

Discourse Marker Augmented Network with Reinforcement Learning for Natural Language Inference

Boyuan Pan[†], Yazheng Yang[‡], Zhou Zhao[‡], Yueting Zhuang[‡], Deng Cai^{†*}, Xiaofei He^{*†}

[†]State Key Lab of CAD&CG, Zhejiang University, Hangzhou, China

[‡]College of Computer Science, Zhejiang University, Hangzhou, China

[‡]Alibaba-Zhejiang University Joint Institute of Frontier Technologies

^{*}Fabu Inc., Hangzhou, China

{panby, yazheng_yang, zhaozhou, yzhuang}@zju.edu.cn

{dengcai, xiaofeihe}@cad.zju.edu.com

Abstract

Natural Language Inference (NLI), also known as Recognizing Textual Entailment (RTE), is one of the most important problems in natural language processing. It requires to infer the logical relationship between two given sentences. While current approaches mostly focus on the interaction architectures of the sentences, in this paper, we propose to transfer knowledge from some important discourse markers to augment the quality of the NLI model. We observe that people usually use some discourse markers such as “so” or “but” to represent the logical relationship between two sentences. These words potentially have deep connections with the meanings of the sentences, thus can be utilized to help improve the representations of them. Moreover, we use reinforcement learning to optimize a new objective function with a reward defined by the property of the NLI datasets to make full use of the labels information. Experiments show that our method achieves the state-of-the-art performance on several large-scale datasets.

1 Introduction

In this paper, we focus on the task of Natural Language Inference (NLI), which is known as a significant yet challenging task for natural language understanding. In this task, we are given two sentences which are respectively called **premise** and **hypothesis**. The goal is to determine whether the logical relationship between them is *entailment*, *neutral*, or *contradiction*.

Recently, performance on NLI(Chen et al., 2017b; Gong et al., 2018; Chen et al., 2017c)

*corresponding author

Premise: A soccer game with multiple males playing.

Hypothesis: Some men are playing a sport.

Label: *Entailment*

Premise: An older and younger man smiling.

Hypothesis: Two men are smiling and laughing at the cats playing on the floor.

Label: *Neutral*

Premise: A black race car starts up in front of a crowd of people

Hypothesis: A man is driving down a lonely road.

Label: *Contradiction*

Table 1: Three examples in SNLI dataset.

has been significantly boosted since the release of some high quality large-scale benchmark datasets such as SNLI(Bowman et al., 2015) and MultiNLI(Williams et al., 2017). Table 1 shows some examples in SNLI. Most state-of-the-art works focus on the interaction architectures between the premise and the hypothesis, while they rarely concerned the discourse relations of the sentences, which is a core issue in natural language understanding.

People usually use some certain set of words to express the discourse relation between two sentences¹. These words, such as “but” or “and”, are denoted as *discourse markers*. These discourse markers have deep connections with the intrinsic relations of two sentences and intuitively correspond to the intent of NLI, such as “but” to “contradiction”, “so” to “entailment”, *etc.*

Very few NLI works utilize this information revealed by discourse markers. Nie et al. (2017) proposed to use discourse markers to help rep-

¹Here sentences mean either the whole sentences or the main clauses of a compound sentence.

represent the meanings of the sentences. However, they represent each sentence by a single vector and directly concatenate them to predict the answer, which is too simple and not ideal for the large-scale datasets.

In this paper, we propose a *Discourse Marker Augmented Network* for natural language inference, where we transfer the knowledge from the existing supervised task: Discourse Marker Prediction (DMP)(Nie et al., 2017), to an integrated NLI model. We first propose a sentence encoder model that learns the representations of the sentences from the DMP task and then inject the encoder to the NLI network. Moreover, because our NLI datasets are manually annotated, each example from the datasets might get several different labels from the annotators although they will finally come to a consensus and also provide a certain label. In consideration of that different confidence level of the final labels should be discriminated, we employ reinforcement learning with a reward defined by the uniformity extent of the original labels to train the model. The contributions of this paper can be summarized as follows.

- Unlike previous studies, we solve the task of the natural language inference via transferring knowledge from another supervised task. We propose the Discourse Marker Augmented Network to combine the learned encoder of the sentences with the integrated NLI model.
- According to the property of the datasets, we incorporate reinforcement learning to optimize a new objective function to make full use of the labels' information.
- We conduct extensive experiments on two large-scale datasets to show that our method achieves better performance than other state-of-the-art solutions to the problem.

2 Task Description

2.1 Natural Language Inference (NLI)

In the natural language inference tasks, we are given a pair of sentences (P, H) , which respectively means the premise and hypothesis. Our goal is to judge whether their logical relationship between their meanings by picking a label from a small set: *entailment* (The hypothesis is definitely a true description of the premise), *neutral*

(The hypothesis might be a true description of the premise), and *contradiction* (The hypothesis is definitely a false description of the premise).

2.2 Discourse Marker Prediction (DMP)

For DMP, we are given a pair of sentences (S_1, S_2) , which is originally the first half and second half of a complete sentence. The model must predict which discourse marker was used by the author to link the two ideas from a set of candidates.

3 Sentence Encoder Model

Following (Nie et al., 2017; Kiros et al., 2015), we use BookCorpus(Zhu et al., 2015) as our training data for discourse marker prediction, which is a dataset of text from unpublished novels, and it is large enough to avoid bias towards any particular domain or application. After preprocessing, we obtain a dataset with the form (S_1, S_2, m) , which means the first half sentence, the last half sentence, and the discourse marker that connected them in the original text. Our goal is to predict the m given S_1 and S_2 .

We first use *Glove*(Pennington et al., 2014) to transform $\{S_t\}_{t=1}^2$ into vectors word by word and subsequently input them to a bi-directional LSTM:

$$\begin{aligned} \vec{\mathbf{h}}_t^i &= \overrightarrow{\text{LSTM}}(\text{Glove}(S_t^i)), i = 1, \dots, |S_t| \\ \overleftarrow{\mathbf{h}}_t^i &= \overleftarrow{\text{LSTM}}(\text{Glove}(S_t^i)), i = |S_t|, \dots, 1 \end{aligned} \quad (1)$$

where $\text{Glove}(w)$ is the embedding vector of the word w from the *Glove* lookup table, $|S_t|$ is the length of the sentence S_t . We apply max pooling on the concatenation of the hidden states from both directions, which provides regularization and shorter back-propagation paths(Collobert and Weston, 2008), to extract the features of the whole sequences of vectors:

$$\begin{aligned} \vec{\mathbf{r}}_t &= \text{Max}_{dim}([\vec{\mathbf{h}}_t^1, \vec{\mathbf{h}}_t^2, \dots, \vec{\mathbf{h}}_t^{|S_t|}]) \\ \overleftarrow{\mathbf{r}}_t &= \text{Max}_{dim}([\overleftarrow{\mathbf{h}}_t^1, \overleftarrow{\mathbf{h}}_t^2, \dots, \overleftarrow{\mathbf{h}}_t^{|S_t|}]) \end{aligned} \quad (2)$$

where Max_{dim} means that the max pooling is performed across each dimension of the concatenated vectors, $[\cdot]$ denotes concatenation. Moreover, we combine the last hidden state from both directions and the results of max pooling to represent our sentences:

$$\mathbf{r}_t = [\vec{\mathbf{r}}_t; \overleftarrow{\mathbf{r}}_t; \vec{\mathbf{h}}_t^{|S_t|}; \overleftarrow{\mathbf{h}}_t^1] \quad (3)$$

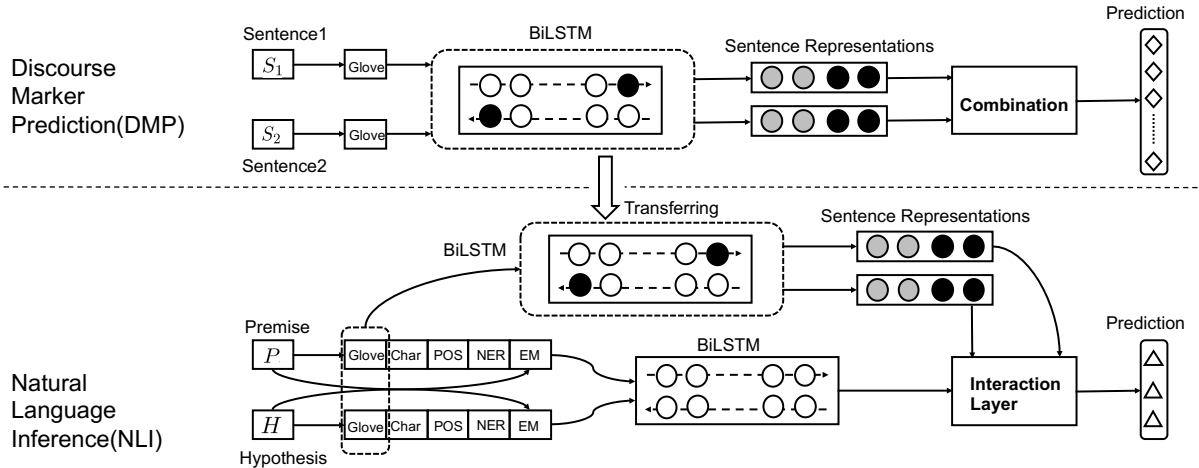


Figure 1: Overview of our Discourse Marker Augmented Network, comprising the part of Discourse Marker Prediction (upper) for pre-training and Natural Language Inference (bottom) to which the learned knowledge will be transferred.

where \mathbf{r}_t is the representation vector of the sentence S_t . To predict the discourse marker between S_1 and S_2 , we combine the representations of them with some linear operation:

$$\mathbf{r} = [\mathbf{r}_1; \mathbf{r}_2; \mathbf{r}_1 + \mathbf{r}_2; \mathbf{r}_1 \odot \mathbf{r}_2] \quad (4)$$

where \odot is elementwise product. Finally we project \mathbf{r} to a vector of label size (the total number of discourse markers in the dataset) and use softmax function to normalize the probability distribution.

4 Discourse Marker Augmented Network

As presented in Figure 1, we show how our Discourse Marker Augmented Network incorporates the learned encoder into the NLI model.

4.1 Encoding Layer

We denote the premise as P and the hypothesis as H . To encode the words, we use the concatenation of following parts:

Word Embedding: Similar to the previous section, we map each word to a vector space by using pre-trained word vectors *GloVe*.

Character Embedding: We apply Convolutional Neural Networks (CNN) over the characters of each word. This approach is proved to be helpful in handling out-of-vocab (OOV) words (Yang et al., 2017).

POS and NER tags: We use the part-of-speech (POS) tags and named-entity recognition (NER)

tags to get syntactic information and entity label of the words. Following (Pan et al., 2017b), we apply the skip-gram model (Mikolov et al., 2013) to train two new lookup tables of POS tags and NER tags respectively. Each word can get its own POS embedding and NER embedding by these lookup tables. This approach represents much better geometrical features than common used one-hot vectors.

Exact Match: Inspired by the machine comprehension tasks (Chen et al., 2017a), we want to know whether every word in P is in H (and H in P). We use three binary features to indicate whether the word can be exactly matched to any question word, which respectively means original form, lowercase and lemma form.

For encoding, we pass all sequences of vectors into a bi-directional LSTM and obtain:

$$\begin{aligned} \mathbf{p}_i &= \text{BiLSTM}(f_{rep}(P_i), \mathbf{p}_{i-1}), i = 1, \dots, n \\ \mathbf{u}_j &= \text{BiLSTM}(f_{rep}(H_j), \mathbf{u}_{j-1}), j = 1, \dots, m \end{aligned} \quad (5)$$

where $f_{rep}(\mathbf{x}) = [\text{Glove}(\mathbf{x}); \text{Char}(\mathbf{x}); \text{POS}(\mathbf{x}); \text{NER}(\mathbf{x}); \text{EM}(\mathbf{x})]$ is the concatenation of the embedding vectors and the feature vectors of the word \mathbf{x} , $n = |P|$, $m = |H|$.

4.2 Interaction Layer

In this section, we feed the results of the encoding layer and the learned sentence encoder into the attention mechanism, which is responsible for linking and fusing information from the premise and the hypothesis words.

We first obtain a similarity matrix $\mathbf{A} \in R^{n \times m}$ between the premise and hypothesis by

$$\mathbf{A}_{ij} = \mathbf{v}_1^\top [\mathbf{p}_i; \mathbf{u}_j; \mathbf{p}_i \circ \mathbf{u}_j; \mathbf{r}_p; \mathbf{r}_h] \quad (6)$$

where \mathbf{v}_1 is the trainable parameter, \mathbf{r}_p and \mathbf{r}_h are sentences representations from the equation (3) learned in the Section 3, which denote the premise and hypothesis respectively. In addition to previous popular similarity matrix, we incorporate the relevance of each word of $P(H)$ to the whole sentence of $H(P)$. Now we use \mathbf{A} to obtain the attentions and the attended vectors in both directions.

To signify the attention of the i -th word of P to every word of H , we use the weighted sum of \mathbf{u}_j by \mathbf{A}_i :

$$\tilde{\mathbf{u}}_i = \sum_j \mathbf{A}_{ij} \cdot \mathbf{u}_j \quad (7)$$

where $\tilde{\mathbf{u}}_i$ is the attention vector of the i -th word of P for the entire H . In the same way, the $\tilde{\mathbf{p}}_j$ is obtained via:

$$\tilde{\mathbf{p}}_j = \sum_i \mathbf{A}_{ij} \cdot \mathbf{p}_i \quad (8)$$

To model the local inference between aligned word pairs, we integrate the attention vectors with the representation vectors via:

$$\begin{aligned} \hat{\mathbf{p}}_i &= f([\mathbf{p}_i; \tilde{\mathbf{u}}_i; \mathbf{p}_i - \tilde{\mathbf{u}}_i; \mathbf{p}_i \odot \tilde{\mathbf{u}}_i]) \\ \hat{\mathbf{u}}_j &= f([\mathbf{u}_j; \tilde{\mathbf{p}}_j; \mathbf{u}_j - \tilde{\mathbf{p}}_j; \mathbf{u}_j \odot \tilde{\mathbf{p}}_j]) \end{aligned} \quad (9)$$

where f is a 1-layer feed-forward neural network with the ReLU activation function, $\hat{\mathbf{p}}_i$ and $\hat{\mathbf{u}}_j$ are local inference vectors. Inspired by (Seo et al., 2016) and (Chen et al., 2017b), we use a modeling layer to capture the interaction between the premise and the hypothesis. Specifically, we use bi-directional LSTMs as building blocks:

$$\begin{aligned} \mathbf{p}_i^M &= \text{BiLSTM}(\hat{\mathbf{p}}_i, \mathbf{p}_{i-1}^M) \\ \mathbf{u}_j^M &= \text{BiLSTM}(\hat{\mathbf{u}}_j, \mathbf{u}_{j-1}^M) \end{aligned} \quad (10)$$

Here, \mathbf{p}_i^M and \mathbf{u}_j^M are the modeling vectors which contain the crucial information and relationship among the sentences.

We compute the representation of the whole sentence by the weighted average of each word:

$$\begin{aligned} \mathbf{p}^M &= \sum_i \frac{\exp(\mathbf{v}_2^\top \mathbf{p}_i^M)}{\sum_{i'} \exp(\mathbf{v}_2^\top \mathbf{p}_{i'}^M)} \mathbf{p}_i^M \\ \mathbf{u}^M &= \sum_j \frac{\exp(\mathbf{v}_3^\top \mathbf{u}_j^M)}{\sum_{j'} \exp(\mathbf{v}_3^\top \mathbf{u}_{j'}^M)} \mathbf{u}_j^M \end{aligned} \quad (11)$$

Label Number	SNLI		MultiNLI	
	Correct	Total	Correct	Total
1	510711	510711	392702	392702
2	0	0	0	0
3	8748	0	3045	0
4	16395	2199	4859	0
5	33179	56123	11743	19647

Table 2: Statistics of the labels of SNLI and MultiNLI. **Total** means the number of examples whose number of annotators is in the left column. **Correct** means the number of examples whose number of correct labels from the annotators is in the left column.

where $\mathbf{v}_2, \mathbf{v}_3$ are trainable vectors. We don't share these parameter vectors in this seemingly parallel structure because there is some subtle difference between the premise and hypothesis, which will be discussed later in Section 5.

4.3 Output Layer

The NLI task requires the model to predict the logical relation from the given set: *entailment*, *neutral* or *contradiction*. We obtain the probability distribution by a linear function with softmax function:

$$\mathbf{d} = \text{softmax}(\mathbf{W}[\mathbf{p}^M; \mathbf{u}^M; \mathbf{p}^M \odot \mathbf{u}^M; \mathbf{r}_p \odot \mathbf{r}_h]) \quad (12)$$

where \mathbf{W} is a trainable parameter. We combine the representations of the sentences computed above with the representations learned from DMP to obtain the final prediction.

4.4 Training

As shown in Table 2, many examples from our datasets are labeled by several people, and the choices of the annotators are not always consistent. For instance, when the label number is 3 in SNLI, "total=0" means that no examples have 3 annotators (maybe more or less); "correct=8748" means that there are 8748 examples whose number of correct labels is 3 (the number of annotators maybe 4 or 5, but some provided wrong labels). Although all the labels for each example will be unified to a final (correct) label, diversity of the labels for a single example indicates the low confidence of the result, which is not ideal to only use the final label to optimize the model.

We propose a new objective function that combines both the log probabilities of the ground-truth label and a reward defined by the property of the datasets for the reinforcement learning. The most widely used objective function for the natural language inference is to minimize the negative log cross-entropy loss:

$$J_{CE}(\Theta) = -\frac{1}{N} \sum_k^N \log(\mathbf{d}_l^k) \quad (13)$$

where Θ are all the parameters to optimize, N is the number of examples in the dataset, \mathbf{d}_l is the probability of the ground-truth label l .

However, directly using the final label to train the model might be difficult in some situations, where the example is confusing and the labels from the annotators are different. For instance, consider an example from the SNLI dataset:

- P : “A smiling costumed woman is holding an umbrella.”
- H : “A happy woman in a fairy costume holds an umbrella.”

The final label is *neutral*, but the original labels from the five annotators are *neural*, *neural*, *entailment*, *contradiction*, *neural*, in which case the relation between “smiling” and “happy” might be under different comprehension. The final label’s confidence of this example is obviously lower than an example that all of its labels are the same. To simulate the thought of human being more closely, in this paper, we tackle this problem by using the REINFORCE algorithm(Williams, 1992) to minimize the negative expected reward, which is defined as:

$$J_{RL}(\Theta) = -\mathbb{E}_{l \sim \pi(l|P,H)}[R(l, \{l^*\})] \quad (14)$$

where $\pi(l|P, H)$ is the previous action policy that predicts the label given P and H , $\{l^*\}$ is the set of annotated labels, and

$$R(l, \{l^*\}) = \frac{\text{number of } l \text{ in } \{l^*\}}{|\{l^*\}|} \quad (15)$$

is the reward function defined to measure the distance to all the ideas of the annotators.

To avoid of overwriting its earlier results and further stabilize training, we use a linear function to integrate the above two objective functions:

$$J(\Theta) = \lambda J_{CE}(\Theta) + (1 - \lambda) J_{RL}(\Theta) \quad (16)$$

where λ is a tunable hyperparameter.

Discourse Marker	Percentage(%)
<i>but</i>	57.12
<i>because</i>	9.41
<i>if</i>	29.78
<i>when</i>	25.32
<i>so</i>	31.01
<i>although</i>	1.76
<i>before</i>	15.52
<i>still</i>	11.29

Table 3: Statistics of discourse markers in our dataset from BookCorpus.

5 Experiments

5.1 Datasets

BookCorpus: We use the dataset from BookCorpus(Zhu et al., 2015) to pre-train our sentence encoder model. We preprocessed and collected discourse markers from BookCorpus as (Nie et al., 2017). We finally curated a dataset of 6527128 pairs of sentences for 8 discourse markers, whose statistics are shown in Table 3.

SNLI: Stanford Natural Language Inference(Bowman et al., 2015) is a collection of more than 570k human annotated sentence pairs labeled for entailment, contradiction, and semantic independence. SNLI is two orders of magnitude larger than all other resources of its type. The premise data is extracted from the captions of the Flickr30k corpus(Young et al., 2014), the hypothesis data and the labels are manually annotated. The original SNLI corpus contains also the other category, which includes the sentence pairs lacking consensus among multiple human annotators. We remove this category and use the same split as in (Bowman et al., 2015) and other previous work.

MultiNLI: Multi-Genre Natural Language Inference(Williams et al., 2017) is another large-scale corpus for the task of NLI. MultiNLI has 433k sentences pairs and is in the same format as SNLI, but it includes a more diverse range of text, as well as an auxiliary test set for cross-genre transfer evaluation. Half of these selected genres appear in training set while the rest are not, creating in-domain (matched) and cross-domain (mismatched) development/test sets.

Method	SNLI	MultiNLI	
		Matched	Mismatched
300D LSTM encoders(Bowman et al., 2016)	80.6	–	–
300D Tree-based CNN encoders(Mou et al., 2016)	82.1	–	–
4096D BiLSTM with max-pooling(Conneau et al., 2017)	84.5	–	–
600D Gumbel TreeLSTM encoders(Choi et al., 2017)	86.0	–	–
600D Residual stacked encoders(Nie and Bansal, 2017)	86.0	74.6	73.6
Gated-Att BiLSTM(Chen et al., 2017d)	–	73.2	73.6
100D LSTMs with attention(Rocktäschel et al., 2016)	83.5	–	–
300D re-read LSTM(Sha et al., 2016)	87.5	–	–
DIIN(Gong et al., 2018)	88.0	78.8	77.8
Biattentive Classification Network(McCann et al., 2017)	88.1	–	–
300D CAFE(Tay et al., 2017)	88.5	78.7	77.9
KIM(Chen et al., 2017b)	88.6	–	–
600D ESIM + 300D Syntactic TreeLSTM(Chen et al., 2017c)	88.6	–	–
DMAN	88.8	78.9	78.2
BiMPM(Ensemble)(Wang et al., 2017)	88.8	–	–
DIIN(Ensemble)(Gong et al., 2018)	88.9	80.0	78.7
KIM(Ensemble)(Chen et al., 2017b)	89.1	–	–
300D CAFE(Ensemble)(Tay et al., 2017)	89.3	80.2	79.0
DMAN(Ensemble)	89.6	80.3	79.4

Table 4: Performance on the SNLI dataset and the MultiNLI dataset. In the top part, we show sentence encoding-based models; In the medium part, we present the performance of integrated neural network models; In the bottom part, we show the results of ensemble models.

5.2 Implementation Details

We use the Stanford CoreNLP toolkit(Manning et al., 2014) to tokenize the words and generate POS and NER tags. The word embeddings are initialized by 300d *Glove*(Pennington et al., 2014), the dimensions of POS and NER embeddings are 30 and 10. The dataset we use to train the embeddings of POS tags and NER tags are the training set given by SNLI. We apply Tensorflow r1.3 as our neural network framework. We set the hidden size as 300 for all the LSTM layers and apply dropout(Srivastava et al., 2014) between layers with an initial ratio of 0.9, the decay rate as 0.97 for every 5000 step. We use the AdaDelta for optimization as described in (Zeiler, 2012) with ρ as 0.95 and ϵ as $1e-8$. We set our batch size as 36 and the initial learning rate as 0.6. The parameter λ in the objective function is set to be 0.2. For DMP task, we use stochastic gradient descent with initial learning rate as 0.1, and we anneal by half each time the validation accuracy is lower than the previous epoch. The number of epochs is set to be 10, and the feedforward dropout rate is 0.2. The

learned encoder in subsequent NLI task is trainable.

5.3 Results

In table 4, we compare our model to other competitive published models on SNLI and MultiNLI. As we can see, our method Discourse Marker Augmented Network (DMAN) clearly outperforms all the baselines and achieves the state-of-the-art results on both datasets.

The methods in the top part of the table are sentence encoding based models. Bowman et al. (2016) proposed a simple baseline that uses LSTM to encode the whole sentences and feed them into a MLP classifier to predict the final inference relationship, they achieve an accuracy of 80.6% on SNLI. Nie and Bansal (2017) test their model on both SNLI and MultiNLI, and achieves competitive results.

In the medium part, we show the results of other neural network models. Obviously, the performance of most of the integrated methods are better than the sentence encoding based models above. Both DIIN(Gong et al., 2018) and

Ablation Model	Accuracy
Only Sentence Encoder Model	83.37
No Sentence Encoder Model	87.24
No Char Embedding	87.95
No POS Embedding	88.76
No NER Embedding	88.71
No Exact Match	88.26
No REINFORCE	88.41
DMAN	88.83

Table 5: Ablations on the SNLI development dataset.

CAFE(Tay et al., 2017) exceed other methods by more than 4% on MultiNLI dataset. However, our DMAN achieves 88.8% on SNLI, 78.9% on matched MultiNLI and 78.2% on mismatched MultiNLI, which are all best results among the baselines.

We present the ensemble results on both datasets in the bottom part of the table 4. We build an ensemble model which consists of 10 single models with the same architecture but initialized with different parameters. The performance of our model achieves 89.6% on SNLI, 80.3% on matched MultiNLI and 79.4% on mismatched MultiNLI, which are all state-of-the-art results.

5.4 Ablation Analysis

As shown in Table 5, we conduct an ablation experiment on SNLI development dataset to evaluate the individual contribution of each component of our model. Firstly we only use the results of the sentence encoder model to predict the answer, in other words, we represent each sentence by a single vector and use dot product with a linear function to do the classification. The result is obviously not satisfactory, which indicates that only using sentence embedding from discourse markers to predict the answer is not ideal in large-scale datasets. We then remove the sentence encoder model, which means we don't use the knowledge transferred from the DMP task and thus the representations r_p and r_h are set to be zero vectors in the equation (6) and the equation (12). We observe that the performance drops significantly to 87.24%, which is nearly 1.5% to our DMAN model, which indicates that the discourse markers have deep connections with the logical relations between two sentences they links. When

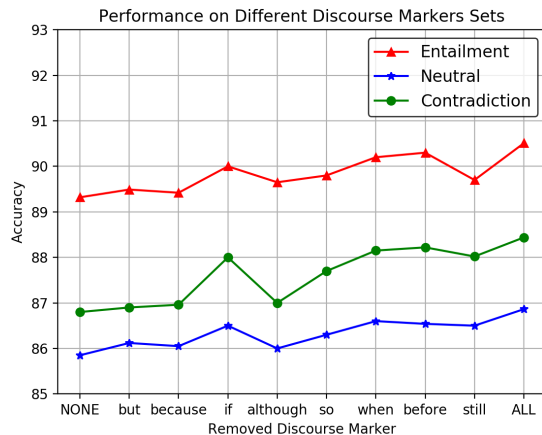


Figure 2: Performance when the sentence encoder is pretrained on different discourse markers sets. “NONE” means the model doesn't use any discourse markers; “ALL” means the model use all the discourse markers.

we remove the character-level embedding and the POS and NER features, the performance drops a lot. We conjecture that those feature tags help the model represent the words as a whole while the char-level embedding can better handle the out-of-vocab (OOV) or rare words. The exact match feature also demonstrates its effectiveness in the ablation result. Finally, we ablate the reinforcement learning part, in other words, we only use the original loss function to optimize the model (set $\lambda = 1$). The result drops about 0.5%, which proves that it is helpful to utilize all the information from the annotators.

5.5 Semantic Analysis

In Figure 2, we show the performance on the three relation labels when the model is pre-trained on different discourse markers sets. In other words, we removed discourse marker from the original set each time and use the rest 7 discourse markers to pre-train the sentence encoder in the DMP task and then train the DMAN. As we can see, there is a sharp decline of accuracy when removing “but”, “because” and “although”. We can intuitively speculate that “but” and “although” have direct connections with the contradiction label (which drops most significantly) while “because” has some links with the entailment label. We observe that some discourse markers such as “if” or “before” contribute much less than other words which have strong logical hints, although they

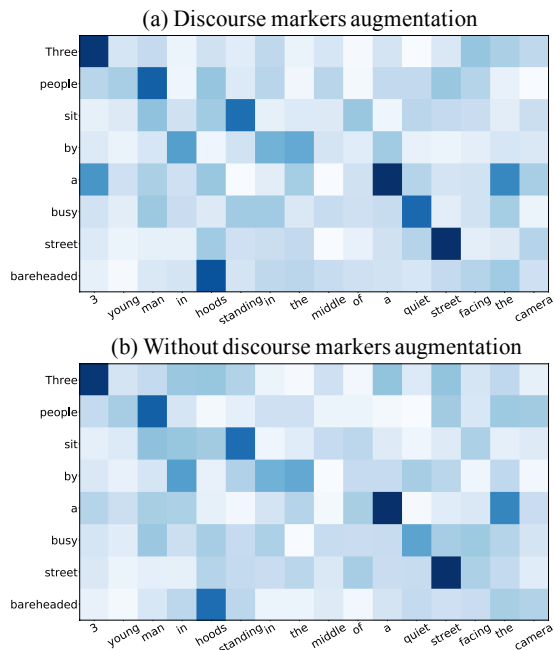


Figure 3: Comparison of the visualized similarity relations.

actually improve the performance of the model. Compared to the other two categories, the “contradiction” label examples seem to benefit the most from the pre-trained sentence encoder.

5.6 Visualization

In Figure 3, we also provide a visualized analysis of the hidden representation from similarity matrix A (computed in the equation (6)) in the situations that whether we use the discourse markers or not. We pick a sentence pair whose premise is “3 young man in hoods standing in the middle of a quiet street facing the camera.” and hypothesis is “Three people sit by a busy street bareheaded.” We observe that the values are highly correlated among the synonyms like “people” with “man”, “three” with “3” in both situations. However, words that might have contradictory meanings like “hoods” with “bareheaded”, “quiet” with “busy” perform worse without the discourse markers augmentation, which conforms to the conclusion that the “contradiction” label examples benefit a lot which is observed in the Section 5.5.

6 Related Work

6.1 Discourse Marker Applications

This work is inspired most directly by the DisSent model and Discourse Prediction Task of Nie et al. (2017), which introduce the use of the discourse

markers information for the pretraining of sentence encoders. They follow (Kiros et al., 2015) to collect a large sentence pairs corpus from Book-Corpus (Zhu et al., 2015) and propose a sentence representation based on that. They also apply their pre-trained sentence encoder to a series of natural language understanding tasks such as sentiment analysis, question-type, entailment, and relatedness. However, all those datasets are provided by Conneau et al. (2017) for evaluating sentence embeddings and are almost all small-scale and are not able to support more complex neural network. Moreover, they represent each sentence by a single vector and directly combine them to predict the answer, which is not able to interact among the words level.

In closely related work, Jernite et al. (2017) propose a model that also leverage discourse relations. However, they manually group the discourse markers into several categories based on human knowledge and predict the category instead of the explicit discourse marker phrase. However, the size of their dataset is much smaller than that in (Nie et al., 2017), and sometimes there has been disagreement among annotators about what exactly is the correct categorization of discourse relations (Hobbs, 1990).

Unlike previous works, we insert the sentence encoder into an integrated network to augment the semantic representation for NLI tasks rather than directly combining the sentence embeddings to predict the relations.

6.2 Natural Language Inference

Earlier research on the natural language inference was based on small-scale datasets (Marelli et al., 2014), which relied on traditional methods such as shallow methods (Glickman et al., 2005), natural logic methods (MacCartney and Manning, 2007), etc. These datasets are either not large enough to support complex deep neural network models or too easy to challenge natural language.

Large and complicated networks have been successful in many natural language processing tasks (Zhu et al., 2017; Chen et al., 2017e; Pan et al., 2017a). Recently, Bowman et al. (2015) released Stanford Natural language Inference (SNLI) dataset, which is a high-quality and large-scale benchmark, thus inspired many significant works (Bowman et al., 2016; Mou et al., 2016; Vendrov et al., 2016; Conneau et al., 2017; Wang

et al., 2017; Gong et al., 2018; McCann et al., 2017; Chen et al., 2017b; Choi et al., 2017; Tay et al., 2017). Most of them focus on the improvement of the interaction architectures and obtain competitive results, while transfer learning from external knowledge is popular as well. Vendrov et al. (2016) incorporated Skipthought (Kiros et al., 2015), which is an unsupervised sequence model that has been proven to generate useful sentence embedding. McCann et al. (2017) proposed to transfer the pre-trained encoder from the neural machine translation (NMT) to the NLI tasks.

Our method combines a pre-trained sentence encoder from the DMP task with an integrated NLI model to compose a novel framework. Furthermore, unlike previous studies, we make full use of the labels provided by the annotators and employ policy gradient to optimize a new objective function in order to simulate the thought of human being.

7 Conclusion

In this paper, we propose *Discourse Marker Augmented Network* for the task of the natural language inference. We transfer the knowledge learned from the discourse marker prediction task to the NLI task to augment the semantic representation of the model. Moreover, we take the various views of the annotators into consideration and employ reinforcement learning to help optimize the model. The experimental evaluation shows that our model achieves the state-of-the-art results on SNLI and MultiNLI datasets. Future works involve the choice of discourse markers and some other transfer learning sources.

8 Acknowledgements

This work was supported in part by the National Nature Science Foundation of China (Grant Nos: 61751307), in part by the grant ZJU Research 083650 of the ZJUI Research Program from Zhejiang University and in part by the National Youth Top-notch Talent Support Program. The experiments are supported by Chengwei Yao in the Experiment Center of the College of Computer Science and Technology, Zhejiang university.

References

Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large an-

notated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 632–642.

Samuel R Bowman, Jon Gauthier, Abhinav Rastogi, Raghav Gupta, Christopher D Manning, and Christopher Potts. 2016. A fast unified model for parsing and sentence understanding. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, volume 1, pages 1466–1477.

Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017a. Reading wikipedia to answer open-domain questions. In *Meeting of the Association for Computational Linguistics (ACL)*, pages 1870–1879.

Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, and Diana Inkpen. 2017b. Natural language inference with external knowledge. *arXiv preprint arXiv:1711.04289*.

Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017c. Enhanced lstm for natural language inference. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, volume 1, pages 1657–1668.

Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017d. Recurrent neural network-based sentence encoder with gated attention for natural language inference. In *Proceedings of the 2nd Workshop on Evaluating Vector Space Representations for NLP*, pages 36–40.

Zheqian Chen, Ben Gao, Huimin Zhang, Zhou Zhao, Haifeng Liu, and Deng Cai. 2017e. User personalized satisfaction prediction via multiple instance deep learning. In *International Conference on World Wide Web (WWW)*, pages 907–915.

Jihun Choi, Kang Min Yoo, and Sang goo Lee. 2017. Learning to compose task-specific tree structures. *The Association for the Advancement of Artificial Intelligence (AAAI)*.

Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning (ICML)*, pages 160–167. ACM.

Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 670–680.

Oren Glickman, Ido Dagan, and Moshe Koppel. 2005. Web based probabilistic textual entailment.

- Yichen Gong, Heng Luo, and Jian Zhang. 2018. [Natural language inference over interaction space](#). In *International Conference on Learning Representations (ICLR)*.
- Jerry R Hobbs. 1990. *Literature and cognition*. 21. Center for the Study of Language (CSLI).
- Yacine Jernite, Samuel R Bowman, and David Sonntag. 2017. Discourse-based objectives for fast unsupervised sentence representation learning. *arXiv preprint arXiv:1705.00557*.
- Ryan Kiros, Yukun Zhu, Ruslan R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. In *Advances in neural information processing systems (NIPS)*, pages 3294–3302.
- Bill MacCartney and Christopher D Manning. 2007. Natural logic for textual inference. In *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*, pages 193–200. Association for Computational Linguistics.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven 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.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, Roberto Zamparelli, et al. 2014. A sick cure for the evaluation of compositional distributional semantic models. In *LREC*, pages 216–223.
- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. 2017. Learned in translation: Contextualized word vectors. In *Advances in Neural Information Processing Systems (NIPS)*, pages 6297–6308.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems (NIPS)*, pages 3111–3119.
- Lili Mou, Rui Men, Ge Li, Yan Xu, Lu Zhang, Rui Yan, and Zhi Jin. 2016. Natural language inference by tree-based convolution and heuristic matching. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 130–136.
- Allen Nie, Erin D Bennett, and Noah D Goodman. 2017. Dissent: Sentence representation learning from explicit discourse relations. *arXiv preprint arXiv:1710.04334*.
- Yixin Nie and Mohit Bansal. 2017. Shortcut-stacked sentence encoders for multi-domain inference. *arXiv preprint arXiv:1708.02312*.
- Boyuan Pan, Hao Li, Zhou Zhao, Deng Cai, and Xiaofei He. 2017a. Keyword-based query comprehending via multiple optimized-demand augmentation. *arXiv preprint arXiv:1711.00179*.
- Boyuan Pan, Hao Li, Zhou Zhao, Bin Cao, Deng Cai, and Xiaofei He. 2017b. Memen: Multi-layer embedding with memory networks for machine comprehension. *arXiv preprint arXiv:1707.09098*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Tim Rocktäschel, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, and Phil Blunsom. 2016. Reasoning about entailment with neural attention. *International Conference on Learning Representations (ICLR)*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.
- Lei Sha, Baobao Chang, Zhifang Sui, and Sujian Li. 2016. Reading and thinking: Re-read lstm unit for textual entailment recognition. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2870–2879.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- Yi Tay, Luu Anh Tuan, and Siu Cheung Hui. 2017. A compare-propagate architecture with alignment factorization for natural language inference. *arXiv preprint arXiv:1801.00102*.
- Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun. 2016. Order-embeddings of images and language. *International Conference on Learning Representations (ICLR)*.
- Zhiguo Wang, Wael Hamza, and Radu Florian. 2017. Bilateral multi-perspective matching for natural language sentences. *International Joint Conference on Artificial Intelligence (IJCAI)*.
- Adina Williams, Nikita Nangia, and Samuel R Bowman. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256.

- Zhilin Yang, Bhuwan Dhingra, Ye Yuan, Junjie Hu, William W. Cohen, and Ruslan Salakhutdinov. 2017. Words or characters? fine-grained gating for reading comprehension. *International Conference on Learning Representations (ICLR)*.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Matthew D Zeiler. 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Yu Zhu, Hao Li, Yikang Liao, Beidou Wang, Ziyu Guan, Haifeng Liu, and Deng Cai. 2017. What to do next: modeling user behaviors by time-lstm. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3602–3608.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision (ICCV)*, pages 19–27.