Event Representation Learning Enhanced with External Commonsense Knowledge

Xiao Ding, Kuo Liao, Ting Liu,[*] **Zhongyang Li, Junwen Duan**
Research Center for Social Computing and Information Retrieval
Harbin Institute of Technology, China
{xding, kliao, tliu, zyli, jwduan}@ir.hit.edu.cn

Abstract

Prior work has proposed effective methods to learn event representations that can capture syntactic and semantic information over text corpus, demonstrating their effectiveness for downstream tasks such as script event prediction. On the other hand, events extracted from raw texts lacks of commonsense knowledge, such as the intents and emotions of the event participants, which are useful for distinguishing event pairs when there are only subtle differences in their surface realizations. To address this issue, this paper proposes to leverage external commonsense knowledge about the intent and sentiment of the event. Experiments on three event-related tasks, i.e., event similarity, script event prediction and stock market prediction, show that our model obtains much better event embeddings for the tasks, achieving 78% improvements on hard similarity task, yielding more precise inferences on subsequent events under given contexts, and better accuracies in predicting the volatilities of the stock market[1].

1 Introduction

Events are a kind of important *objective* information of the world. Structuralizing and representing such information as machine-readable knowledge are crucial to artificial intelligence (Li et al., 2018b, 2019). The main idea is to learn distributed representations for structured events (i.e. event embeddings) from text, and use them as the basis to induce textual features for downstream applications, such as script event prediction and stock market prediction.

Parameterized additive models are among the most widely used for learning distributed event representations in prior work (Granroth-Wilding

---

[*]Corresponding author
[1]The code and data are available on https://github.com/MagiaSN/CommonsenseERL_EMNLP_2019.
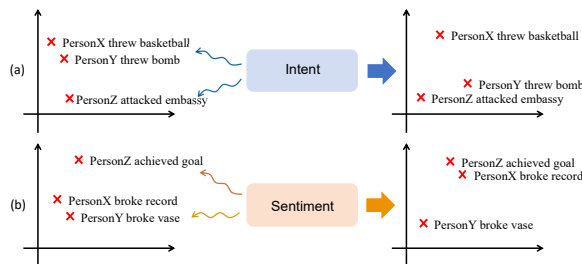

Figure 1: Intent and sentiment enhanced event embeddings can distinguish distinct events even with high lexical overlap, and find similar events even with low lexical overlap.

and Clark, 2016; Modi, 2016), which passes the concatenation or addition of event arguments' word embeddings to a parameterized function. The function maps the summed vectors into an event embedding space. Furthermore, Ding et al. (2015) and Weber et al. (2018) propose using neural tensor networks to perform semantic composition of event arguments, which can better capture the interactions between event arguments.

This line of work only captures shallow event semantics, which is not capable of distinguishing events with subtle differences. On the one hand, the obtained event embeddings cannot capture the relationship between events that are syntactically or semantically similar, if they do not share similar word vectors. For example, as shown in Figure 1 (a), "*PersonX threw bomb*" and "*PersonZ attacked embassy*". On the other hand, two events with similar word embeddings may have similar embeddings despite that they are quite unrelated, for example, as shown in Figure 1 (b), "*PersonX broke record*" and "*PersonY broke vase*". Note that in this paper, similar events generally refer to events with strong semantic relationships rather than just the same events.

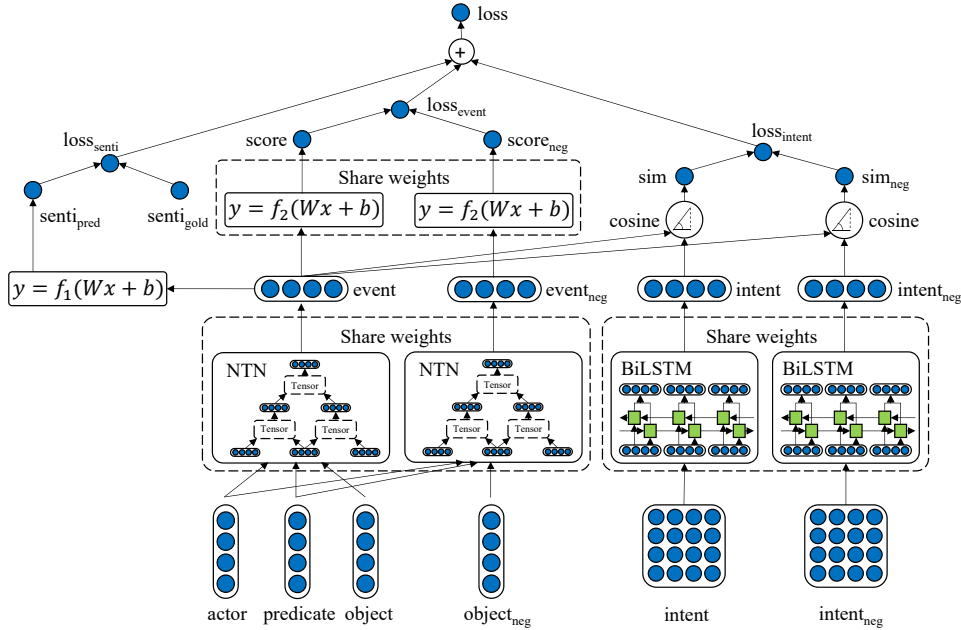One important reason for the problem is the lack

Figure 2: Architecture of the joint embedding model. $event_{neg}$ refers to the corrupted event tuple, which is derived by replacing each word of the event object with a random word in our dictionary. $intent_{neg}$ is the incorrect intent for the given event, which is randomly selected from the annotated dataset.

of the external commonsense knowledge about the mental state of event participants when learning the *objective* event representations. In Figure 1 (a), two event participants "*PersonY*" and "*PersonZ*" may carry out a terrorist attack, and hence, they have the same intent: "*to bloodshed*", which can help representation learning model maps two events into the neighbor vector space. In Figure 1 (b), a change to a single argument leads to a large semantic shift in the event representations, as the change of an argument can result in different emotions of event participants. Who "*broke the record*" is likely to be happy, while, who "*broke a vase*" may be sad. Hence, intent and sentiment can be used to learn more fine-grained semantic features for event embeddings.

Such commonsense knowledge is not explicitly expressed but can be found in a knowledge base such as Event2Mind (Rashkin et al., 2018) and ATOMIC (Sap et al., 2019). Thus, we aim to incorporate the external commonsense knowledge, i.e., **intent** and **sentiment**, into the learning process to generate better event representations. Specifically, we propose a simple and effective model to jointly embed events, intents and emotions into the same vector space. A neural tensor network is used to learn baseline event embeddings, and we define a corresponding loss function to incorporate intent and sentiment information.

Extensive experiments show that incorporating external commonsense knowledge brings promising improvements to event embeddings, achieving 78% and 200% improvements on hard similarity small and big dataset, respectively. With better embeddings, we can achieve superior performances on script event prediction and stock market prediction compared to state-of-the-art baseline methods.

2 Commonsense Knowledge Enhanced Event Representations

The joint embedding framework is shown in Figure 2. We begin by introducing the baseline event embedding learning model, which serves as the basis of the proposed framework. Then, we show how to model intent and sentiment information. Subsequently, we describe the proposed joint model by integrating intent and sentiment into the original objective function to help learn high-quality event representations, and introduce the training details.

2.1 Low-Rank Tensors for Event Embedding

The goal of event embedding is to learn low-dimension dense vector representations for event tuples $E = (A, P, O)$, where P is the action or predicate, A is the actor or subject and O is the object on which the action is performed. Event

4895

Figure 3: Baseline event-embedding model.

Figure 4: An illustration of low-rank neural tensor network for learning event embeddings.

embedding models compound vector representations over its predicate and arguments representations. The challenge is that the composition models should be effective for learning the interactions between the predicate and the argument. Simple additive transformations are incompetent.

We follow Ding et al. (2015) modelling such informative interactions through tensor composition. The architecture of neural tensor network (NTN) for learning event embeddings is shown in Figure 3, where the bilinear tensors are used to explicitly model the relationship between the actor and the action, and that between the object and the action.

The inputs of NTN are the word embeddings of A, P and O, and the outputs are event embeddings. We initialized our word representations using publicly available d-dimensional ($d = 100$) GloVe vectors (Pennington et al., 2014). As most event arguments consist of several words, we represent the actor, action and object as the average of their word embeddings, respectively.

From Figure 3, $S_1 \in \mathbb{R}^d$ is computed by:

$$S_1 = f\left(A^T T_1^{[1:k]} P + W \begin{bmatrix} A \\ P \end{bmatrix} + b\right)$$
$$g(S_1) = g(A, P) = U^T S_1 \tag{1}$$

where $T_1^{[1:k]} \in \mathbb{R}^{d \times d \times k}$ is a tensor, which is a set of k matrices, each with $d \times d$ dimensions. The bilinear tensor product $A^T T_1^{[1:k]} P$ is a vector $r \in \mathbb{R}^k$, where each entry is computed by one slice of the tensor ($r_i = A^T T_1^{[i]} P, i = 1, \cdots, k$). The other parameters are a standard feed-forward neural network, where $W \in \mathbb{R}^{k \times 2d}$ is the weight matrix, $b \in \mathbb{R}^k$ is the bias vector, $U \in \mathbb{R}^k$ is a

hyper-parameter and $f = tanh$ is a standard non-linearity applied element-wise. S_2 and C in Figure 3 are computed in the same way as S_1.

One problem with tensors is *curse of dimensionality*, which limits the wide application of tensors in many areas. It is therefore essential to approximate tensors of higher order in a compressed scheme, for example, a low-rank tensor decomposition. To decrease the number of parameters in standard neural tensor network, we make low-rank approximation that represents each matrix by two low-rank matrices plus diagonal, as illustrated in Figure 4. Formally, the parameter of the i-th slice is $T_{appr}^{[i]} = T^{[i_1]} \times T^{[i_2]} + diag(t^{[i]})$, where $T^{[i_1]} \in \mathbb{R}^{d \times n}$, $T^{[i_2]} \in \mathbb{R}^{n \times d}$, $t^{[i]} \in \mathbb{R}^d$, n is a hyper-parameter, which is used for adjusting the degree of tensor decomposition. The output of neural tensor layer is formalized as follows.

$$S_1 = f\left(A^T [T_{appr}]_1^{[1:k]} P + W \begin{bmatrix} A \\ P \end{bmatrix} + b\right), \tag{2}$$

where $[T_{appr}]_1^{[1:k]}$ is the low-rank tensor that defines multiple low-rank bilinear layers. k is the slice number of neural tensor network which is also equal to the output length of S_1.

We assume that event tuples in the training data should be scored higher than corrupted tuples, in which one of the event arguments is replaced with a random argument. Formally, the corrupted event tuple is $E^r = (A^r, P, O)$, which is derived by replacing each word in A with a random word w^r in our dictionary $\mathcal{D}$ (which contains all the words in the training data) to obtain a corrupted counterpart A^r. We calculate the *margin loss* of the two event tuples as:

$$\mathcal{L}_{\mathcal{E}} = loss(E, E^r) = \max(0, 1 - g(E) + g(E^r)) + \lambda \|\Phi\|_2^2, \tag{3}$$

where $\Phi = (T_1, T_2, T_3, W, b)$ is the set of model parameters. The standard L_2 regularization is used, for which the weight λ is set as 0.0001. The

algorithm goes over the training set for multiple iterations. For each training instance, if the loss $loss(E, E^r) = \max(0, 1 - g(E) + g(E^r))$ is equal to zero, the online training algorithm continues to process the next event tuple. Otherwise, the parameters are updated to minimize the loss using back-propagation (Rumelhart et al., 1985).

2.2 Intent Embedding

Intent embedding refers to encoding the event participants' intents into event vectors, which is mainly used to explain why the actor performed the action. For example, given two events *"PersonX threw basketball"* and *"PersonX threw bomb"*, there are only subtle differences in their surface realizations, however, the intents are totally different. *"PersonX threw basketball"* is just for fun, while *"PersonX threw bomb"* could be a terrorist attack. With the intents, we can easily distinguish these superficial similar events.

One challenge for incorporating intents into event embeddings is that we should have a large-scale labeled dataset, which annotated the event and its actor's intents. Recently, Rashkin et al. (2018) and Sap et al. (2019) released such valuable commonsense knowledge dataset (ATOMIC), which consists of 25,000 event phrases covering a diverse range of daily-life events and situations. For example, given an event *"PersonX drinks coffee in the morning"*, the dataset labels *PersonX*'s likely intent is *"PersonX wants to stay awake"*.

We notice that the intents labeled in ATOMIC is a sentence. Hence, intent embedding is actually a sentence representation learning task. Among various neural networks for encoding sentences, bi-directional LSTMs (BiLSTM) (Hochreiter and Schmidhuber, 1997) have been a dominant method, giving state-of-the-art results in language modelling (Peters et al., 2018) and syntactic parsing (Dozat and Manning, 2016).

We use BiLSTM model to learn intent representations. BiLSTM consists of two LSTM components, which process the input in the forward left-to-right and the backward right-to-left directions, respectively. In each direction, the reading of input words is modelled as a recurrent process with a single hidden state. Given an initial value, the state changes its value recurrently, each time consuming an incoming word.

Take the forward LSTM component for example. Denoting the initial state as $\overrightarrow{h}^0$, which is a

model parameter, it reads the input word representations $\boldsymbol{x}_0, \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$, and the recurrent state transition step for calculating $\overrightarrow{h}^1, \ldots, \overrightarrow{h}^{n+1}$ is defined as Graves and Schmidhuber (2005).

The backward LSTM component follows the same recurrent state transition process as the forward LSTM component. Starting from an initial state $\overleftarrow{h}^{n+1}$, which is a model parameter, it reads the input $\boldsymbol{x}_n, \boldsymbol{x}_{n-1}, \ldots, \boldsymbol{x}_0$, changing its value to $\overleftarrow{h}^n, \overleftarrow{h}^{n-1}, \ldots, \overleftarrow{h}^0$, respectively.

The BiLSTM model uses the concatenated value of $\overrightarrow{h}^t$ and $\overleftarrow{h}^t$ as the hidden vector for w_t:

$$\boldsymbol{h}^t = [\overrightarrow{h}^t; \overleftarrow{h}^t] \tag{4}$$

A single hidden vector representation $\boldsymbol{v}_i$ of the input intent can be obtained by concatenating the last hidden states of the two LSTMs:

$$\boldsymbol{v}_i = [\overrightarrow{h}^{n+1}; \overleftarrow{h}^0] \tag{5}$$

In the training process, we calculate the similarity between a given event vector $\boldsymbol{v}_e$ and its related intent vector $\boldsymbol{v}_i$. For effectively training the model, we devise a ranking type loss function as follows:

$$\mathcal{L}_\mathcal{I} = max(0, 1 - cosine(\boldsymbol{v}_e, \boldsymbol{v}_i) + cosine(\boldsymbol{v}_e, \boldsymbol{v}_i')) \tag{6}$$

where $\boldsymbol{v}_i'$ is the incorrect intent for $\boldsymbol{v}_e$, which is randomly selected from the annotated dataset.

2.3 Sentiment Embedding

Sentiment embedding refers to encoding the event participants' emotions into event vectors, which is mainly used to explain how does the actor feel after the event. For example, given two events *"PersonX broke record"* and *"PersonX broke vase"*, there are only subtle differences in their surface realizations, however, the emotions of *PersonX* are totally different. After *"PersonX broke record"*, *PersonX* may be feel happy, while after *"PersonX broke vase"*, *PersonX* could be feel sad. With the emotions, we can also effectively distinguish these superficial similar events.

We also use ATOMIC (Sap et al., 2019) as the event sentiment labeled dataset. In this dataset, the sentiment of the event is labeled as words. For example, the sentiment of *"PersonX broke vase"* is labeled as "(*sad, be regretful, feel sorry, afraid*)". We use SenticNet (Cambria et al., 2018) to normalize these emotion words ($W = \{w_1, w_2, \ldots, w_n\}$) as the positive (labeled as 1) or

the negative (labeled as -1) sentiment. The sentiment polarity of the event P_e is dependent on the polarity of the labeled emotion words P_W: $P_e = 1$, if $\sum_i P_{w_i} > 0$, or $P_e = -1$, if $\sum_i P_{w_i} < 0$. We use the softmax binary classifier to learn sentiment enhanced event embeddings. The input of the classifier is event embeddings, and the output is its sentiment polarity (positive or negative). The model is trained in a supervised manner by minimizing the cross entropy error of the sentiment classification, whose loss function is given below.

$$\mathcal{L}_{\mathcal{S}} = - \sum_{x_e \in C} \sum_{l \in L} p_l^g(x_e) \cdot log(p_l(x_e)) \qquad (7)$$

where C means all training instances, L is the collection of sentiment categories, x_e means an event vector, $p_l(x_e)$ is the probability of predicting x_e as class l, $p_l^g(x_e)$ indicates whether class l is the correct sentiment category, whose value is 1 or -1.

2.4 Joint Event, Intent and Sentiment Embedding

Given a training event corpus with annotated intents and emotions, our model jointly minimizes a linear combination of the loss functions on events, intents and sentiment:

$$\mathcal{L} = \alpha \mathcal{L}_{\mathcal{E}} + \beta \mathcal{L}_{\mathcal{I}} + \gamma \mathcal{L}_{\mathcal{S}} \qquad (8)$$

where $\alpha, \beta, \gamma \in [0, 1]$ are model parameters to weight the three loss functions.

We use the New York Times Gigaword Corpus (LDC2007T07) for pre-training event embeddings. Event triples are extracted based on the Open Information Extraction technology (Schmitz et al., 2012). We initialize the word embedding layer with 100 dimensional pre-trained GloVe vectors (Pennington et al., 2014), and fine-tune initialized word vectors during our model training. We use Adagrad (Duchi et al., 2011) for optimizing the parameters with initial learning rate 0.001 and batch size 128.

3 Experiments

We compare the performance of intent and sentiment powered event embedding model with state-of-the-art baselines on three tasks: event similarity, script event prediction and stock prediction.

3.1 Baselines

We compare the performance of our approach against a variety of event embedding models de-

veloped in recent years. These models can be categorized into three groups:

- **Averaging Baseline (Avg)** This represents each event as the average of the constituent word vectors using pre-trained GloVe embeddings (Pennington et al., 2014).

- **Compositional Neural Network (Comp. NN)** The event representation in this model is computed by feeding the concatenation of the subject, predicate, and object embedding into a two layer neural network (Modi and Titov, 2013; Modi, 2016; Granroth-Wilding and Clark, 2016).

- **Element-wise Multiplicative Composition (EM Comp.)** This method simply concatenates the element-wise multiplications between the verb and its subject/object.

- **Neural Tensor Network** This line of work use tensors to learn the interactions between the predicate and its subject/object (Ding et al., 2015; Weber et al., 2018). According to the different usage of tensors, we have three baseline methods: **Role Factor Tensor** (Weber et al., 2018) which represents the predicate as a tensor, **Predicate Tensor** (Weber et al., 2018) which uses two tensors learning the interactions between the predicate and its subject, and the predicate and its object, respectively, **NTN** (Ding et al., 2015), which we used as the baseline event embedding model in this paper, and **KGEB** (Ding et al., 2016), which incorporates knowledge graph information in NTN.

3.2 Event Similarity Evaluation

3.2.1 Hard Similarity Task

We first follow Weber et al. (2018) evaluating our proposed approach on the hard similarity task. The goal of this task is that similar events should be close to each other in the same vector space, while dissimilar events should be far away with each other. To this end, Weber et al. (2018) created two types of event pairs, one with events that should be close to each other but have very little lexical overlap (e.g., police catch robber / authorities apprehend suspect), and another with events that should be farther apart but have high overlap (e.g., police catch robber / police catch disease).

Method	Hard Similarity (*Accuracy %*)		Transitive Sentence Similarity (ρ)
	Small Dataset	Big Dataset	
Avg	5.2	13.7	0.67
Comp. NN	33.0	18.9	0.63
EM Comp.	33.9	18.7	0.57
Role Factor Tensor	43.5	20.7	0.64
Predicate Tensor	41.0	25.6	0.63
KGEB	52.6	49.8	0.61
NTN	40.0	37.0	0.60
NTN+Int	65.2	58.1	0.67
NTN+Senti	54.8	52.2	0.61
NTN+Int+Senti	**77.4**	**62.8**	**0.74**

Table 1: Experimental results on hard similarity dataset and transitive sentence similarity dataset. The small dataset (230 event pairs) of hard similarity task from Weber et al. (2018), and the big dataset (2,000 event pairs) is annotated by us. The best results are in bold.

The labeled dataset contains 230 event pairs (115 pairs each of similar and dissimilar types). Three different annotators were asked to give the similarity/dissimilarity rankings, of which only those the annotators agreed upon completely were kept. For each event representation learning method, we obtain the cosine similarity score of the pairs, and report the fraction of cases where the similar pair receives a higher cosine value than the dissimilar pair (we use *Accuracy* $\in [0, 1]$ denoting it). To evaluate the robustness of our approach, we extend this dataset to 1,000 event pairs (similar and dissimilar events each account for 50%), and we will release this dataset to the public.

3.2.2 Transitive Sentence Similarity

Except for the hard similarity task, we also evaluate our approach on the transitive sentence similarity dataset (Kartsaklis and Sadrzadeh, 2014), which contains 108 pairs of transitive sentences: short phrases containing a single subject, object and verb (e.g., agent sell property). It also has another dataset which consists of 200 sentence pairs. In this dataset, the sentences to be compared are constructed using the same subject and object and semantically correlated verbs, such as 'spell' and 'write'; for example, 'pupils write letters' is compared with 'pupils spell letters'. As this dataset is not suitable for our task, we only evaluate our approach and baselines on 108 sentence pairs.

Every pair is annotated by a human with a similarity score from 1 to 7. For example, pairs such as (design, reduce, amount) and (company, cut, cost) are annotated with a high similarity score, while pairs such as (wife, pour, tea) and (worker, join, party) are given low similarity scores. Since each pair has several annotations, we use the average

annotator score as the gold score[2]. To evaluate the cosine similarity given by each model and the annotated similarity score, we use the Spearman's correlation ($\rho \in [-1, 1]$).

3.2.3 Results

Experimental results of hard similarity and transitive sentence similarity are shown in Table 1. We find that:

(1) Simple averaging achieved competitive performance in the task of transitive sentence similarity, while performed very badly in the task of hard similarity. This is mainly because hard similarity dataset is specially created for evaluating the event pairs that should be close to each other but have little lexical overlap and that should be farther apart but have high lexical overlap. Obviously, on such dataset, simply averaging word vectors which is incapable of capturing the semantic interactions between event arguments, cannot achieve a sound performance.

(2) Tensor-based compositional methods (NTN, KGEB, Role Factor Tensor and Predicate Tensor) outperformed parameterized additive models (Comp. NN and EM Comp.), which shows that tensor is capable of learning the semantic composition of event arguments.

(3) Our commonsense knowledge enhanced event representation learning approach outperformed all baseline methods across all datasets (achieving 78% and 200% improvements on hard similarity small and big dataset, respectively, compared to previous SOTA method), which indicates that commonsense knowledge is useful for distinguishing distinct events.

---

[2]To directly compare with baseline methods (Weber et al., 2018), this paper compares with averaged annotator scores, other than comparing with every annotator scores.

Event1	Event 2	oScore	mScore	Event1	Event 2	oScore	mScore
man clears test	he passed exam	-0.08	0.40	man passed car	man passed exam	0.81	0.12
he grind corn	cook chops beans	0.31	0.81	he grind corn	he grind teeth	0.89	0.36
he made meal	chef cooked pasta	0.51	0.85	chef cooked pasta	chef cooked books	0.89	0.45
farmer load truck	person packs car	0.58	0.83	farmer load truck	farmer load gun	0.93	0.55
player scored goal	she carried team	0.19	0.44	she carried bread	she carried team	0.59	0.09

Table 2: Case study of the cosine similarity score changes with incorporating the intent and sentiment. oScore is the original cosine similarity score without intent and sentiment, and mScore is the modified cosine similarity score with intent and sentiment.

Methods	Acc (%)
SGNN	52.45
SGNN+Int	**53.93**
SGNN+Senti	53.57
SGNN+Int+Senti	53.88
SGNN+PairLSTM	52.71
SGNN+EventComp	54.15
SGNN+EventComp+PairLSTM	54.93
SGNN+PairLSTM+Int+Senti	54.14
SGNN+EventComp+Int+Senti	55.08
SGNN+EventComp+PairLSTM+Int+Senti	**56.03**

Table 3: Results of script event prediction on the test set. The improvement is significant at $p < 0.05$. Acc is short for Accuracy.

Figure 5: Experimental results on S&P 500 index prediction. "+Int" means that we encode the intent information into the original event embeddings.

3.2.4 Case Study

To further analyse the effects of intents and emotions on the event representation learning, we present case studies in Table 2, which directly shows the changes of similarity scores before and after incorporating intent and sentiment. For example, the original similarity score of two events "*chef cooked pasta*" and "*chef cooked books*" is very high (0.89) as they have high lexical overlap. However, their intents differ greatly. The intent of "*chef cooked pasta*" is "*to hope his customer enjoying the delicious food*", while the intent of "*chef cooked books*" is "*to falsify their financial statements*". Enhanced with the intents, the similarity score of the above two events dramatically drops to 0.45. For another example, as the event pair "*man clears test*" and "*he passed exam*" share the same sentiment polarity, their similarity score is boosted from -0.08 to 0.40.

3.3 Script Event Prediction

Event is a kind of important real-world knowledge. Learning effective event representations can be benefit for numerous applications. Script event prediction (Chambers and Jurafsky, 2008) is a challenging event-based commonsense reasoning task, which is defined as giving an existing event context, one needs to choose the most reasonable subsequent event from a candidate list.

Following Li et al. (2018a), we evaluate on the standard multiple choice narrative cloze (MCNC) dataset (Granroth-Wilding and Clark, 2016). As SGNN proposed by Li et al. (2018a) achieved state-of-the-art performances for this task, we use the framework of SGNN, and only replace their input event embeddings with our intent and sentiment-enhanced event embeddings.

Wang et al. (2017) and Li et al. (2018a) showed that script event prediction is a challenging problem, and even 1% of accuracy improvement is very difficult. Experimental results shown in Table 3 demonstrate that we can achieve more than 1.5% improvements in single model comparison and more than 1.4% improvements in multi-model integration comparison, just by replacing the input embeddings, which confirms that better event understanding can lead to better inference results. An interesting result is that the event embeddings only incorporated with intents achieved the best result against other baselines. This confirms that capturing people's intents is helpful to infer their next plan. In addition, we notice that the event embeddings only incorporated with sentiment also achieve better performance than SGNN. This is mainly because the emotional consistency does also contribute to predicate the subsequent event.

3.4 Stock Market Prediction

It has been shown that news events influence the trends of stock price movements (Luss and d'Aspremont, 2012). As news events affect human decisions and the volatility of stock prices is influenced by human trading, it is reasonable to say that events can influence the stock market.

In this section, we compare with several event-driven stock market prediction baseline methods: (1) **Word**, Luss and d'Aspremont (2012) use bag-of-words represent news events for stock prediction; (2) **Event**, Ding et al. (2014) represent events by subject-predicate-object triples for stock prediction; (3) **NTN**, Ding et al. (2015) learn contlnues event vectors for stock prediction; (4) **KGEB**, Ding et al. (2016) incorporate knowledge graph into event vectors for stock prediction.

Experimental results are shown in Figure 5. We find that knowledge-driven event embedding is a competitive baseline method, which incorporates world knowledge to improve the performances of event embeddings on the stock prediction. Sentiment is often discussed in predicting stock market, as positive or negative news can affect people's trading decision, which in turn influences the movement of stock market. In this study, we empirically show that event emotions are effective for improving the performance of stock prediction (+2.4%).

4   Related Work

Recent advances in computing power and NLP technology enables more accurate models of events with structures. Using open information extraction to obtain structured events representations, we find that the actor and object of events can be better captured (Ding et al., 2014). For example, a structured representation of the event above can be (Actor = *Microsoft*, Action = *sues*, Object = *Barnes & Noble*). They report improvements on stock market prediction using their structured representation instead of words as features.

One disadvantage of structured representations of events is that they lead to increased sparsity, which potentially limits the predictive power. Ding et al. (2015) propose to address this issue by representing structured events using *event embeddings*, which are dense vectors. The goal of event representation learning is that similar events should be embedded close to each other in the same vector space, and distinct events should be farther from each other.

Previous work investigated compositional models for event embeddings. Granroth-Wilding and Clark (2016) concatenate predicate and argument embeddings and feed them to a neural network to generate an event embedding. Event embeddings are further concatenated and fed through another neural network to predict the coherence between the events. Modi (2016) encodes a set of events in a similar way and use that to incrementally predict the next event – first the argument, then the predicate and then next argument. Pichotta and Mooney (2016) treat event prediction as a sequence to sequence problem and use RNN based models conditioned on event sequences in order to predict the next event. These three works all model narrative chains, that is, event sequences in which a single entity (the protagonist) participates in every event. Hu et al. (2017) also apply an RNN approach, applying a new hierarchical LSTM model in order to predict events by generating descriptive word sequences. This line of work combines the words in these phrases by the passing the concatenation or addition of their word embeddings to a parameterized function that maps the summed vector into event embedding space. The additive nature of these models makes it difficult to model subtle differences in an event's surface form.

To address this issue, Ding et al. (2015), and Weber et al. (2018) propose tensor-based composition models, which combine the subject, predicate and object to produce the final event representation. The models capture multiplicative interactions between these elements and are thus able to make large shifts in event semantics with only small changes to the arguments.

However, previous work mainly focuses on the nature of the event and lose sight of external commonsense knowledge, such as the intent and sentiment of event participants. This paper proposes to encode intent and sentiment into event embeddings, such that we can obtain a kind of more powerful event representations.

5   Conclusion

Understanding events requires effective representations that contain commonsense knowledge. High-quality event representations are valuable for many NLP downstream applications. This paper proposed a simple and effective framework to incorporate commonsense knowledge into the learn-

ing process of event embeddings. Experimental results on event similarity, script event prediction and stock prediction showed that commonsense knowledge enhanced event embeddings can improve the quality of event representations and benefit the downstream applications.

References

Erik Cambria, Soujanya Poria, Devamanyu Hazarika, and Kenneth Kwok. 2018. Senticnet 5: discovering conceptual primitives for sentiment analysis by means of context embeddings. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018.*

Nathanael Chambers and Dan Jurafsky. 2008. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797. Association for Computational Linguistics.

Xiao Ding, Yue Zhang, Ting Liu, and Junwen Duan. 2014. Using structured events to predict stock price movement: An empirical investigation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1415–1425, Doha, Qatar. Association for Computational Linguistics.

Xiao Ding, Yue Zhang, Ting Liu, and Junwen Duan. 2015. Deep learning for event-driven stock prediction. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 2327–2333.

Xiao Ding, Yue Zhang, Ting Liu, and Junwen Duan. 2016. Knowledge-driven event embedding for stock prediction. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2133–2142.

Timothy Dozat and Christopher D Manning. 2016. Deep biaffine attention for neural dependency parsing. *arXiv preprint arXiv:1611.01734.*

John C Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159.

Mark Granroth-Wilding and Stephen Christopher Clark. 2016. What happens next? event prediction using a compositional neural network model. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 2727–2733.

Alex Graves and Jürgen Schmidhuber. 2005. Framewise phoneme classification with bidirectional lstm and other neural network architectures. *Neural Networks*, 18(5-6):602–610.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

Linmei Hu, Juanzi Li, Liqiang Nie, Xiao-Li Li, and Chao Shao. 2017. What happens next? future subevent prediction using contextual hierarchical lstm. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3450–3456.

Dimitri Kartsaklis and Mehrnoosh Sadrzadeh. 2014. A study of entanglement in a categorical framework of natural language. *Electronic Proceedings in Theoretical Computer Science*, 172.

Zhongyang Li, Xiao Ding, and Ting Liu. 2018a. Constructing narrative event evolutionary graph for script event prediction. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.*, pages 4201–4207.

Zhongyang Li, Xiao Ding, and Ting Liu. 2018b. Generating reasonable and diversified story ending using sequence to sequence model with adversarial training. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1033–1043, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Zhongyang Li, Xiao Ding, and Ting Liu. 2019. Story ending prediction by transferable bert. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 1800–1806.

Ronny Luss and Alexandre d'Aspremont. 2012. Predicting abnormal returns from news using text classification. *Quantitative Finance*, (doi:10.1080/14697688.2012.672762):1–14.

Ashutosh Modi. 2016. Event embeddings for semantic script modeling. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 75–83.

Ashutosh Modi and Ivan Titov. 2013. Learning semantic script knowledge with event embeddings. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*.

Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.

Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 2227–2237.

Karl Pichotta and Raymond J Mooney. 2016. Learning statistical scripts with lstm recurrent neural networks. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 2800–2806.

Hannah Rashkin, Maarten Sap, Emily Allaway, Noah A. Smith, and Yejin Choi. 2018. Event2mind: Commonsense inference on events, intents, and reactions. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 463–473. Association for Computational Linguistics.

David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. 1985. Learning internal representations by error propagation. Technical report, DTIC Document.

Maarten Sap, Ronan LeBras, Emily Allaway, Chandra Bhagavatula, Nicholas Lourie, Hannah Rashkin, Brendan Roof, Noah A Smith, and Yejin Choi. 2019. Atomic: An atlas of machine commonsense for if-then reasoning. pages 3027–3035.

Michael Schmitz, Robert Bart, Stephen Soderland, Oren Etzioni, et al. 2012. Open language learning for information extraction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 523–534. Association for Computational Linguistics.

Zhongqing Wang, Yue Zhang, and Ching-Yun Chang. 2017. Integrating order information and event relation for script event prediction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 57–67. Association for Computational Linguistics.

Noah Weber, Niranjan Balasubramanian, and Nathanael Chambers. 2018. Event representations with tensor-based compositions. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018.*