

Improved Relation Extraction with Feature-Rich Compositional Embedding Models

Matthew R. Gormley^{1*} and Mo Yu^{2*} and Mark Dredze¹

¹Human Language Technology Center of Excellence

Center for Language and Speech Processing

Johns Hopkins University, Baltimore, MD, 21218

²Machine Intelligence and Translation Lab

Harbin Institute of Technology, Harbin, China

gflfof@gmail.com, {mrg, mdredze}@cs.jhu.edu

Abstract

Compositional embedding models build a representation (or embedding) for a linguistic structure based on its component word embeddings. We propose a Feature-rich Compositional Embedding Model (FCM) for relation extraction that is expressive, generalizes to new domains, and is easy-to-implement. The key idea is to combine both (unlexicalized) hand-crafted features with learned word embeddings. The model is able to directly tackle the difficulties met by traditional compositional embeddings models, such as handling arbitrary types of sentence annotations and utilizing global information for composition. We test the proposed model on two relation extraction tasks, and demonstrate that our model outperforms both previous compositional models and traditional feature rich models on the ACE 2005 relation extraction task, and the SemEval 2010 relation classification task. The combination of our model and a log-linear classifier with hand-crafted features gives state-of-the-art results. We made our implementation available for general use¹.

1 Introduction

Two common NLP feature types are lexical properties of words and unlexicalized linguistic/structural interactions between words. Prior work on relation extraction has extensively studied how to design such features by combining *discrete* lexical properties (e.g. the identity of a word,

its lemma, its morphological features) with aspects of a word’s linguistic context (e.g. whether it lies between two entities or on a dependency path between them). While these help learning, they make generalization to unseen words difficult. An alternative approach to capturing lexical information relies on *continuous* word embeddings² as representative of words but generalizable to new words. Embedding features have improved many tasks, including NER, chunking, dependency parsing, semantic role labeling, and relation extraction (Miller et al., 2004; Turian et al., 2010; Koo et al., 2008; Roth and Woodsend, 2014; Sun et al., 2011; Plank and Moschitti, 2013; Nguyen and Grishman, 2014). Embeddings can capture lexical information, but alone they are insufficient: in state-of-the-art systems, they are used alongside features of the broader linguistic context.

In this paper, we introduce a compositional model that combines unlexicalized linguistic context and word embeddings for relation extraction, a task in which contextual feature construction plays a major role in generalizing to unseen data. Our model allows for the composition of embeddings with arbitrary linguistic structure, as expressed by hand crafted features. In the following sections, we begin with a precise construction of compositional embeddings using word embeddings in conjunction with unlexicalized features. Various feature sets used in prior work (Turian et al., 2010; Nguyen and Grishman, 2014; Hermann et al., 2014; Roth and Woodsend, 2014) are cap-

²Such embeddings have a long history in NLP, including term-document frequency matrices and their low-dimensional counterparts obtained by linear algebra tools (LSA, PCA, CCA, NNMF), Brown clusters, random projections and vector space models. Recently, neural networks / deep learning have provided several popular methods for obtaining such embeddings.

*Gormley and Yu contributed equally.

¹<https://github.com/mgormley/pacaya>

Class	M_1	M_2	Sentence Snippet
(1) ART(M_1, M_2)	a man	a taxicab	A man <i>driving</i> what appeared to be a taxicab
(2) PART-WHOLE(M_1, M_2)	the southern suburbs	Baghdad	direction of the southern suburbs <i>of</i> Baghdad
(3) PHYSICAL(M_2, M_1)	the united states	284 people	in the united states , 284 people <i>died</i>

Table 1: Examples from ACE 2005. In (1) the word “driving” is a strong indicator of the relation ART^3 between M_1 and M_2 . A feature that depends on the embedding for this context word could generalize to other lexical indicators of the same relation (e.g. “operating”) that don’t appear with ART during training. But lexical information alone is insufficient; relation extraction requires the identification of lexical roles: *where* a word appears structurally in the sentence. In (2), the word “of” between “suburbs” and “Baghdad” suggests that the first entity is part of the second, yet the earlier occurrence after “direction” is of no significance to the relation. Even finer information can be expressed by a word’s role on the dependency path between entities. In (3) we can distinguish the word “died” from other irrelevant words that don’t appear between the entities.

tured as special cases of this construction. Adding these compositional embeddings directly to a standard log-linear model yields a special case of our full model. We then treat the word embeddings as parameters giving rise to our powerful, efficient, and easy-to-implement *log-bilinear model*. The model capitalizes on arbitrary types of linguistic annotations by better utilizing features associated with substructures of those annotations, including global information. We choose features to promote different properties and to distinguish different functions of the input words.

The full model involves three stages. First, it decomposes the annotated sentence into substructures (i.e. a word and associated annotations). Second, it extracts features for each substructure (word), and combines them with the word’s embedding to form a *substructure embedding*. Third, we sum over substructure embeddings to form a composed *annotated sentence embedding*, which is used by a final softmax layer to predict the output label (relation).

The result is a state-of-the-art relation extractor for unseen domains from ACE 2005 (Walker et al., 2006) and the relation classification dataset from SemEval-2010 Task 8 (Hendrickx et al., 2010).

Contributions This paper makes several contributions, including:

1. We introduce the FCM, a new compositional embedding model for relation extraction.
2. We obtain the best reported results on ACE-2005 for coarse-grained relation extraction in the cross-domain setting, by combining FCM with a log-linear model.
3. We obtain results on on SemEval-2010 Task 8 competitive with the best reported results.

Note that other work has already been published that builds on the FCM, such as Hashimoto et al. (2015), Nguyen and Grishman (2015), dos Santos

³In ACE 2005, ART refers to a relation between a person and an artifact; such as a user, owner, inventor, or manufacturer relationship

et al. (2015), Yu and Dredze (2015) and Yu et al. (2015). Additionally, we have extended FCM to incorporate a low-rank embedding of the features (Yu et al., 2015), which focuses on fine-grained relation extraction for ACE and ERE. This paper obtains better results than the low-rank extension on ACE coarse-grained relation extraction.

2 Relation Extraction

In relation extraction we are given a sentence as input with the goal of identifying, for all pairs of entity mentions, what relation exists between them, if any. For each pair of entity mentions in a sentence S , we construct an instance (y, \mathbf{x}) , where $\mathbf{x} = (M_1, M_2, S, A)$. $S = \{w_1, w_2, \dots, w_n\}$ is a sentence of length n that expresses a relation of type y between two entity mentions M_1 and M_2 , where M_1 and M_2 are sequences of words in S . A is the associated annotations of sentence S , such as part-of-speech tags, a dependency parse, and named entities. We consider directed relations: for a relation type Rel , $y = Rel(M_1, M_2)$ and $y' = Rel(M_2, M_1)$ are different relations. Table 1 shows ACE 2005 relations, and has a strong label bias towards negative examples. We also consider the task of relation *classification* (SemEval), where the number of negative examples is artificially reduced.

Embedding Models Word embeddings and compositional embedding models have been successfully applied to a range of NLP tasks, however the applications of these embedding models to relation extraction are still limited. Prior work on relation classification (e.g. SemEval 2010 Task 8) has focused on short sentences with at most one relation per sentence (Socher et al., 2012; Zeng et al., 2014). For relation extraction, where negative examples abound, prior work has assumed that only the named entity boundaries and not their types were available (Plank and Moschitti, 2013; Nguyen et al., 2015). Other work has as-

sumed that the order of two entities in a relation are given while the relation type itself is unknown (Nguyen and Grishman, 2014; Nguyen and Grishman, 2015). The standard relation extraction task, as adopted by ACE 2005 (Walker et al., 2006), uses long sentences containing multiple named entities with known types⁴ and unknown relation directions. We are the first to apply neural language model embeddings to this task.

Motivation and Examples Whether a word is indicative of a relation depends on multiple properties, which may relate to its context within the sentence. For example, whether the word is in-between the entities, on the dependency path between them, or to their left or right may provide additional complementary information. Illustrative examples are given in Table 1 and provide the motivation for our model. In the next section, we will show how we develop informative representations capturing both the semantic information in word embeddings and the contextual information expressing a word’s role relative to the entity mentions. We are the first to incorporate all of this information at once. The closest work is that of Nguyen and Grishman (2014), who use a log-linear model for relation extraction with embeddings as features for only the entity heads. Such embedding features are insensitive to the broader contextual information and, as we show, are not sufficient to elicit the word’s role in a relation.

3 A Feature-rich Compositional Embedding Model for Relations

We propose a general framework to construct an embedding of a sentence with annotations on its component words. While we focus on the relation extraction task, the framework applies to any task that benefits from both embeddings and typical hand-engineered lexical features.

3.1 Combining Features with Embeddings

We begin by describing a precise method for constructing substructure embeddings and annotated sentence embeddings from existing (usually unlexicalized) features and embeddings. Note that these embeddings can be included directly in a log-linear model as features—doing so results in

⁴Since the focus of this paper is relation extraction, we adopt the evaluation setting of prior work which uses gold named entities to better facilitate comparison.

a special case of our full model presented in the next subsection.

An annotated sentence is first decomposed into substructures. The type of substructures can vary by task; for relation extraction we consider one substructure per word⁵. For each substructure in the sentence we have a hand-crafted feature vector f_{w_i} and a dense embedding vector e_{w_i} . We represent each substructure as the outer product \otimes between these two vectors to produce a matrix, herein called a **substructure embedding**: $h_{w_i} = f_{w_i} \otimes e_{w_i}$. The features f_{w_i} are based on the local context in S and annotations in A , which can include global information about the annotated sentence. These features allow the model to promote different properties and to distinguish different functions of the words. Feature engineering can be task specific, as relevant annotations can change with regards to each task. In this work we utilize unlexicalized binary features common in relation extraction. Figure 1 depicts the construction of a sentence’s substructure embeddings.

We further sum over the substructure embeddings to form an **annotated sentence embedding**:

$$e_x = \sum_{i=1}^n f_{w_i} \otimes e_{w_i} \quad (1)$$

When both the hand-crafted features and word embeddings are treated as inputs, as has previously been the case in relation extraction, this annotated sentence embedding can be used directly as the features of a log-linear model. In fact, we find that the feature sets used in prior work for many other NLP tasks are special cases of this simple construction (Turian et al., 2010; Nguyen and Grishman, 2014; Hermann et al., 2014; Roth and Woodsend, 2014). This highlights an important connection: when the word embeddings are constant, our constructions of substructure and annotated sentence embeddings are just specific forms of polynomial (specifically quadratic) feature combination—hence their commonality in the literature. Our experimental results suggest that such a construction is more powerful than directly including embeddings into the model.

3.2 The Log-Bilinear Model

Our full log-bilinear model first forms the substructure and annotated sentence embeddings from

⁵We use words as substructures for relation extraction, but use the general terminology to maintain model generality.




Figure 1: Example construction of substructure embeddings. Each substructure is a word w_i in S , augmented by the target entity information and related information from annotation A (e.g. a dependency tree). We show the factorization of the annotated sentence into substructures (left), the concatenation of the substructure embeddings for the sentence (middle), and a single substructure embedding from that concatenation (right). The annotated sentence embedding (not shown) would be the sum of the substructure embeddings, as opposed to their concatenation.

the previous subsection. The model uses its parameters to score the annotated sentence embedding and uses a softmax to produce an output label. We call the entire model the **Feature-rich Compositional Embedding Model (FCM)**.

Our task is to determine the label y (relation) given the instance $\mathbf{x} = (M_1, M_2, S, A)$. We formulate this as a probability.

$$P(y|\mathbf{x}; T, \mathbf{e}) = \frac{\exp(\sum_{i=1}^n T_y \odot (f_{w_i} \otimes e_{w_i}))}{Z(\mathbf{x})} \quad (2)$$

where \odot is the ‘matrix dot product’ or Frobenious inner product of the two matrices. The normalizing constant which sums over all possible output labels $y' \in L$ is given by $Z(\mathbf{x}) = \sum_{y' \in L} \exp(\sum_{i=1}^n T_{y'} \odot (f_{w_i} \otimes e_{w_i}))$. The parameters of the model are the word embeddings \mathbf{e} for each word type and a list of weight matrix $T = [T_y]_{y \in L}$ which is used to score each label y . The model is log-bilinear⁶ (i.e. log-quadratic) since we recover a log-linear model by fixing either \mathbf{e} or T . We study both the full log-bilinear and the log-linear model obtained by fixing the word embeddings.

3.3 Discussion of the Model

Substructure Embeddings Similar words (i.e. those with similar embeddings) with similar functions in the sentence (i.e. those with similar features) will have similar matrix representations. To understand our selection of the outer product, consider the example in Fig. 1. The word “driving” can indicate the *ART* relation if it appears on the

dependency path between M_1 and M_2 . Suppose the third feature in f_{w_i} indicates this on-path feature. Our model can now learn parameters which give the third row a high weight for the *ART* label. Other words with embeddings similar to “driving” that appear on the dependency path between the mentions will similarly receive high weight for the *ART* label. On the other hand, if the embedding is similar but is not on the dependency path, it will have 0 weight. Thus, our model generalizes its model parameters across words with similar embeddings only when they share similar functions in the sentence.

Smoothed Lexical Features Another intuition about the selection of outer product is that it is actually a smoothed version of traditional lexical features used in classical NLP systems. Consider a lexical feature $f = u \wedge w$, which is a conjunction (logic-and) between non-lexical property u and lexical part (word) w . If we represent w as a one-hot vector, then the outer product exactly recovers the original feature f . Then if we replace the one-hot representation with its word embedding, we get the current form of our FCM. Therefore, our model can be viewed as a smoothed version of lexical features, which keeps the expressive strength, and uses embeddings to generalize to low frequency features.

Time Complexity Inference in FCM is much faster than both CNNs (Collobert et al., 2011) and RNNs (Socher et al., 2013b; Bordes et al., 2012). FCM requires $O(snd)$ products on average with sparse features, where s is the average number of per-word non-zero feature values, n is the length of the sentence, and d is the dimension of word embedding. In contrast, CNNs and RNNs usually

⁶Other popular log-bilinear models are the log-bilinear language models (Mnih and Hinton, 2007; Mikolov et al., 2013).

have complexity $O(C \cdot nd^2)$, where C is a model dependent constant.

4 Hybrid Model

We present a hybrid model which combines the FCM with an existing log-linear model. We do so by defining a new model:

$$p_{\text{FCM}+\text{loglin}}(y|x) = \frac{1}{Z} p_{\text{FCM}}(y|x) p_{\text{loglin}}(y|x) \quad (3)$$

The log-linear model has the usual form: $p_{\text{loglin}}(y|x) \propto \exp(\boldsymbol{\theta} \cdot \mathbf{f}(x, y))$, where $\boldsymbol{\theta}$ are the model parameters and $\mathbf{f}(x, y)$ is a vector of features. The integration treats each model as a providing a score which we multiply together. The constant Z ensures a normalized distribution.

5 Training

FCM training optimizes a cross-entropy objective:

$$\ell(D; T, \mathbf{e}) = \sum_{(\mathbf{x}, y) \in D} \log P(y|\mathbf{x}; T, \mathbf{e})$$

where D is the set of all training data and \mathbf{e} is the set of word embeddings. To optimize the objective, for each instance (y, \mathbf{x}) we perform stochastic training on the loss function $\ell = \ell(y, \mathbf{x}; T, \mathbf{e}) = \log P(y|\mathbf{x}; T, \mathbf{e})$. The gradients of the model parameters are obtained by backpropagation (i.e. repeated application of the chain rule). We define the vector $\mathbf{s} = [\sum_i T_y \odot (f_{w_i} \otimes e_{w_i})]_{1 \leq y \leq L}$, which yields

$$\frac{\partial \ell}{\partial \mathbf{s}} = \left[(I[y' = y] - P(y'|\mathbf{x}; T, \mathbf{e}))_{1 \leq y' \leq L} \right]^T,$$

where the indicator function $I[x]$ equals 1 if x is true and 0 otherwise. We have the following gradients: $\frac{\partial \ell}{\partial T} = \frac{\partial \ell}{\partial \mathbf{s}} \otimes \sum_{i=1}^n f_{w_i} \otimes e_{w_i}$, which is equivalent to:

$$\frac{\partial \ell}{\partial T_{y'}} = (I[y = y'] - P(y'|\mathbf{x}; T, \mathbf{e})) \cdot \sum_{i=1}^n f_{w_i} \otimes e_{w_i}.$$

When we treat the word embeddings as parameters (i.e. the log-bilinear model), we also fine-tune the word embeddings with the FCM model:

$$\frac{\partial \ell}{\partial e_w} = \sum_{i=1}^n \left(\left(\sum_y \frac{\partial \ell}{\partial \mathbf{s}_y} T_y \right) \cdot f_i \right) \cdot I[w_i = w].$$

As is common in deep learning, we initialize these embeddings from an neural language model and then *fine-tune* them for our supervised task. The training process for the hybrid model (§ 4) is also easily done by backpropagation since each sub-model has separate parameters.

Set	Template
HeadEmb	$\{I[i = h_1], I[i = h_2]\}$ $(w_i \text{ is head of } M_1/M_2) \times \{\phi, t_{h_1}, t_{h_2}, t_{h_1} \oplus t_{h_2}\}$
Context	$I[i = h_1 \pm 1]$ (left/right token of w_{h_1}) $I[i = h_2 \pm 1]$ (left/right token of w_{h_2})
In-between	$I[i > h_1] \& I[i < h_2]$ (in between) $\times \{\phi, t_{h_1}, t_{h_2}, t_{h_1} \oplus t_{h_2}\}$
On-path	$I[w_i \in P]$ (on path) $\times \{\phi, t_{h_1}, t_{h_2}, t_{h_1} \oplus t_{h_2}\}$

Table 2: Feature sets used in FCM.

6 Experimental Settings

Features Our FCM features (Table 2) use a feature vector f_{w_i} over the word w_i , the two target entities M_1, M_2 , and their dependency path. Here h_1, h_2 are the indices of the two head words of M_1, M_2 , \times refers to the Cartesian product between two sets, t_{h_1} and t_{h_2} are entity types (named entity tags for ACE 2005 or WordNet supertags for SemEval 2010) of the head words of two entities, and ϕ stands for the empty feature. \oplus refers to the conjunction of two elements. The *In-between* features indicate whether a word w_i is in between two target entities, and the *On-path* features indicate whether the word is on the dependency path, on which there is a set of words P , between the two entities.

We also use the target entity type as a feature. Combining this with the basic features results in more powerful compound features, which can help us better distinguish the functions of word embeddings for predicting certain relations. For example, if we have a person and a vehicle, we know it will be more likely that they have an *ART* relation. For the *ART* relation, we introduce a corresponding weight vector, which is closer to lexical embeddings similar to the embedding of “drive”.

All linguistic annotations needed for features (POS, chunks⁷, parses) are from Stanford CoreNLP (Manning et al., 2014). Since SemEval does not have gold entity types we obtained WordNet and named entity tags using Ciaramita and Altun (2006). For all experiments we use 200-d word embeddings trained on the NYT portion of the Gigaword 5.0 corpus (Parker et al., 2011), with word2vec (Mikolov et al., 2013). We use the CBOW model with negative sampling (15 negative words). We set a window size $c=5$, and remove types occurring less than 5 times.

Models We consider several methods. (1) FCM in isolation without fine-tuning. (2) FCM in isolation with fine-tuning (i.e. trained as a log-bilinear

⁷Obtained from the constituency parse using the CONLL 2000 chunking converter (Perl script).

model). (3) A log-linear model with a rich binary feature set from Sun et al. (2011) (Baseline)—this consists of all the baseline features of Zhou et al. (2005) plus several additional carefully-chosen features that have been highly tuned for ACE-style relation extraction over years of research. We exclude the Country gazetteer and WordNet features from Zhou et al. (2005). The two remaining methods are hybrid models that integrate FCM as a sub-model within the log-linear model (§ 4). We consider two combinations. (4) The feature set of Nguyen and Grishman (2014) obtained by using the embeddings of heads of two entity mentions (+HeadOnly). (5) Our full FCM model (+FCM). All models use L2 regularization tuned on dev data.

6.1 Datasets and Evaluation

ACE 2005 We evaluate our relation extraction system on the English portion of the ACE 2005 corpus (Walker et al., 2006).⁸ There are 6 domains: Newswire (nw), Broadcast Conversation (bc), Broadcast News (bn), Telephone Speech (cts), Usenet Newsgroups (un), and Weblogs (wl). Following prior work we focus on the domain adaptation setting, where we train on one set (the union of the news domains (bn+nw), tune hyperparameters on a dev domain (half of bc) and evaluate on the remainder (cts, wl, and the remainder of bc) (Plank and Moschitti, 2013; Nguyen and Grishman, 2014). We assume that gold entity spans *and* types are available for train and test. We use all pairs of entity mentions to yield 43,518 total relations in the training set. We report precision, recall, and F1 for relation extraction. While it is not our focus, for completeness we include results with unknown entity types following Plank and Moschitti (2013) (Appendix 1).

SemEval 2010 Task 8 We evaluate on the SemEval 2010 Task 8 dataset⁹ (Hendrickx et al., 2010) to compare with other compositional models and highlight the advantages of FCM. This task is to determine the relation type (or no relation) between two entities in a sentence. We adopt the setting of Socher et al. (2012). We use 10-fold

⁸Many relation extraction systems evaluate on the ACE 2004 corpus (Mitchell et al., 2005). Unfortunately, the most common convention is to use 5-fold cross validation, treating the entirety of the dataset as both train *and* evaluation data. Rather than continuing to overfit this data by perpetuating the cross-validation convention, we instead focus on ACE 2005.

⁹http://docs.google.com/View?docid=dfvx49s_36c28v9pmw

cross validation on the training data to select hyperparameters and do regularization by early stopping. The learning rates for FCM with/without fine-tuning are 5e-3 and 5e-2 respectively. We report macro-F1 and compare to previously published results.

7 Results

ACE 2005 Despite FCM’s (1) simple feature set, it is competitive with the log-linear baseline (3) on out-of-domain test sets (Table 3). In the typical gold entity spans and types setting, both Plank and Moschitti (2013) and Nguyen and Grishman (2014) found that they were unable to obtain improvements by adding embeddings to baseline feature sets. By contrast, we find that on all domains the combination baseline + FCM (5) obtains the highest F1 and significantly outperforms the other baselines, *yielding the best reported results for this task*. We found that fine-tuning of embeddings (2) did not yield improvements on our out-of-domain development set, in contrast to our results below for SemEval. We suspect this is because fine-tuning allows the model to overfit the training domain, which then hurts performance on the unseen ACE test domains. Accordingly, Table 3 shows only the log-linear model.

Finally, we highlight an important contrast between FCM (1) and the log-linear model (3): the latter uses over 50 feature templates based on a POS tagger, dependency parser, chunker, and constituency parser. FCM uses only a dependency parse but still obtains better results (Avg. F1).

SemEval 2010 Task 8 Table 4 shows FCM compared to the best reported results from the SemEval-2010 Task 8 shared task and several other compositional models.

For the FCM we considered two feature sets. We found that using NE tags instead of WordNet tags helps with fine-tuning but hurts without. This may be because the set of WordNet tags is larger making the model more expressive, but also introduces more parameters. When the embeddings are fixed, they can help to better distinguish different functions of embeddings. But when fine-tuning, it becomes easier to over-fit. Alleviating over-fitting is a subject for future work (§ 9).

With either WordNet or NER features, FCM achieves better performance than the RNN and MVRNN. With NER features and fine-tuning, it outperforms a CNN (Zeng et al., 2014) and also

Model	bc			cts			wl			Avg. F1
	P	R	F1	P	R	F1	P	R	F1	
(1) FCM only (ST)	66.56	57.86	61.90	65.62	44.35	52.93	57.80	44.62	50.36	55.06
(3) Baseline (ST)	74.89	48.54	58.90	74.32	40.26	52.23	63.41	43.20	51.39	54.17
(4) + HeadOnly (ST)	70.87	50.76	59.16	71.16	43.21	53.77	57.71	42.92	49.23	54.05
(5) + FCM (ST)	74.39	55.35	63.48	74.53	45.01	56.12	65.63	47.59	55.17	58.26

Table 3: Comparison of models on ACE 2005 out-of-domain test sets. Baseline + HeadOnly is our reimplementation of the features of Nguyen and Grishman (2014).

Classifier	Features	F1
SVM (Rink and Harabagiu, 2010) (Best in SemEval2010)	POS, prefixes, morphological, WordNet, dependency parse, Levin classed, ProBank, FrameNet, NomLex-Plus, Google n-gram, paraphrases, TextRunner	82.2
RNN	word embedding, syntactic parse	74.8
RNN + linear	word embedding, syntactic parse, POS, NER, WordNet	77.6
MVRNN	word embedding, syntactic parse	79.1
MVRNN + linear	word embedding, syntactic parse, POS, NER, WordNet	82.4
CNN (Zeng et al., 2014)	word embedding, WordNet	82.7
CR-CNN (log-loss)	word embedding	82.7
CR-CNN (ranking-loss)	word embedding	84.1
RelEmb (word2vec embedding)	word embedding	81.8
RelEmb (task-spec embedding)	word embedding	82.8
RelEmb (task-spec embedding) + linear	word embedding, dependency paths, WordNet, NE	83.5
DepNN	word embedding, dependency paths	82.8
DepNN + linear	word embedding, dependency paths, WordNet, NER	83.6
(1) FCM (log-linear)	word embedding, dependency parse, WordNet	82.0
	word embedding, dependency parse, NER	81.4
(2) FCM (log-bilinear)	word embedding, dependency parse, WordNet	82.5
	word embedding, dependency parse, NER	83.0
(5) FCM (log-linear) + linear (Hybrid)	word embedding, dependency parse, WordNet	83.1
	word embedding, dependency parse, NER	83.4

Table 4: Comparison of FCM with previously published results for SemEval 2010 Task 8.

the combination of an embedding model and a traditional log-linear model (RNN/MVRNN + linear) (Socher et al., 2012). As with ACE, FCM uses less linguistic resources than many close competitors (Rink and Harabagiu, 2010).

We also compared to concurrent work on enhancing the compositional models with task-specific information for relation classification, including Hashimoto et al. (2015) (RelEmb), which trained task-specific word embeddings, and dos Santos et al. (2015) (CR-CNN), which proposed a task-specific ranking-based loss function. Our Hybrid methods (FCM + linear) get comparable results to theirs. Note that their base compositional model results without any task-specific enhancements, i.e. RelEmb with word2vec embeddings and CR-CNN with log-loss, are still lower than the best FCM result. We believe that FCM can be also improved with these task-specific enhancements, e.g. replacing the word embeddings to the task-specific ones from (Hashimoto et al., 2015) increases the result to 83.7% (see §7.2 for details). We leave the application of ranking-based loss to future work.

Finally, a concurrent work (Liu et al., 2015) proposes DepNN, which builds representations for the dependency path (and its attached subtrees) between two entities by applying recursive and convolutional neural networks successively. Compared to their model, our FCM achieves comparable results. Of note, our FCM and the RelEmb are also the most efficient models among all above compositional models since they have linear time complexity with respect to the dimension of embeddings.

7.1 Effects of the embedding sub-models

We next investigate the effects of different types of features on FCM using ablation tests on ACE 2005 (Table 5.) We focus on FCM alone with the feature templates of Table 2. Additionally, we show results of using *only* the head embedding features from Nguyen and Grishman (2014) (HeadOnly). Not surprisingly, the HeadOnly model performs poorly (F1 score = 14.30%), showing the importance of our rich binary feature set. Among all the features templates, removing HeadEmb results in the largest degradation. The second most im-

Feature Set	Prec	Rec	F1
HeadOnly	31.67	9.24	14.30
FCM	69.17	56.73	62.33
-HeadEmb	66.06	47.00	54.92
-Context	70.89	55.27	62.11
-In-between	66.39	51.86	58.23
-On-path	69.23	53.97	60.66
FCM-EntityTypes	71.33	34.68	46.67

Table 5: Ablation test of FCM on development set.

portant feature template is `In-between`, while `Context` features have little impact. Removing all entity type features (t_{h_i}) does significantly worse than the full model, showing the value of our entity type features.

7.2 Effects of the word embeddings

Good word embeddings are critical for both FCM and other compositional models. In this section, we show the results of FCM with embeddings used to initialize other recent state-of-the-art models. Those embeddings include the 300- d baseline embeddings trained on English Wikipedia (w2v-enwiki-d300) and the 100- d task-specific embeddings (task-specific-d100)¹⁰ from the RelEmb paper (Hashimoto et al., 2015), the 400- d embeddings from the CR-CNN paper (dos Santos et al., 2015). Moreover, we list the best result (DepNN) in Liu et al. (2015), which uses the same embeddings as ours. Table 6 shows the effects of word embeddings on FCM and provides relative comparisons between FCM and the other state-of-the-art models. We use the same hyperparameters and number of iterations in Table 4.

The results show that using different embeddings to initialize FCM can improve F1 beyond our previous results. We also find that increasing the dimension of the word embeddings does not necessarily lead to better results due to the problem of over-fitting (e.g. w2v-enwiki-d400 vs. w2v-enwiki-d300). With the same initial embeddings, FCM usually gets better results without any changes to the hyperparameters than the competing model, further confirming the advantage of FCM at the model-level as discussed under Table 4. The only exception is the DepNN model, which gets better result than FCM on the same embeddings. The task-specific embeddings from (Hashimoto et al., 2015) leads to the best performance (an improvement of 0.7%). This observa-

¹⁰In the task-specific setting, FCM will represent entity words and context words with separate sets of embeddings.

Embeddings	Model	F1
w2v-enwiki-d300	RelEmb	81.8
	(2) FCM (log-bilinear)	83.4
task-specific-d100	RelEmb	82.8
	RelEmb+linear (2) FCM (log-bilinear)	83.5 83.7
w2v-enwiki-d400	CR-CNN	82.7
	(2) FCM (log-bilinear)	83.0
w2v-nyt-d200	DepNN	83.6
	(2) FCM (log-bilinear)	83.0

Table 6: Evaluation of FCMs with different word embeddings on SemEval 2010 Task 8.

tion suggests that the other compositional models may also benefit from the work of Hashimoto et al. (2015).

8 Related Work

Compositional Models for Sentences In order to build a representation (embedding) for a sentence based on its component word embeddings and structural information, recent work on compositional models (stemming from the deep learning community) has designed model structures that mimic the structure of the input. For example, these models could take into account the order of the words (as in Convolutional Neural Networks (CNNs)) (Collobert et al., 2011) or build off of an input tree (as in Recursive Neural Networks (RNNs) or the Semantic Matching Energy Function) (Socher et al., 2013b; Bordes et al., 2012).

While these models work well on sentence-level representations, the nature of their designs also limits them to fixed types of substructures from the annotated sentence, such as chains for CNNs and trees for RNNs. Such models cannot capture arbitrary combinations of linguistic annotations available for a given task, such as word order, dependency tree, and named entities used for relation extraction. Moreover, these approaches ignore the differences in functions between words appearing in different roles. This does not suit more general substructure labeling tasks in NLP, e.g. these models cannot be directly applied to relation extraction since they will output the same result for any pair of entities in a same sentence.

Compositional Models with Annotation Features To tackle the problem of traditional compositional models, Socher et al. (2012) made the RNN model specific to relation extraction tasks by working on the minimal sub-tree which spans the two target entities. However, these specializations

to relation extraction does not generalize easily to other tasks in NLP. There are two ways to achieve such specialization in a more general fashion:

1. *Enhancing Compositional Models with Features.* A recent trend enhances compositional models with annotation features. Such an approach has been shown to significantly improve over pure compositional models. For example, Hermann et al. (2014) and Nguyen and Grishman (2014) gave different weights to words with different syntactic context types or to entity head words with different argument IDs. Zeng et al. (2014) use concatenations of embeddings as features in a CNN model, according to their positions relative to the target entity mentions. Belinkov et al. (2014) enrich embeddings with linguistic features before feeding them forward to a RNN model. Socher et al. (2013a) and Hermann and Blunsom (2013) enhanced RNN models by refining the transformation matrices with phrase types and CCG super tags.

2. *Engineering of Embedding Features.* A different approach to combining traditional linguistic features and embeddings is hand-engineering features with word embeddings and adding them to log-linear models. Such approaches have achieved state-of-the-art results in many tasks including NER, chunking, dependency parsing, semantic role labeling, and relation extraction (Miller et al., 2004; Turian et al., 2010; Koo et al., 2008; Roth and Woodsend, 2014; Sun et al., 2011; Plank and Moschitti, 2013). Roth and Woodsend (2014) considered features similar to ours for semantic role labeling.

However, in prior work both of above approaches are only able to utilize limited information, usually one property for each word. Yet there may be different useful properties of a word which can contribute to the performances of the task. By contrast, our FCM can easily utilize these features without changing the model structures.

In order to better utilize the dependency annotations, recently work built their models according to the dependency paths (Ma et al., 2015; Liu et al., 2015), which share similar motivations to the usage of `On-path` features in our work.

Task-Specific Enhancements for Relation Classification An orthogonal direction of improving compositional models for relation classification is to enhance the models with task-specific information. For example, Hashimoto et al. (2015) trained

task-specific word embeddings, and dos Santos et al. (2015) proposed a ranking-based loss function for relation classification.

9 Conclusion

We have presented FCM, a new compositional model for deriving sentence-level and substructure embeddings from word embeddings. Compared to existing compositional models, FCM can easily handle arbitrary types of input and handle global information for composition, while remaining easy to implement. We have demonstrated that FCM alone attains near state-of-the-art performances on several relation extraction tasks, and in combination with traditional feature based log-linear models it obtains state-of-the-art results.

Our next steps in improving FCM focus on enhancements based on task-specific embeddings or loss functions as in Hashimoto et al. (2015; dos Santos et al. (2015)). Moreover, as the model provides a general idea for representing both sentences and sub-structures in language, it has the potential to contribute useful components to various tasks, such as dependency parsing, SRL and paraphrasing. Also as kindly pointed out by one anonymous reviewer, our FCM can be applied to the TAC-KBP (Ji et al., 2010) tasks, by replacing the training objective to a multi-instance multi-label one (e.g. Surdeanu et al. (2012)). We plan to explore the above applications of FCM in the future.

Acknowledgments

We thank the anonymous reviewers for their comments, and Nicholas Andrews, Francis Ferraro, and Benjamin Van Durme for their input. We thank Kazuma Hashimoto, Cícero Nogueira dos Santos, Bing Xiang and Bowen Zhou for sharing their word embeddings and many helpful discussions. Mo Yu is supported by the China Scholarship Council and by NSFC 61173073.

References

- Yonatan Belinkov, Tao Lei, Regina Barzilay, and Amir Globerson. 2014. Exploring compositional architectures and word vector representations for prepositional phrase attachment. *Transactions of the Association for Computational Linguistics*, 2:561–572.
- Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. 2012. A semantic matching en-

- ergy function for learning with multi-relational data. *Machine Learning*, pages 1–27.
- Massimiliano Ciaramita and Yasemin Altun. 2006. Broad-coverage sense disambiguation and information extraction with a supersense sequence tagger. In *EMNLP2006*, pages 594–602, July.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *JMLR*, 12:2493–2537.
- Cicero dos Santos, Bing Xiang, and Bowen Zhou. 2015. Classifying relations by ranking with convolutional neural networks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 626–634, Beijing, China, July. Association for Computational Linguistics.
- Kazuma Hashimoto, Pontus Stenetorp, Makoto Miwa, and Yoshimasa Tsuruoka. 2015. Task-oriented learning of word embeddings for semantic relation classification. *arXiv preprint arXiv:1503.00095*.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2010. Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. In *Proceedings of SemEval-2 Workshop*.
- Karl Moritz Hermann and Phil Blunsom. 2013. The role of syntax in vector space models of compositional semantics. In *Association for Computational Linguistics*, pages 894–904.
- Karl Moritz Hermann, Dipanjan Das, Jason Weston, and Kuzman Ganchev. 2014. Semantic frame identification with distributed word representations. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1448–1458, Baltimore, Maryland, June. Association for Computational Linguistics.
- Heng Ji, Ralph Grishman, Hoa Trang Dang, Kira Griffitt, and Joe Ellis. 2010. Overview of the tac 2010 knowledge base population track. In *Third Text Analysis Conference (TAC 2010)*.
- Terry Koo, Xavier Carreras, and Michael Collins. 2008. Simple semi-supervised dependency parsing. In *Proceedings of ACL-08: HLT*, pages 595–603, Columbus, Ohio, June. Association for Computational Linguistics.
- Qi Li and Heng Ji. 2014. Incremental joint extraction of entity mentions and relations. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 402–412, Baltimore, Maryland, June. Association for Computational Linguistics.
- Yang Liu, Furu Wei, Sujian Li, Heng Ji, Ming Zhou, and Houfeng WANG. 2015. A dependency-based neural network for relation 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 285–290, Beijing, China, July. Association for Computational Linguistics.
- Mingbo Ma, Liang Huang, Bowen Zhou, and Bing Xiang. 2015. Dependency-based convolutional neural networks for sentence embedding. 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 174–179, Beijing, China, July. Association for Computational Linguistics.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. *arXiv preprint arXiv:1310.4546*.
- Scott Miller, Jethran Guinness, and Alex Zamanian. 2004. Name tagging with word clusters and discriminative training. In Susan Dumais, Daniel Marcu, and Salim Roukos, editors, *HLT-NAACL 2004: Main Proceedings*. Association for Computational Linguistics.
- Alexis Mitchell, Stephanie Strassel, Shudong Huang, and Ramez Zakhary. 2005. Ace 2004 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*.
- Andriy Mnih and Geoffrey Hinton. 2007. Three new graphical models for statistical language modelling. In *Proceedings of the 24th international conference on Machine learning*, pages 641–648. ACM.
- Thien Huu Nguyen and Ralph Grishman. 2014. Employing word representations and regularization for domain adaptation of relation extraction. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 68–74, Baltimore, Maryland, June. Association for Computational Linguistics.
- Thien Huu Nguyen and Ralph Grishman. 2015. Relation extraction: Perspective from convolutional neural networks. In *Proceedings of NAACL Workshop on Vector Space Modeling for NLP*.
- Thien Huu Nguyen, Barbara Plank, and Ralph Grishman. 2015. Semantic representations for do-

- main adaptation: A case study on the tree kernel-based method for relation extraction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 635–644, Beijing, China, July. Association for Computational Linguistics.
- Robert Parker, David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2011. English gigaword fifth edition, june. *Linguistic Data Consortium, LDC2011T07*.
- Barbara Plank and Alessandro Moschitti. 2013. Embedding semantic similarity in tree kernels for domain adaptation of relation extraction. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1498–1507, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Bryan Rink and Sanda Harabagiu. 2010. Utd: Classifying semantic relations by combining lexical and semantic resources. In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 256–259, Uppsala, Sweden, July. Association for Computational Linguistics.
- Michael Roth and Kristian Woodsend. 2014. Composition of word representations improves semantic role labelling. In *EMNLP*.
- Richard Socher, Brody Huval, Christopher D. Manning, and Andrew Y. Ng. 2012. Semantic compositionality through recursive matrix-vector spaces. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1201–1211, Jeju Island, Korea, July. Association for Computational Linguistics.
- Richard Socher, John Bauer, Christopher D Manning, and Andrew Y Ng. 2013a. Parsing with compositional vector grammars. In *In Proceedings of the ACL conference*. Citeseer.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013b. Recursive deep models for semantic compositionality over a sentiment treebank. In *Empirical Methods in Natural Language Processing*, pages 1631–1642.
- Ang Sun, Ralph Grishman, and Satoshi Sekine. 2011. Semi-supervised relation extraction with large-scale word clustering. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 521–529, Portland, Oregon, USA, June. Association for Computational Linguistics.
- Mihai Surdeanu, Julie Tibshirani, Ramesh Nallapati, and Christopher D Manning. 2012. Multi-instance multi-label learning for relation extraction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 455–465. Association for Computational Linguistics.
- Joseph Turian, Lev Ratinov, and Yoshua Bengio. 2010. Word representations: a simple and general method for semi-supervised learning. In *Association for Computational Linguistics*, pages 384–394.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. ACE 2005 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*.
- Mo Yu and Mark Dredze. 2015. Learning composition models for phrase embeddings. *Transactions of the Association for Computational Linguistics*, 3:227–242.
- Mo Yu, Matthew R. Gormley, and Mark Dredze. 2015. Combining word embeddings and feature embeddings for fine-grained relation extraction. In *Proceedings of NAACL*.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2335–2344, Dublin, Ireland, August. Dublin City University and Association for Computational Linguistics.
- GuoDong Zhou, Jian Su, Jie Zhang, and Min Zhang. 2005. Exploring various knowledge in relation extraction. In *Association for Computational Linguistics*, pages 427–434.

Appendix 1: Experiments on ACE 2005 where Gold Entity Types Are Unknown

Experimental Settings: For comparison with prior work (Plank and Moschitti, 2013), we (1) generate relation instances from all pairs of entities within each sentence with three or fewer intervening entity mentions—labeling those pairs with no relation as negative instances, (2) use gold entity spans (but not types) at train and test time, and (3) evaluate on the 7 coarse relation types, ignoring the subtypes. In the training set, 35,990 total relations are annotated of which only 3,658 are non-nil relations. We did not match the number of tokens they reported in the `cts` and `wl` domains. Therefore, in this section we only report the results on the test set of `bc` domain. We will leave experiments on additional domains in future work.

We run the same models as in §7 on this task. Here the FCM does not use entity type features. Plank and Moschitti (2013) also use Brown clusters and word vectors learned by latent-semantic analysis (LSA). In order to make a fair comparison with their method, we also report the FCM result using Brown clusters (prefixes of length 5) of entity heads as entity types. Furthermore, we report non-comparable settings using WordNet super-sense tags of entity heads as types. The WordNet features were also used in their paper but not as substitution of entity types. We use the same toolkit to get the WordNet tags as in §6. The Brown clusters are from (Koo et al., 2008)¹¹.

Results: Table 7 shows the results under the low-resource setting. When no entity types are available, the performance of our FCM only model greatly decreases to 48.15%, which is consistent with our observation in the ablation tests. The baseline model also relies heavily on the entity types. After we remove all the hand-engineering features that contain entity type information, the performance of our baseline model drop to 40.62%, even lower than the reduced FCM only model.

The combination of baseline model and head embeddings (Baseline + HeadOnly) greatly improve the results. This is consistent with the observation in Nguyen and Grishman (2014) that when the gold entity types are unknown, information of the entity heads provided by their embed-

dings will play a more important role. Combination of the baseline and FCM (Baseline + FCM) also achieves improvement but not significantly better than Baseline + HeadOnly. A possible explanation is that FCM becomes less efficient on using context word embeddings when the entity type information is unavailable. In this situation the head embeddings provided by FCM become the dominating contribution to the baseline model, making the model have similar behavior as the Baseline + HeadOnly method.

Finally, we find Brown clusters can help FCM when entity types are unknown. Although the performance is still not significantly better than Baseline + HeadOnly, it outperforms all the results in Plank and Moschitti (2013) as a *single model*, and with the *same source of features*. WordNet super-sense tags further improves FCM, and achieves the best reported results on this low-resource setting. These results are encouraging since it shows FCM may be more useful under the end-to-end setting where predictions of both entity mentions and relation mentions are required in place of predicting relation based on gold tags (Li and Ji, 2014).

Recently Nguyen et al. (2015) proposed a novel way of applying embeddings to tree-kernels. From the results, our best single model achieves comparable result with their best single system, while their combination method is slightly better than ours. This suggests that we may benefit more from combining the usages of multiple word representations; and we will investigate it in future work.

Model	bc		
	P	R	F1
PM'13 (Brown)	54.4	43.4	48.3
PM'13 (LSA)	53.9	45.2	49.2
PM'13 (Combination)	55.3	43.1	48.5
(1) FCM only	53.7	43.7	48.2
(3) Baseline	59.4	30.9	40.6
(4) + HeadOnly	64.9	41.3	50.5
(5) + FCM	65.5	41.5	50.8
(1) FCM only w/ Brown	64.6	40.2	49.6
(1) FCM only w/WordNet	64.0	43.2	51.6
Linear+Emb	46.5	49.3	47.8
Tree-kernel+Emb (Single)	57.6	46.6	51.5
Tree-kernel+Emb (Combination)	58.5	47.3	52.3

Table 7: Comparison of models on ACE 2005 out-of-domain test sets for the low-resource setting, where the gold entity spans are known but entity types are unknown. PM'13 is the results reported in Plank and Moschitti (2013). “Linear+Emb” is the implementation of our method (4) in (Nguyen et al., 2015). The “Tree-kernel+Emb” methods are the enrichments of tree-kernels with embeddings proposed by Nguyen et al. (2015).

¹¹<http://people.csail.mit.edu/maestro/papers/bllip-clusters.gz>