category hierarchical structure. The resulting entity and category vectors effectively capture meaningful semantic relatedness between entities and categories.

We train the category and entity vectors on Wikipedia, and evaluate our methods from two applications: **concept categorization** (Baroni and Lenci, 2010) and **dataless hierarchical classification** (Song and Roth, 2014). Our method achieves state-of-the-art performance on both tasks.

Figure 1 shows an overview of the research elements in this paper. In summary, we make the following contributions:

- We incorporate category information into entity embeddings with the proposed CE model to get entity and category embeddings simultaneously.
- We leverage category hierarchies and develop the HCE model to enhance the embedding quality.
- We propose a new concept categorization method based on nearest neighbor classification. Our method avoids the granularity disparity issue of categories which has plagued traditional clustering-based methods.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

This work has been partially funded by ARO award #W911NF1210416

¹<https://www.wikipedia.org>




Figure 1: The organization of research elements comprising this paper.

- We construct a new concept categorization dataset from Wikipedia.
- We show strong potential of utilizing entity embeddings on dataless classification.

2 Related Work

Entity embedding method is based on the analysis of distributional semantics, e.g., the recently proposed skip-gram model (Mikolov et al., 2013) that learns to predict context words given target words. This model tries to maximize the average log likelihood of the context words so that the embeddings encode meaningful semantics. For instance, entity hierarchy embedding (Hu et al., 2015) extends it to predict context entities given target entity in KBs. (Yamada et al., 2016) proposed a method to jointly embed words and entities by optimizing word-word, entity-word, and entity-entity predicting models. Our models extend this line of research by incorporating hierarchical category information to embed categories and entities in a joint semantic space.

Relational embedding, also known as knowledge graph embedding, is a family of methods to represent entities as vectors and relations as operations applied to entities such that certain properties are preserved (Paccanaro and Hinton, 2001; Bordes et al., 2011; Nickel et al., 2012; Bordes et al., 2013; Neelakantan and Chang, 2015). For instance, the linear relational embedding (Paccanaro and Hinton, 2001) applies a relation to an entity based on matrix-vector multiplication, while TransE (Bordes et al., 2013) simplifies the operation to vector addition. To derive the embedding representation, they minimize a global loss function considering all (entity, relation, entity) triplets so that the embeddings encode meaningful semantics. Our approach differs from this line as we learn from raw text in addition to knowledge graphs, and methodologically use probabilistic models instead of transition-based models.

Category embedding has been widely explored within different context. (Weinberger and Chapelle, 2009) initially proposed a taxonomy embedding to achieve document categorization and derive the embeddings of taxonomies of topics in form of topic prototypes. Furthermore, (Hwang and Sigal, 2014) proposed a discriminative learning framework that can give category embeddings by approximating each category embeddings as a sum of its direct super-category plus a sparse combination of attributes. These methods have primarily targeted on documents/object classification rather than entity representations.

Several recent methods are proposed to extend word representation to phrases (Yin and Schütze, 2014; Yu and Dredze, 2015; Passos et al., 2014). However, they do not use structured knowledge to derive phrase representations.

3 Joint Embedding of Categories and Entities

In order to find representations for categories and entities that can capture their semantic relatedness, we use existing hierarchical categories and entities labeled with these categories, and explore two methods: 1) **Category Embedding model** (CE Model): it replaces the entities in the context with their directly labeled categories to build categories' context; 2) **Hierarchical Category Embedding** (HCE Model): it further incorporates all ancestor categories of the context entities to utilize the hierarchical information.

3.1 Category Embedding (CE) Model

Our category embedding (CE) model is based on the Skip-gram word embedding model (Mikolov et al., 2013). The skip-gram model aims at generating word representations that are good at predicting *context* words surrounding a *target* word in a sliding window. Previous work (Hu et al., 2015) extends the entity’s context to the whole article that describes the entity and acquires a set of entity pairs $\mathbf{D} = \{(e_t, e_c)\}$, where e_t denotes the *target* entity and e_c denotes the *context* entity.

Our CE model extends those approaches by incorporating category information. In KBs such as Wikipedia, category hierarchies are usually given as DAG or tree structures, and entities are categorized into one or more categories as leaves. Thus, in KBs, each entity e_t is labeled with one or more categories (c_1, c_2, \dots, c_k) , $k \geq 1$ and described by an article containing other *context* entities (see Data in Figure 1).

To learn embeddings of entities and categories simultaneously, we adopt a method that incorporates the labeled categories into the entities when predicting the context entities, similar to TWE-1 model (Liu et al., 2015) which incorporates topic information with words to predict context words. For example, if e_t is the *target* entity in the document, its labeled categories (c_1, c_2, \dots, c_k) would be combined with the entity e_t to predict the context entities like e_{c1} and e_{c2} (see CE Model in Figure 1). For each target-context entity pair (e_t, e_c) , the probability of e_c being context of e_t is defined as the following softmax:

$$P(e_c|e_t) = \frac{\exp(e_t \cdot e_c)}{\sum_{e \in \mathbf{E}} \exp(e_t \cdot e)}, \quad (1)$$

where \mathbf{E} denotes the set of all entity vectors, and \exp is the exponential function. For convenience, here we abuse the notation of e_t and e_c to denote a target entity vector and a context entity vector respectively.

Similar to TWE-1 model, We learn the target and context vectors by maximizing the average log probability:

$$L = \frac{1}{|\mathbf{D}|} \sum_{(e_c, e_t) \in \mathbf{D}} \left[\log P(e_c|e_t) + \sum_{c_i \in \mathbf{C}(e_t)} \log P(e_c|c_i) \right], \quad (2)$$

where \mathbf{D} is the set of all entity pairs and we abuse the notation of c_i to denote a category vector, and $\mathbf{C}(e_t)$ denotes the categories of entity e_t .

3.2 Hierarchical Category Embedding(HCE) Model

From the design above, we can get the embeddings of all categories and entities in KBs without capturing the semantics of hierarchical structure of categories. In a category hierarchy, the categories at lower layers will cover fewer but more specific concepts than categories at upper layers. To capture this feature, we extend the CE model to further incorporate the ancestor categories of the target entity when predicting the context entities (see HCE Model in Figure 1). If a category is near an entity, its ancestor categories would also be close to that entity. On the other hand, an increasing distance of the category from the target entity would decrease the power of that category in predicting context entities. Therefore, given the target entity e_t and the context entity e_c , the objective is to maximize the following weighted average log probability:

$$L = \frac{1}{|\mathbf{D}|} \sum_{(e_c, e_t) \in \mathbf{D}} \left[\log P(e_c|e_t) + \sum_{c_i \in \mathbf{A}(e_t)} w_i \log P(e_c|c_i) \right], \quad (3)$$

where $\mathbf{A}(e_t)$ represents the set of ancestor categories of entity e_t , and w_i is the weight of each category in predicting the context entity. To implement the intuition that a category is more relevant to its closer ancestor, for example, “NBC Mystery Movies” is more relevant to “Mystery Movies” than “Entertainment”, we set $w_i \propto \frac{1}{l(c_c, c_i)}$ where $l(c_c, c_i)$ denotes the average number of steps going down from category c_i to category c_c , and it is constrained with $\sum_i w_i = 1$.

Figure 2 presents the results of our HCE model for DOTA-all data set (see Section 5.1.1). The visualization shows that our embedding method is able to clearly separate entities into distinct categories.




Figure 2: Category and entity embedding visualization of the DOTA-all data set (see Section 5.1.1). We use t-SNE (Van der Maaten and Hinton, 2008) algorithms to map vectors into a 2-dimensional space. Labels with the same color are entities belonging to the same category. Labels surrounded by a box are categories vectors.

3.3 Learning

Learning CE and HCE models follows the optimization scheme of skip-gram model (Mikolov et al., 2013). We use negative sampling to reformulate the objective function, which is then optimized through stochastic gradient descent (SGD).

Specifically, the likelihood of each context entity of a target entity is defined with the softmax function in Eq. 1, which iterates over all entities. Thus, it is computationally intractable. We apply the standard negative sampling technique to transform the objective function in equation (3) to equation (4) below and then optimize it through SGD:

$$L = \sum_{(e_c, e_t) \in \mathbf{D}} \left[\log \sigma(e_c \cdot e_t) + \sum_{c_i \in \mathbf{A}(e_t)} w_i \log \sigma(e_c \cdot c_i) \right] + \sum_{(e'_c, e_t) \in \mathbf{D}'} \left[\log \sigma(-e'_c \cdot e_t) + \sum_{c_i \in \mathbf{A}(e_t)} w_i \log \sigma(-e'_c \cdot c_i) \right], \quad (4)$$

where \mathbf{D}' is the set of negative sample pairs and $\sigma(x) = 1/(1 + \exp(-x))$ is the sigmoid function.

4 Applications

We apply our category and entity embedding to two applications: concept categorization and dataless hierarchical classification.

4.1 Concept Categorization

Concept² categorization, also known as *concept learning* or *noun categorization*, is a process of assigning a concept to one candidate category, given a set of concepts and candidate categories. Traditionally, concept categorization is achieved by **concept clustering** due to the lack of category representations. Since our model can generate representations of categories, we propose a new method of using **nearest neighbor (NN) classification** to directly categorize each concept to a certain category.

The **concept clustering** is defined as: given a set of concepts like *dog*, *cat*, *apple* and the corresponding gold standard categorizations like *animal*, *fruit*, apply a word space model to project all the concepts to

²In this paper, concept and entity denote the same thing.

a semantic space and perform clustering. The clustering results can be evaluated by comparing with the gold standard categorizations. Since we have representation of categories, we propose **nearest neighbor (NN) classification** to categorize concepts by directly comparing concept vectors with candidate category vectors. Precisely, given a set of concepts \mathbf{E} and a set of candidate categories \mathbf{C} , we convert all concepts to concept vectors and all candidate categories to category vectors. Then we use the equation $c = \operatorname{argmin}_{c_i \in \mathbf{C}} \|c_i - e\|$ to assign the concept vector e with category c .

Since purity works as a standard evaluation metric for clustering (Rothenhäusler and Schütze, 2009), to compare our model with the concept clustering, we also use purity to measure our model’s performance. Specifically, **purity** is defined as:

$$\text{purity}(\mathbf{\Omega}, \mathbf{G}) = \frac{1}{n} \sum_k \max_j |\omega_k \cap g_j|, \quad (5)$$

where $\mathbf{\Omega}$ denotes a clustering solution, \mathbf{G} is a set of gold standard classes, n is the number of instances, k is the cluster index and j is the class index, ω_k represents the set of labels in a cluster and g_j is the set of labels in a class. A higher purity indicates better model performance.

4.2 Dataless Classification

Dataless classification uses the similarity between documents and labels in an enriched semantic space to determine in which category the given document is. It has been proved that explicit semantic analysis (ESA) (Gabrilovich and Markovitch, 2007) has shown superior performance on dataless classification (Song and Roth, 2014). ESA uses a bag-of-entities retrieved from Wikipedia to represent the text. For example, the document “*Jordan plays basketball*” can be represented as a ESA sparse vector of $\{Michael\ Jordan:48.3, Basketball:29.8, Air\ Jordan: 28.8, Outline\ of\ basketball: 28.5, Chicago\ Bulls: 23.6\}$ in an ESA implementation. After converting documents and short label descriptions to ESA vectors, a nearest neighbor classifier is applied to assign labels for each document.

Due to sparsity problem of ESA vectors, sparse vector densification is introduced to augment the similarity calculation between two ESA vectors (Song and Roth, 2015). This is achieved by considering pairwise similarity between entities in ESA vectors. However, they simply use word2vec (Mikolov et al., 2013) to derive entity representation. We extend it by directly applying entity embeddings and show the potential of entity embeddings to improve dataless classification.

We use averaged F_1 scores to measure the performance of all methods (Yang, 1999). Let TP_i, FP_i, FN_i denote the true-positive, false-positive and false-negative values for the i th label in label set T , the micro- and macro-averaged F_1 scores are defined as: $MicroF_1 = 2\bar{P} * \bar{R} / (\bar{P} + \bar{R})$ and $MacroF_1 = \frac{1}{|T|} \sum_i^{|T|} \frac{2P_i * R_i}{P_i + R_i}$, where $P_i = TP_i / (TP_i + FP_i)$ and $R_i = TP_i / (TP_i + FN_i)$ are precision and recall for i th label, $\bar{P} = \sum_i^{|T|} TP_i / \sum_i^{|T|} (TP_i + FP_i)$ and $\bar{R} = \sum_i^{|T|} TP_i / \sum_i^{|T|} (TP_i + FN_i)$ are averaged precision and recall for all labels.

5 Experiments

In the experiments, we use the dataset collected from Wikipedia on Dec. 1, 2015³ as the training data. We preprocess the category hierarchy by pruning administrative categories and deleting bottom-up edges to construct a DAG. The final version of data contains 5,373,165 entities and 793,856 categories organized as a DAG with a maximum depth of 18. The root category is “main topic classifications”. We train category and entity vectors in dimensions of 50, 100, 200, 250, 300, 400, 500, with batch size $B = 500$ and negative sample size $k = 10$.

With the training dataset defined above, we conduct experiments on two applications: concept categorization and dataless hierarchical classification.

³<https://dumps.wikimedia.org/wikidatawiki/20151201/>

5.1 Concept Categorization

5.1.1 Datasets

There are two datasets used in this experiment. The first one is the **Battig** test set introduced by (Baroni and Lenci, 2010), which includes 83 concepts from 10 categories. The **Battig** test set only contains single-word concepts without any multiple-word concepts (e.g., “table tennis”). Hence, using this dataset restricts the power of concept categorization to single-word level. We use this dataset because it has been used as a benchmark for most previous approaches for concept categorization.

Due to the limitations of the **Battig** test set, we construct a new entity categorization dataset **DOTA** (Dataset Of enTity cAtegorization) with 450 entities categorized into 15 categories (refer to Appendix A). All the categories and entities are extracted from Wikipedia, so the resulting dataset does not necessarily contains only single-word entities. Thus, the dataset can be split into two parts, **DOTA-single** that contains 300 single-word entities categorized into 15 categories and **DOTA-mult** that contains 150 multiple-word entities categorized into the same 15 categories. We design the **DOTA** dataset based on the following principles:

Coverage vs Granularity: Firstly, the dataset should cover at least one category of Wikipedia’s main topics including “Culture”, “Geography”, “Health”, “Mathematics”, “Nature”, “People”, “Philosophy”, “Religion”, “Society” and “Technology”. Secondly, categories should be in different granularity, from large categories (e.g., “philosophy”) to small categories (e.g., “dogs”). Large categories are ones that are located within 5 layers away from the root, medium categories are 6-10 layers away from the root, while small categories have distance of 11-18 to the root. Our dataset consists of 1/3 large categories, 1/3 medium categories, and 1/3 small categories.

Single-Words vs Multiple-Words: Previous concept categorization datasets only contain single-words. However, some concepts are multiple-words and cannot be simply represented by single-words. For example, the concept “hot dog” is very different from the concept “dog”. Therefore, we make each category of the dataset contain 10 multiple-word entities and 20 single-word entities.

5.1.2 Baselines

We compare our entity and category embeddings with the following state of the art word and entity embeddings.

WE_{Mikolov} (Mikolov et al., 2013): We trained word embeddings with Mikolov’s word2vec toolkit⁴ on the same Wikipedia corpus as ours (1.7 billion tokens) and then applied the Skip-gram model with negative sample size of 10 and window size of 5 in dimensionality of 50, 100, 200, 250, 300, 400, 500.

WE_{Senna} (Collobert et al., 2011): We downloaded this 50-dimension word embedding⁵ trained on Wikipedia over 2 months. We use this embedding as a baseline because it is also trained on Wikipedia.

Given the above semantic representation of words, we derive each entity/category embedding by averaging all word vectors among each entity/category label. If the label is a phrase that has been embedded (e.g., Mikolov’s embedding contains some common phrases), we direct use the phrase embedding.

HEE (Hu et al., 2015): This Hierarchical Entity Embedding method uses the whole Wikipedia hierarchy to train entity embeddings and distance metrics. We used the tools provided by the authors to train entity embeddings on the same Wikipedia corpus as ours. We set the batch size $B = 500$, the initial learning rate $\eta = 0.1$ and decrease it by a factor of 5 whenever the objective value does not increase, and the negative sample size $k = 5$ as suggested by the authors. This method doesn’t provide a way of learning category embeddings. Therefore, it cannot be used in the **nearest neighbor classification** method.

TransE (Bordes et al., 2013): It is a state of the art relational embedding method introduced in Section. 2, which embeds entities and relations at the same time. It can be extended to derive category embeddings, since categories can be seen as a special entity type. To make comparisons fair enough, we adopt three different versions of TransE to derive entity and category embedding to compare with our methods. **TransE₁**: Use entities and categories as two entity types and the direct-labeled-category relation. So the triplets we have is in the form of (entity, direct-labeled-category, category). **TransE₂**: Add

⁴<https://code.google.com/archive/p/word2vec/>

⁵<http://ronan.collobert.com/senna/>

to TransE₁ with the hierarchy structure of wikipedia categories. Namely, add the super-category relation between categories. Therefore we have new triplets in the form of (category, super-category, category). **TransE₃**: To make it use full information as ours, we extend TransE₂ to utilize context entity relationship described in Section. 3.1. Therefore we add new triplets in the form of (entity, context-entity, entity).

HC: To further evaluate the advantage of utilizing category hierarchy in training entity and category embedding, we also compare our Hierarchical Category Embedding (HCE) model with our Category Embedding (CE) model that has no hierarchical information.

5.1.3 Results

In the experiments, we used scikit-learn (Pedregosa et al., 2011) to perform clustering. We tested k-means and hierarchical clustering with different distance metrics (euclidean, cosine) and linkage criterion (ward, complete, average). All these choices along with the vector dimensionality are treated as our models’ hyper-parameters. For selecting hyper-parameters, we randomly split the Battig and Dota datasets to 50% of validation data and 50% of test data, evenly across all categories. We trained all the embeddings (except SENNA) on the same Wikipedia dump and tuned hyper-parameters on the validation set. For experiments on Dota dataset, since the ground truth is contained in our Wikipedia corpus, we deleted all category-entity links contained in Dota dataset from our category hierarchy to train HEE, TransE and HCE embeddings to make comparison fair enough.

Table. 1a shows the experimental results of the **concept clustering** method. It is clear that **hierarchical category embedding** (HCE) model outperforms other methods in all datasets. Our model achieves a purity of 89% on Battig and 89% on DOTA-all.

	Battig	DOTA-single	DOTA-mult	DOTA-all
WE _{Senna} (50,50)	0.74	0.61	0.43	0.45
WE _{Mikolov} (400,400)	0.86	0.83	0.73	0.78
HEE (400,400)	0.82	0.83	0.80	0.81
TransE ₁ (300,250)	0.67	0.71	0.68	0.69
TransE ₂ (400,300)	0.73	0.78	0.75	0.76
TransE ₃ (200,400)	0.43	0.53	0.50	0.51
CE (300,200)	0.84	0.86	0.83	0.85
HCE (400,400)	0.89	0.92	0.88	0.89

(a) Purity of **concept clustering** method

	Battig	DOTA-single	DOTA-mult	DOTA-all
WE _{Senna} (50,50)	0.44	0.52	0.32	0.45
WE _{Mikolov} (400,400)	0.74	0.74	0.67	0.72
HEE	-	-	-	-
TransE ₁ (300,250)	0.66	0.72	0.69	0.71
TransE ₂ (400,300)	0.75	0.80	0.77	0.79
TransE ₃ (300,400)	0.46	0.55	0.52	0.54
CE (200,200)	0.79	0.89	0.85	0.88
HCE (400,400)	0.87	0.93	0.91	0.92

(b) Purity of **NN classification** method

Table 1: Purity of **nearest neighbor (NN) classification** and **concept clustering** methods with different embeddings on Battig and DOTA datasets. The two numbers given in each parentheses are the vector dimensionality among {50,100,200,250,300,400,500} that produce the best results on Battig and Dota validation set respectively. All the purity scores are calculated under such choices of hyper-parameters on the test set.

For word embeddings like Mikolov and Senna, the performance drops a lot from single-word entity categorization to multiple-word entity categorization, because these embeddings mainly contain single words. To get the embeddings of multiple-word, we use the mean word vectors to denote multiple-word embeddings. However, the meaning of a multiple-word is not simply the aggregation of the meaning of the words it contains. The good performance of HEE shows the high quality of its entity embeddings. By incorporating category hierarchy information, TransE₂ gets an advantage over TransE₁. This advantage can also be observed when we incorporate category hierarchy information into CE. However, further incorporating of context entities deteriorates the performance of TransE. This may due to the nature of transition-based models that assume relationships as transitions among entities. They get much noise when introduced with many context entity triplets along with only one entity-to-context-entity relation type.

Table.1b shows the experimental results of the **nearest neighbor (NN) classification** method. The results indicate the feasibility of using category vectors to directly predict the concept categories without clustering entities. Our model achieves a purity of 87% on Battig and 92% on DOTA-all.

5.2 Dataless Classification

5.2.1 Datasets

20Newsgroups Data(20NG):The 20 newsgroups dataset (Lang, 1995) contains about 20,000 newsgroups documents evenly categorized to 20 newsgroups, and further categorized to six super-classes. We use the same label description provided by (Song and Roth, 2014).

RCV1 Dataset: The RCV1 dataset (Lewis et al., 2004) contains 804,414 manually labeled newswire documents, and categorized with respect to three controlled vocabularies: industries, topics and regions. We use topics as our hierarchical classification problem. There are 103 categories including all nodes except for the root in the hierarchy, and the maximum depth is 4. To ease the computational cost of comparison, we follow the chronological split proposed in (Lewis et al., 2004) to use the first 23,149 documents marked as training samples in the dataset. The dataset also provides the name and the description of each category label.

5.2.2 Implementation Details

The baseline of similarity measure between each label and each document is cosine similarity between corresponding ESA vectors. We use ESA with 500 entities, and augment similarity measure by plugging in different embeddings introduced in Section. 5.1.2 with Hungarian sparse vector densification method described in (Song and Roth, 2015). The Hungarian method is a combinatorial optimization algorithm aiming to find an optimal assignment matching the two sides of a bipartite graph on a one-to-one basis. We average over all cosine similarities of entity pairs produced by the Hungarian method to produce similarity between ESA vectors. In addition, we cut off all entity similarities below 0.85 and map them to zero empirically.

For classification, We use the bottom-up pure dataless hierarchical classification algorithm which proved to be superior in (Song and Roth, 2014) and set the threshold δ to be 0.95 empirically.

5.2.3 Results

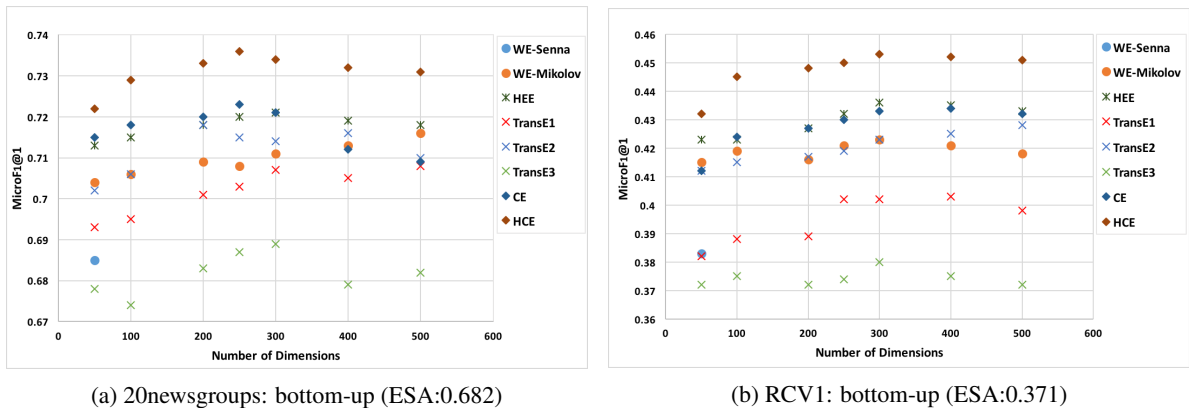


Figure 3: MicroF₁@1 results of ESA augmented with different embeddings for dataless hierarchical classification. It is clear the ESA densified with HCE performs best on dataless hierarchical classification.

We perform dataless hierarchical classification using the same ESA settings, densified with different embedding methods for similarity calculation. Figure. 3 presents the classification performance with different embedding methods in dimensionality of $\{50,100,200,250,300,400,500\}$. It is clear that ESA densified with HCE stably outperform other competitive methods on these two datasets, which indicates the high quality of entity embeddings derived by HCE.

Besides, we observe that entity embedding methods such as TransE₂, HEE and HCE perform better than word embedding methods. This is because ESA represents text as bag-of-entities, and entity embeddings successfully encode entity similarities. Compared with bag-of-words baselines in (Song and Roth, 2014), it is clear that bag-of-entities models can better capture semantic relatedness between documents and short category descriptions.

6 Conclusion

In this paper, we proposed a framework to learn entity and category embeddings to capture semantic relatedness between entities and categories. This framework can incorporate taxonomy hierarchy from large scale knowledge bases. Experiments on both concept categorization and dataless hierarchical classification indicate the potential usage of category and entity embeddings on more other NLP applications.

References

- Marco Baroni and Alessandro Lenci. 2010. Distributional memory: A general framework for corpus-based semantics. *Computational Linguistics*, 36(4):673–721.
- Kurt 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 2008 ACM SIGMOD international conference on Management of data*, pages 1247–1250. ACM.
- Antoine Bordes, Jason Weston, Ronan Collobert, and Yoshua Bengio. 2011. Learning structured embeddings of knowledge bases. In *Conference on Artificial Intelligence*, number EPFL-CONF-192344.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems*, pages 2787–2795.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *The Journal of Machine Learning Research*, 12:2493–2537.
- Christiane Fellbaum. 1998. *WordNet*. Wiley Online Library.
- Evgeniy Gabrilovich and Shaul Markovitch. 2007. Computing semantic relatedness using wikipedia-based explicit semantic analysis. In *IJCAI*, volume 7, pages 1606–1611.
- Siddharth Gopal and Yiming Yang. 2013. Recursive regularization for large-scale classification with hierarchical and graphical dependencies. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 257–265. ACM.
- Zhiting Hu, Poyao Huang, Yuntian Deng, Yingkai Gao, and Eric P Xing. 2015. Entity hierarchy embedding. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP)*, volume 1, pages 1292–1300.
- Zhiting Hu, Gang Luo, Mrinmaya Sachan, Eric Xing, and Zaiqing Nie. 2016. Grounding topic models with knowledge bases. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*.
- Sung Ju Hwang and Leonid Sigal. 2014. A unified semantic embedding: Relating taxonomies and attributes. In *Advances in Neural Information Processing Systems*, pages 271–279.
- Ken Lang. 1995. Newsweeder: Learning to filter netnews. In *Proceedings of the 12th international conference on machine learning*, pages 331–339.
- David D Lewis, Yiming Yang, Tony G Rose, and Fan Li. 2004. Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research*, 5(Apr):361–397.
- Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. In *AAAI*, pages 2181–2187.
- Yang Liu, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. 2015. Topical word embeddings.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Arvind Neelakantan and Ming-Wei Chang. 2015. Inferring missing entity type instances for knowledge base completion: New dataset and methods. *arXiv preprint arXiv:1504.06658*.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2012. Factorizing yago: scalable machine learning for linked data. In *Proceedings of the 21st international conference on World Wide Web*, pages 271–280. ACM.

- Alberto Paccanaro and Geoffrey E Hinton. 2001. Learning distributed representations of concepts using linear relational embedding. *Knowledge and Data Engineering, IEEE Transactions on*, 13(2):232–244.
- Alexandre Passos, Vineet Kumar, and Andrew McCallum. 2014. Lexicon infused phrase embeddings for named entity resolution. *arXiv preprint arXiv:1404.5367*.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in python. *The Journal of Machine Learning Research*, 12:2825–2830.
- Klaus Rothenhäusler and Hinrich Schütze. 2009. Unsupervised classification with dependency based word spaces. In *Proceedings of the Workshop on Geometrical Models of Natural Language Semantics*, pages 17–24. Association for Computational Linguistics.
- Yangqiu Song and Dan Roth. 2014. On dataless hierarchical text classification. In *AAAI*, pages 1579–1585.
- Yangqiu Song and Dan Roth. 2015. Unsupervised sparse vector densification for short text similarity. *Proc. North Am. Chapter Assoc. Computat. Linguistics*, pages 1275–1280.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(2579-2605):85.
- Nakul Verma, Dhruv Mahajan, Sundararajan Sellamanickam, and Vinod Nair. 2012. Learning hierarchical similarity metrics. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2280–2287. IEEE.
- Kilian Q Weinberger and Olivier Chapelle. 2009. Large margin taxonomy embedding for document categorization. In *Advances in Neural Information Processing Systems*, pages 1737–1744.
- Ikuya Yamada, Hiroyuki Shindo, Hideaki Takeda, and Yoshiyasu Takefuji. 2016. Joint learning of the embedding of words and entities for named entity disambiguation. *arXiv preprint arXiv:1601.01343*.
- Yiming Yang. 1999. An evaluation of statistical approaches to text categorization. *Information retrieval*, 1(1-2):69–90.
- Wenpeng Yin and Hinrich Schütze. 2014. An exploration of embeddings for generalized phrases. In *ACL (Student Research Workshop)*, pages 41–47.
- Mo Yu and Mark Dredze. 2015. Learning composition models for phrase embeddings. *Transactions of the Association for Computational Linguistics*, 3:227–242.
- Hao Zhang, Zhiting Hu, Yuntian Deng, Mrinmaya Sachan, Zhicheng Yan, and Eric P Xing. 2016. Learning concept taxonomies from multi-modal data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.

A The DOTA dataset: 300 single-word entities and 150 multi-word entities from 15 Wikipedia Categories

Category	Entities
beverages	juice, beer, milk, coffee, tea, cocktail, wine, liqueur, sake, vodka, mead, sherry, brandy, gin, rum, latte, whisky, cider, gose, rompopo, orange juice, masala chai, green tea, black tea, herbal tea, coconut milk, corn syrup, soy milk, rose water, hyeonmi cha
sports	bowling, football, aerobics, hockey, karate, korfbal, handball, floorball, skiing, cycling, racing, softball, shooting, netball, snooker, powerlifting, jumping, wallball, volleyball, snowboarding, table tennis, floor hockey, olympic sports, wheelchair basketball, crab soccer, indoor soccer, table football, roller skating, vert skating, penny football
emotions	love, anxiety, empathy, fear, envy, loneliness, shame, anger, annoyance, happiness, jealousy, apathy, resentment, frustration, belongingness, sympathy, pain, worry, hostility, sadness, broken heart, panic disorder, sexual desire, falling in love, emotional conflict, learned helplessness, chronic stress, anxiety sensitivity, mental breakdown, bike rage
weather	cloud, wind, thunderstorm, fog, snow, wave, blizzard, sunlight, tide, virga, lightning, cyclone, whirlwind, sunset, dust, frost, flood, thunder, supercooling, fahrenheit, acid rain, rain and snow mixed, cumulus cloud, winter storm, blowing snow, geomagnetic storm, blood rain, fire whirl, pulse storm, dirty thunderstorm
landforms	lake, waterfall, stream, river, wetland, marsh, valley, pond, sandstone, mountain, cave, swamp, ridge, plateau, cliff, grassland, glacier, hill, bay, island, glacial lake, drainage basin, river delta, stream bed, vernal pool, salt marsh, proglacial lake, mud volcano, pit crater, lava lake
trees	wood, oak, pine, evergreen, willow, vine, shrub, birch, beech, maple, pear, fir, pinales, lauraceae, sorbus, buxus, acacia, rhamnaceae, fagales, sycamore, alhambra creek, alstonia boonei, atlantic hazelwood, bee tree, blood banana, datun sahib, druid oak, new year tree, heart pine, fan palm
algebra	addition, multiplication, exponentiation, tetration, polynomial, calculus, permutation, subgroup, integer, monomial, bijection, homomorphism, determinant, sequence, permanent, homotopy, subset, factorization, associativity, commutativity, real number, abstract algebra, convex set, prime number, complex analysis, natural number, complex number, lie algebra, identity matrix, set theory
geometry	trigonometry, circle, square, polyhedron, surface, sphere, cube, icosahedron, hemipolyhedron, digon, midpoint, centroid, octadecagon, curvature, curve, zonohedron, cevian, orthant, cuboctahedron, midsphere, regular polygon, uniform star polyhedron, isogonal figure, icosahedral symmetry, hexagonal bipyramid, snub polyhedron, homothetic center, geometric shape, bragg plane, affine plane
fish	goldfish, gourami, koi, cobitidae, tetra, goby, danio, wrasse, acanthuridae, anchovy, carp, catfish, cod, eel, flatfish, perch, pollock, salmon, triggerfish, herring, cave catfish, coachwhip ray, dwarf cichlid, moray eel, coastal fish, scissortail rasbora, flagtail pipefish, armoured catfish, hawaiian flagtail, pelagic fish
dogs	spaniel, foxhound, bloodhound, beagle, pekingese, weimaraner, collie, terrier, poodle, puppy, otterhound, labradoodle, puggle, eurasier, drever, brindle, schnoodle, bandog, leonberger, cockapoo, golden retriever, tibetan terrier, bull terrier, welsh springer spaniel, hunting dog, bearded collie, picardy spaniel, afghan hound, brittany dog, redbone coonhound
music	jazz, blues, song, choir, opera, rhythm, lyrics, melody, harmony, concert, comedy, violin, drum, piano, drama, cello, composer, musician, drummer, pianist, hip hop, classical music, electronic music, folk music, dance music, musical instrument, disc jockey, popular music, sheet music, vocal music
politics	democracy, law, government, liberalism, justice, policy, rights, utilitarianism, election, capitalism, ideology, egalitarianism, debate, regime, globalism, authoritarianism, monarchism, anarchism, communism, individualism, freedom of speech, political science, public policy, civil society, international law, social contract, election law, social justice, global justice, group conflict
philosophy	ethics, logic, ontology, aristotle, plato, rationalism, platonism, relativism, existence, truth, positivism, meta-logic, subjectivism, idealism, materialism, aesthetics, probabilism, monism, truth, existence, western philosophy, contemporary philosophy, cognitive science, logical truth, ancient philosophy, universal mind, visual space, impossible world, theoretical philosophy, internal measurement
linguistics	syntax, grammar, semantics, lexicon, speech, phonetics, vocabulary, phoneme, lexicography, language, pragmatics, orthography, terminology, pronoun, noun, verb, pronunciation, lexicology, metalinguistics, paleolinguistics, language death, historical linguistics, dependency grammar, noun phrase, comparative linguistics, word formation, cognitive semantics, syntactic structures, auxiliary verb, computational semantics
vehicles	truck, car, aircraft, minibus, motorcycle, microvan, bicycle, tractor, microcar, van, ship, helicopter, airplane, towing, velomobile, rocket, train, bus, gyrocar, cruiser, container ship, school bus, road train, tow truck, audi a6, garbage truck, hydrogen tank, light truck, compressed air car, police car