

DiSAN: Directional Self-Attention Network for RNN/CNN-Free Language Understanding

Tao Shen,[†] Jing Jiang,[†] Tianyi Zhou,[‡] Shirui Pan,[†] Guodong Long,[†] Chengqi Zhang[†]

[†]Centre of Artificial Intelligence, FEIT, University of Technology Sydney

[‡]Paul G. Allen School of Computer Science & Engineering, University of Washington
tao.shen@student.uts.edu.au, tianyizh@uw.edu

{guodong.long, jing.jiang, shirui.pan, chengqi.zhang}@uts.edu.au

Abstract

Recurrent neural nets (RNN) and convolutional neural nets (CNN) are widely used on NLP tasks to capture the long-term and local dependencies, respectively. Attention mechanisms have recently attracted enormous interest due to their highly parallelizable computation, significantly less training time, and flexibility in modeling dependencies. We propose a novel attention mechanism in which the attention between elements from input sequence(s) is directional and multi-dimensional (i.e., feature-wise). A light-weight neural net, “Directional Self-Attention Network (DiSAN)”, is then proposed to learn sentence embedding, based solely on the proposed attention without any RNN/CNN structure. DiSAN is only composed of a directional self-attention with temporal order encoded, followed by a multi-dimensional attention that compresses the sequence into a vector representation. Despite its simple form, DiSAN outperforms complicated RNN models on both prediction quality and time efficiency. It achieves the best test accuracy among all sentence encoding methods and improves the most recent best result by 1.02% on the Stanford Natural Language Inference (SNLI) dataset, and shows state-of-the-art test accuracy on the Stanford Sentiment Treebank (SST), Multi-Genre natural language inference (MultiNLI), Sentences Involving Compositional Knowledge (SICK), Customer Review, MPQA, TREC question-type classification and Subjectivity (SUBJ) datasets.

1 Introduction

Context dependency plays a significant role in language understanding and provides critical information to natural language processing (NLP) tasks. For different tasks and data, researchers often switch between two types of deep neural network (DNN): recurrent neural network (RNN) with sequential architecture capturing long-range dependencies (e.g., long short-term memory (LSTM) (Hochreiter and Schmidhuber 1997) and gated recurrent unit (GRU) (Chung et al. 2014)), and convolutional neural network (CNN) (Kim 2014) whose hierarchical structure is good at extracting local or position-invariant features. However, which network to choose in practice is an open question, and the choice relies largely on the empirical knowledge.

Recent works have found that equipping RNN or CNN with an attention mechanism can achieve state-of-the-art

performance on a large number of NLP tasks, including neural machine translation (Bahdanau, Cho, and Bengio 2015; Luong, Pham, and Manning 2015), natural language inference (Liu et al. 2016), conversation generation (Shang, Lu, and Li 2015), question answering (Hermann et al. 2015; Sukhbaatar et al. 2015), machine reading comprehension (Seo et al. 2017), and sentiment analysis (Kokkinos and Potamianos 2017). The attention uses a hidden layer to compute a categorical distribution over elements from the input sequence to reflect their importance weights. It allows RNN/CNN to maintain a variable-length memory, so that elements from the input sequence can be selected by their importance/relevance and merged into the output. In contrast to RNN and CNN, the attention mechanism is trained to capture the dependencies that make significant contributions to the task, regardless of the distance between the elements in the sequence. It can thus provide complementary information to the distance-aware dependencies modeled by RNN/CNN. In addition, computing attention only requires matrix multiplication, which is highly parallelizable compared to the sequential computation of RNN.

In a very recent work (Vaswani et al. 2017), an attention mechanism is solely used to construct a sequence to sequence (seq2seq) model that achieves a state-of-the-art quality score on the neural machine translation (NMT) task. The seq2seq model, “Transformer”, has an encoder-decoder structure that is only composed of stacked attention networks, without using either recurrence or convolution. The proposed attention, “multi-head attention”, projects the input sequence to multiple subspaces, then applies scaled dot-product attention to its representation in each subspace, and lastly concatenates their output. By doing this, it can combine different attentions from multiple subspaces. This mechanism is used in Transformer to compute both the context-aware features inside the encoder/decoder and the bottleneck features between them.

The attention mechanism has more flexibility in sequence length than RNN/CNN, and is more task/data-driven when modeling dependencies. Unlike sequential models, its computation can be easily and significantly accelerated by existing distributed/parallel computing schemes. However, to the best of our knowledge, a neural net entirely based on attention has not been designed for other NLP tasks except NMT, especially those that cannot be cast into a seq2seq problem.

Compared to RNN, a disadvantage of most attention mechanisms is that the temporal order information is lost, which however might be important to the task. This explains why positional encoding is applied to the sequence before being processed by the attention in Transformer. How to model order information within an attention is still an open problem.

The goal of this paper is to develop a unified and RNN/CNN-free attention network that can be generally utilized to learn the sentence encoding model for different NLP tasks, such as natural language inference, sentiment analysis, sentence classification and semantic relatedness. We focus on the sentence encoding model because it is a basic module of most DNNs used in the NLP literature.

We propose a novel attention mechanism that differs from previous ones in that it is 1) multi-dimensional: the attention w.r.t. each pair of elements from the source(s) is a vector, where each entry is the attention computed on each feature; and 2) directional: it uses one or multiple positional masks to model the asymmetric attention between two elements. We compute feature-wise attention since each element in a sequence is usually represented by a vector, e.g., word/character embedding (Kim et al. 2016), and attention on different features can contain different information about dependency, thus to handle the variation of contexts around the same word. We apply positional masks to attention distribution since they can easily encode prior structure knowledge such as temporal order and dependency parsing. This design mitigates the weakness of attention in modeling order information, and takes full advantage of parallel computing.

We then build a light-weight and RNN/CNN-free neural network, “Directional Self-Attention Network (DiSAN)”, for sentence encoding. This network relies entirely on the proposed attentions and does not use any RNN/CNN structure. In DiSAN, the input sequence is processed by directional (forward and backward) self-attentions to model context dependency and produce context-aware representations for all tokens. Then, a multi-dimensional attention computes a vector representation of the entire sequence, which can be passed into a classification/regression module to compute the final prediction for a particular task. Unlike Transformer, neither stacking of attention blocks nor an encoder-decoder structure is required. The simple architecture of DiSAN leads to fewer parameters, less computation and easier parallelization.

In experiments¹, we compare DiSAN with the currently popular methods on various NLP tasks, e.g., natural language inference, sentiment analysis, sentence classification, etc. DiSAN achieves the highest test accuracy on the Stanford Natural Language Inference (SNLI) dataset among sentence-encoding models and improves the currently best result by 1.02%. It also shows the state-of-the-art performance on the Stanford Sentiment Treebank (SST), Multi-Genre natural language inference (MultiNLI), SICK, Customer Review, MPQA, SUBJ and TREC question-type classification datasets. Meanwhile, it has fewer parameters and exhibits much higher computation efficiency than the mod-

els it outperforms, e.g., LSTM and tree-based models.

Annotation: 1) Lowercase denotes a vector; 2) bold lowercase denotes a sequence of vectors (stored as a matrix); and 3) uppercase denotes a matrix or a tensor.

2 Background

2.1 Sentence Encoding

In the pipeline of NLP tasks, a sentence is denoted by a sequence of discrete tokens (e.g., words or characters) $\mathbf{v} = [v_1, v_2, \dots, v_n]$, where v_i could be a one-hot vector whose dimension length equals the number of distinct tokens N . A pre-trained token embedding (e.g., word2vec (Mikolov et al. 2013b) or GloVe (Pennington, Socher, and Manning 2014)) is applied to \mathbf{v} and transforms all discrete tokens to a sequence of low-dimensional dense vector representations $\mathbf{x} = [x_1, x_2, \dots, x_n]$ with $x_i \in \mathbb{R}^{d_e}$. This pre-process can be written as $\mathbf{x} = W^{(e)}\mathbf{v}$, where word embedding weight matrix $W^{(e)} \in \mathbb{R}^{d_e \times N}$ and $\mathbf{x} \in \mathbb{R}^{d_e \times n}$.

Most DNN sentence-encoding models for NLP tasks take \mathbf{x} as the input and further generate a vector representation u_i for each x_i by context fusion. Then a sentence encoding is obtained by mapping the sequence $\mathbf{u} = [u_1, u_2, \dots, u_n]$ to a single vector $s \in \mathbb{R}^d$, which is used as a compact encoding of the entire sentence in NLP problems.

2.2 Attention

The attention is proposed to compute an alignment score between elements from two sources. In particular, given the token embeddings of a source sequence $\mathbf{x} = [x_1, x_2, \dots, x_n]$ and the vector representation of a query q , attention computes the alignment score between x_i and q by a compatibility function $f(x_i, q)$, which measures the dependency between x_i and q , or the attention of q to x_i . A softmax function then transforms the scores $[f(x_i, q)]_{i=1}^n$ to a probability distribution $p(z|\mathbf{x}, q)$ by normalizing over all the n tokens of \mathbf{x} . Here z is an indicator of which token in \mathbf{x} is important to q on a specific task. That is, large $p(z = i|\mathbf{x}, q)$ means x_i contributes important information to q . The above process can be summarized by the following equations.

$$a = [f(x_i, q)]_{i=1}^n, \quad (1)$$

$$p(z|\mathbf{x}, q) = \text{softmax}(a). \quad (2)$$

Specifically,

$$p(z = i|\mathbf{x}, q) = \frac{\exp(f(x_i, q))}{\sum_{i=1}^n \exp(f(x_i, q))}. \quad (3)$$

The output of this attention mechanism is a weighted sum of the embeddings for all tokens in \mathbf{x} , where the weights are given by $p(z|\mathbf{x}, q)$. It places large weights on the tokens important to q , and can be written as the expectation of a token sampled according to its importance, i.e.,

$$s = \sum_{i=1}^n p(z = i|\mathbf{x}, q)x_i = \mathbb{E}_{i \sim p(z|\mathbf{x}, q)}(x_i), \quad (4)$$

where $s \in \mathbb{R}^{d_e}$ can be used as the sentence encoding of \mathbf{x} .

Additive attention (or multi-layer perceptron attention) (Bahdanau, Cho, and Bengio 2015; Shang, Lu, and Li

¹Codes and pre-trained models for experiments can be found at <https://github.com/taoshen58/DiSAN>

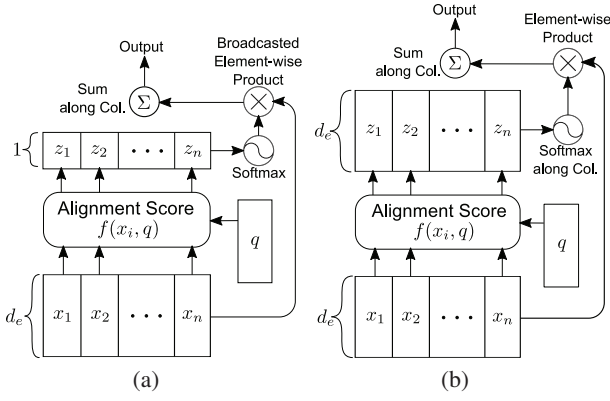


Figure 1: (a) Traditional (additive/multiplicative) attention and (b) multi-dimensional attention. z_i denotes alignment score $f(x_i, q)$, which is a scalar in (a) but a vector in (b).

2015) and multiplicative attention (or dot-product attention) (Vaswani et al. 2017; Sukhbaatar et al. 2015; Rush, Chopra, and Weston 2015) are the two most commonly used attention mechanisms. They share the same and unified form of attention introduced above, but are different in the compatibility function $f(x_i, q)$. Additive attention is associated with

$$f(x_i, q) = w^T \sigma(W^{(1)}x_i + W^{(2)}q), \quad (5)$$

where $\sigma(\cdot)$ is an activation function and $w \in \mathbb{R}^{d_e}$ is a weight vector. Multiplicative attention uses inner product or cosine similarity for $f(x_i, q)$, i.e.,

$$f(x_i, q) = \langle W^{(1)}x_i, W^{(2)}q \rangle. \quad (6)$$

In practice, additive attention often outperforms multiplicative one in prediction quality, but the latter is faster and more memory-efficient due to optimized matrix multiplication.

2.3 Self-Attention

Self-Attention is a special case of the attention mechanism introduced above. It replaces q with a token embedding x_j from the source input itself. It relates elements at different positions from a single sequence by computing the attention between each pair of tokens, x_i and x_j . It is very expressive and flexible for both long-range and local dependencies, which used to be respectively modeled by RNN and CNN. Moreover, it has much faster computation speed and fewer parameters than RNN. In recent works, we have already witnessed its success across a variety of NLP tasks, such as reading comprehension (Hu, Peng, and Qiu 2017) and neural machine translation (Vaswani et al. 2017).

3 Two Proposed Attention Mechanisms

In this section, we introduce two novel attention mechanisms, multi-dimensional attention in Section 3.1 (with two extensions to self-attention in Section 3.2) and directional self-attention in Section 3.3. They are the main components of DiSAN and may be of independent interest to other neural nets for other NLP problems in which an attention is needed.

3.1 Multi-dimensional Attention

Multi-dimensional attention is a natural extension of additive attention (or MLP attention) at the feature level. Instead of computing a single scalar score $f(x_i, q)$ for each token x_i as shown in Eq.(5), multi-dimensional attention computes a feature-wise score vector for x_i by replacing weight vector w in Eq.(5) with a matrix W , i.e.,

$$f(x_i, q) = W^T \sigma \left(W^{(1)}x_i + W^{(2)}q \right), \quad (7)$$

where $f(x_i, q) \in \mathbb{R}^{d_e}$ is a vector with the same length as x_i , and all the weight matrices $W, W^{(1)}, W^{(2)} \in \mathbb{R}^{d_e \times d_e}$. We further add two bias terms to the parts in and out activation $\sigma(\cdot)$, i.e.,

$$f(x_i, q) = W^T \sigma \left(W^{(1)}x_i + W^{(2)}q + b^{(1)} \right) + b. \quad (8)$$

We then compute a categorical distribution $p(z_k|x, q)$ over all the n tokens for each feature $k \in [d_e]$. A large $p(z_k = i|x, q)$ means that feature k of token i is important to q .

We apply the same procedure Eq.(1)-(3) in traditional attention to the k^{th} dimension of $f(x_i, q)$. In particular, for each feature $k \in [d_e]$, we replace $f(x_i, q)$ with $[f(x_i, q)]_k$, and change z to z_k in Eq.(1)-(3). Now each feature k in each token i has an importance weight $P_{ki} \triangleq p(z_k = i|x, q)$. The output s can be written as

$$s = \left[\sum_{i=1}^n P_{ki} \mathbf{x}_{ki} \right]_{k=1}^{d_e} = \left[\mathbb{E}_{i \sim p(z_k|x, q)} (\mathbf{x}_{ki}) \right]_{k=1}^{d_e}. \quad (9)$$

We give an illustration of traditional attention and multi-dimensional attention in Figure 1. In the rest of this paper, we will ignore the subscript k which indexes feature dimension for simplification if no confusion is possible. Hence, the output s can be written as an element-wise product $s = \sum_{i=1}^n P_{ki} \odot x_i$

Remark: The word embedding usually suffers from the polysemy in natural language. Since traditional attention computes a single importance score for each word based on the word embedding, it cannot distinguish the meanings of the same word in different contexts. Multi-dimensional attention, however, computes a score for each feature of each word, so it can select the features that can best describe the word’s specific meaning in any given context, and include this information in the sentence encoding output s .

3.2 Two types of Multi-dimensional Self-attention

When extending multi-dimension to self-attentions, we have two variants of multi-dimensional attention. The first one, called multi-dimensional “**token2token**” self-attention, explores the dependency between x_i and x_j from the same source \mathbf{x} , and generates context-aware coding for each element. It replaces q with x_j in Eq.(8), i.e.,

$$f(x_i, x_j) = W^T \sigma \left(W^{(1)}x_i + W^{(2)}x_j + b^{(1)} \right) + b. \quad (10)$$

Similar to P in vanilla multi-dimensional attention, we compute a probability matrix $P^j \in \mathbb{R}^{d_e \times n}$ for each x_j such that $P_{ki}^j \triangleq p(z_k = i|x, x_j)$. The output for x_j is

$$s_j = \sum_{i=1}^n P_{ki}^j \odot x_i \quad (11)$$

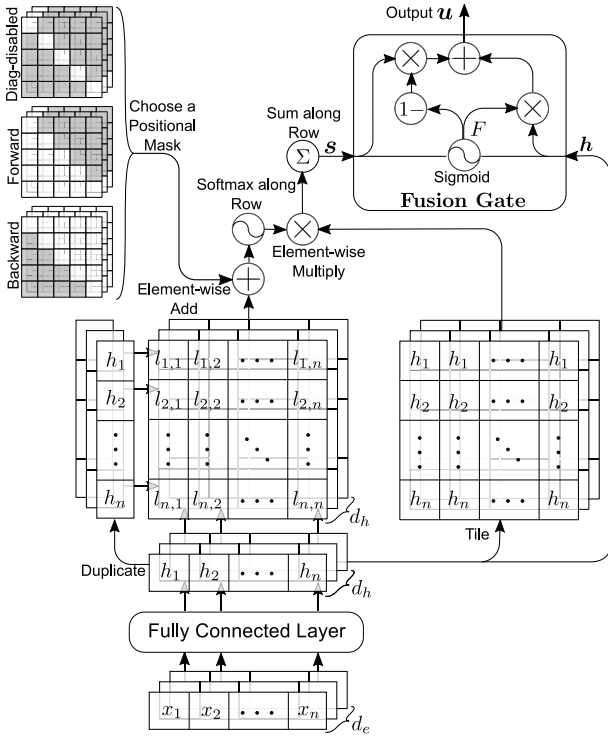


Figure 2: Directional self-attention (DiSA) mechanism. Here, we use $l_{i,j}$ to denote $f(h_i, h_j)$ in Eq. (15).

The output of token2token self-attention for all elements from \mathbf{x} is $\mathbf{s} = [s_1, s_2, \dots, s_n] \in \mathbb{R}^{d_e \times n}$.

The second one, multi-dimensional “source2token” self-attention, explores the dependency between x_i and the entire sequence \mathbf{x} , and compresses the sequence \mathbf{x} into a vector. It removes q from Eq.(8), i.e.,

$$f(x_i) = W^T \sigma \left(W^{(1)} x_i + b^{(1)} \right) + b. \quad (12)$$

The probability matrix is defined as $P_{ki} \triangleq p(z_k = i | \mathbf{x})$ and is computed in the same way as P in vanilla multi-dimensional attention. The output \mathbf{s} is also same, i.e.,

$$\mathbf{s} = \sum_{i=1}^n P_{\cdot i} \odot x_i \quad (13)$$

We will use these two types (i.e., token2token and source2token) of multi-dimensional self-attention in different parts of our sentence encoding model, DiSAN.

3.3 Directional Self-Attention

Directional self-attention (DiSA) is composed of a fully connected layer whose input is the token embeddings \mathbf{x} , a “masked” multi-dimensional token2token self-attention block to explore the dependency and temporal order, and a fusion gate to combine the output and input of the attention block. Its structure is shown in Figure 2. It can be used as either a neural net or a module to compose a large network.

In DiSA, we first transform the input sequence $\mathbf{x} = [x_1, x_2, \dots, x_n]$ to a sequence of hidden state $\mathbf{h} = [h_1, h_2, \dots, h_n]$ by a fully connected layer, i.e.,

$$\mathbf{h} = \sigma_h \left(W^{(h)} \mathbf{x} + b^{(h)} \right), \quad (14)$$

where $\mathbf{x} \in \mathbb{R}^{d_e \times n}$, $\mathbf{h} \in \mathbb{R}^{d_h \times n}$, $W^{(h)}$ and $b^{(h)}$ are the learnable parameters, and $\sigma_h(\cdot)$ is an activation function.

We then apply multi-dimensional token2token self-attention to \mathbf{h} , and generate context-aware vector representations \mathbf{s} for all elements from the input sequence. We make two modifications to Eq.(10) to reduce the number of parameters and make the attention directional.

First, we set W in Eq.(10) to a scalar c and divide the part in $\sigma(\cdot)$ by c , and we use $\tanh(\cdot)$ for $\sigma(\cdot)$, which reduces the number of parameters. In experiments, we always set $c = 5$, and obtain stable output.

Second, we apply a positional mask to Eq.(10), so the attention between two elements can be asymmetric. Given a mask $M \in \{0, -\infty\}^{n \times n}$, we set bias b to a constant vector $M_{ij} \mathbf{1}$ in Eq.(10), where $\mathbf{1}$ is an all-one vector. Hence, Eq.(10) is modified to

$$f(h_i, h_j) = c \cdot \tanh \left([W^{(1)} h_i + W^{(2)} h_j + b^{(1)}] / c \right) + M_{ij} \mathbf{1}. \quad (15)$$

To see why a mask can encode directional information, let us consider a case in which $M_{ij} = -\infty$ and $M_{ji} = 0$, which results in $[f(h_i, h_j)]_k = -\infty$ and unchanged $[f(h_j, h_i)]_k$. Since the probability $p(z_k = i | \mathbf{x}, x_j)$ is computed by softmax, $[f(h_i, h_j)]_k = -\infty$ leads to $p(z_k = i | \mathbf{x}, x_j) = 0$. This means that there is no attention of x_j to x_i on feature k . On the contrary, we have $p(z_k = j | \mathbf{x}, x_i) > 0$, which means that attention of x_i to x_j exists on feature k . Therefore, prior structure knowledge such as temporal order and dependency parsing can be easily encoded by the mask, and explored in generating sentence encoding. This is an important feature of DiSA that previous attention mechanisms do not have.

For self-attention, we usually need to disable the attention of each token to itself (Hu, Peng, and Qiu 2017). This is the same as applying a diagonal-disabled (i.e., diag-disabled) mask such that

$$M_{ij}^{diag} = \begin{cases} 0, & i \neq j \\ -\infty, & i = j \end{cases} \quad (16)$$

Moreover, we can use masks to encode temporal order information into attention output. In this paper, we use two masks, i.e., forward mask M^{fw} and backward mask M^{bw} ,

$$M_{ij}^{fw} = \begin{cases} 0, & i < j \\ -\infty, & otherwise \end{cases} \quad (17)$$

$$M_{ij}^{bw} = \begin{cases} 0, & i > j \\ -\infty, & otherwise \end{cases} \quad (18)$$

In forward mask M^{fw} , there is the only attention of later token j to early token i , and vice versa in backward mask. We show these three positional masks in Figure 3.

Given input sequence \mathbf{x} and a mask M , we compute $f(x_i, x_j)$ according to Eq.(15), and follow the standard procedure of multi-dimensional token2token self-attention to

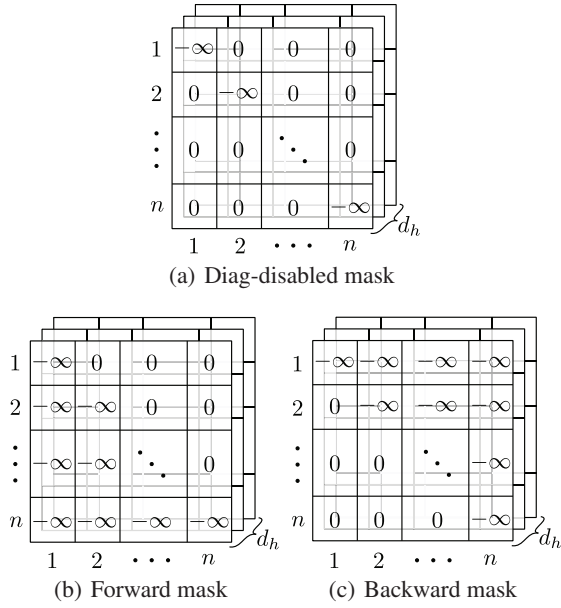


Figure 3: Three positional masks: (a) is the diag-disabled mask M^{diag} ; (b) and (c) are forward mask M^{fw} and backward mask M^{bw} , respectively.

compute the probability matrix P^j for each $j \in [n]$. Each output s_j in \mathbf{s} is computed as in Eq.(11).

The final output $\mathbf{u} \in \mathbb{R}^{d_h \times n}$ of DiSA is obtained by combining the output \mathbf{s} and the input \mathbf{h} of the masked multi-dimensional token2token self-attention block. This yields a temporal order encoded and context-aware vector representation for each element/token. The combination is accomplished by a dimension-wise fusion gate, i.e.,

$$F = \text{sigmoid} \left(W^{(f1)} \mathbf{s} + W^{(f2)} \mathbf{h} + b^{(f)} \right) \quad (19)$$

$$\mathbf{u} = F \odot \mathbf{h} + (1 - F) \odot \mathbf{s} \quad (20)$$

where $W^{(f1)}, W^{(f2)} \in \mathbb{R}^{d_h \times d_h}$ and $b^{(f)} \in \mathbb{R}^{d_h}$ are the learnable parameters of the fusion gate.

4 Directional Self-Attention Network

We propose a light-weight network, ‘‘Directional Self-Attention Network (DiSAN)’’, for sentence encoding. Its architecture is shown in Figure 4.

Given an input sequence of token embedding \mathbf{x} , DiSAN firstly applies two parameter-untied DiSA blocks with forward mask M^{fw} Eq.(17) and M^{bw} Eq.(18), respectively. The feed-forward procedure is given in Eq.(14)-(15) and Eq.(19)-(20). Their outputs are denoted by $\mathbf{u}^{fw}, \mathbf{u}^{bw} \in \mathbb{R}^{d_h \times n}$. We concatenate them vertically as $[\mathbf{u}^{fw}; \mathbf{u}^{bw}] \in \mathbb{R}^{2d_h \times n}$, and use this concatenated output as input to a multi-dimensional source2token self-attention block, whose output $\mathbf{s}^{disan} \in \mathbb{R}^{2d_h}$ computed by Eq.(12)-(13) is the final sentence encoding result of DiSAN.

Remark: In DiSAN, forward/backward DiSA blocks work as context fusion layers. And the multi-dimensional source2token self-attention compresses the sequence into a

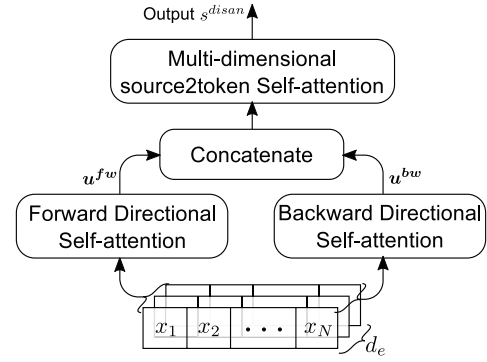


Figure 4: Directional self-attention network (DiSAN)

single vector. The idea of using both forward and backward attentions is inspired by Bi-directional LSTM (Bi-LSTM) (Graves, Jaitly, and Mohamed 2013), in which forward and backward LSTMs are used to encode long-range dependency from different directions. In Bi-LSTM, LSTM combines the context-aware output with the input by multi-gate. The fusion gate used in DiSA shares the similar motivation. However, DiSAN has fewer parameters, simpler structure and better efficiency.

5 Experiments

In this section, we first apply DiSAN to natural language inference and sentiment analysis tasks. DiSAN achieves the state-of-the-art performance and significantly better efficiency than other baseline methods on benchmark datasets for both tasks. We also conduct experiments on other NLP tasks and DiSAN also achieves state-of-the-art performance.

Training Setup: We use cross-entropy loss plus L2 regularization penalty as optimization objective. We minimize it by Adadelata (Zeiler 2012) (an optimizer of mini-batch SGD) with batch size of 64. We use Adadelata rather than Adam (Kingma and Ba 2015) because in our experiments, DiSAN optimized by Adadelata can achieve more stable performance than Adam optimized one. Initial learning rate is set to 0.5. All weight matrices are initialized by Glorot Initialization (Glorot and Bengio 2010), and the biases are initialized with 0. We initialize the word embedding in \mathbf{x} by 300D GloVe 6B pre-trained vectors (Pennington, Socher, and Manning 2014). The Out-of-Vocabulary words in training set are randomly initialized by uniform distribution between $(-0.05, 0.05)$. The word embeddings are fine-tuned during the training phrase. We use Dropout (Srivastava et al. 2014) with keep probability 0.75 for language inference and 0.8 for sentiment analysis. The L2 regularization decay factors γ are 5×10^{-5} and 10^{-4} for language inference and sentiment analysis, respectively. Note that the dropout keep probability and γ varies with the scale of corresponding dataset. Hidden units number d_h is set to 300. Activation functions $\sigma(\cdot)$ are *ELU* (exponential linear unit) (Clevert, Unterthiner, and Hochreiter 2016) if not specified. All models are implemented with TensorFlow² and run on sin-

²<https://www.tensorflow.org>

Model Name	$ \theta $	T(s)/epoch	Train Accu(%)	Test Accu(%)
Unlexicalized features (Bowman et al. 2015)			49.4	50.4
+ Unigram and bigram features (Bowman et al. 2015)			99.7	78.2
100D LSTM encoders (Bowman et al. 2015)	0.2m		84.8	77.6
300D LSTM encoders (Bowman et al. 2016)	3.0m		83.9	80.6
1024D GRU encoders (Vendrov et al. 2016)	15m		98.8	81.4
300D Tree-based CNN encoders (Mou et al. 2016)	3.5m		83.3	82.1
300D SPINN-PI encoders (Bowman et al. 2016)	3.7m		89.2	83.2
600D Bi-LSTM encoders (Liu et al. 2016)	2.0m		86.4	83.3
300D NTI-SLSTM-LSTM encoders (Munkhdalai and Yu 2017b)	4.0m		82.5	83.4
600D Bi-LSTM encoders+intra-attention (Liu et al. 2016)	2.8m		84.5	84.2
300D NSE encoders (Munkhdalai and Yu 2017a)	3.0m		86.2	84.6
Word Embedding with additive attention	0.45m	216	82.39	79.81
Word Embedding with s2t self-attention	0.54m	261	86.22	83.12
Multi-head with s2t self-attention	1.98m	345	89.58	84.17
Bi-LSTM with s2t self-attention	2.88m	2080	90.39	84.98
DiSAN without directions	2.35m	592	90.18	84.66
Directional self-attention network (DiSAN)	2.35m	587	91.08	85.62

Table 1: Experimental results for different methods on SNLI. $|\theta|$: the number of parameters (excluding word embedding part). **T(s)/epoch**: average time (second) per epoch. **Train Accu(%)** and **Test Accu(%)**: the accuracy on training and test set.

gle Nvidia GTX 1080Ti graphic card.

5.1 Natural Language Inference

The goal of Natural Language Inference (NLI) is to reason the semantic relationship between a premise sentence and a corresponding hypothesis sentence. The possible relationship could be *entailment*, *neutral* or *contradiction*. We compare different models on a widely used benchmark, Stanford Natural Language Inference (SNLI)³ (Bowman et al. 2015) dataset, which consists of 549,367/9,842/9,824 (train/dev/test) premise-hypothesis pairs with labels.

Following the standard procedure in Bowman et al. (2016), we launch two sentence encoding models (e.g., DiSAN) with tied parameters for the premise sentence and hypothesis sentence, respectively. Given the output encoding s^p for the premise and s^h for the hypothesis, the representation of relationship is the concatenation of s^p , s^h , $s^p - s^h$ and $s^p \odot s^h$, which is fed into a 300D fully connected layer and then a 3-unit output layer with softmax to compute a probability distribution over the three types of relationship.

For thorough comparison, besides the neural nets proposed in previous works of NLI, we implement five extra neural net baselines to compare with DiSAN. They help us to analyze the improvement contributed by each part of DiSAN and to verify that the two attention mechanisms proposed in Section 3 can improve other networks.

- **Word Embedding with additive attention.**
- **Word Embedding with s2t self-attention:** DiSAN with DiSA blocks removed.
- **Multi-head with s2t self-attention:** Multi-head attention (Vaswani et al. 2017) (8 heads, each has 75 hidden units) with source2token self-attention. The positional encoding

method used in Vaswani et al. (2017) is applied to the input sequence to encode temporal information. We find our experiments show that multi-head attention is sensitive to hyperparameters, so we adjust keep probability of dropout from 0.7 to 0.9 with step 0.05 and report the best result.

- **Bi-LSTM with s2t self-attention:** a multi-dimensional source2token self-attention block is applied to the output of Bi-LSTM (300D forward + 300D backward LSTMs).
- **DiSAN without directions:** DiSAN with the forward/backward masks M^{fw} and M^{bw} replaced with two diag-disabled masks M^{diag} , i.e., DiSAN without forward/backward order information.

Compared to the results from the official leaderboard of SNLI in Table 1, DiSAN outperforms previous works and improves the best latest test accuracy (achieved by a memory-based NSE encoder network) by a remarkable margin of 1.02%. DiSAN surpasses the RNN/CNN based models with more complicated architecture and more parameters by large margins, e.g., +2.32% to Bi-LSTM, +1.42% to Bi-LSTM with additive attention. It even outperforms models with the assistance of a semantic parsing tree, e.g., +3.52% to Tree-based CNN, +2.42% to SPINN-PI.

In the results of the five baseline methods and DiSAN at the bottom of Table 1, we demonstrate that making attention multi-dimensional (feature-wise) or directional brings substantial improvement to different neural nets. First, a comparison between the first two models shows that changing token-wise attention to multi-dimensional/feature-wise attention leads to 3.31% improvement on a word embedding based model. Also, a comparison between the third baseline and DiSAN shows that DiSAN can substantially outperform multi-head attention by 1.45%. Moreover, a comparison between the fourth baseline and DiSAN shows that the DiSA block can even outperform Bi-LSTM layer in context encoding, improving test accuracy by 0.64%. A comparison

³<https://nlp.stanford.edu/projects/snli/>

between the fifth baseline and DiSAN shows that directional self-attention with forward and backward masks (with temporal order encoded) can bring 0.96% improvement.

Additional advantages of DiSAN shown in Table 1 are its fewer parameters and compelling time efficiency. It is $\times 3$ faster than widely used Bi-LSTM model. Compared to other models with competitive performance, e.g., 600D Bi-LSTM encoders with intra-attention (2.8M), 300D NSE encoders (3.0M) and 600D Bi-LSTM encoders with multi-dimensional attention (2.88M), DiSAN only has 2.35M parameters.

5.2 Sentiment Analysis

Model	Test Accu
MV-RNN (Socher et al. 2013)	44.4
RNTN (Socher et al. 2013)	45.7
Bi-LSTM (Li et al. 2015)	49.8
Tree-LSTM (Tai, Socher, and Manning 2015)	51.0
CNN-non-static (Kim 2014)	48.0
CNN-Tensor (Lei, Barzilay, and Jaakkola 2015)	51.2
NCSL (Teng, Vo, and Zhang 2016)	51.1
LR-Bi-LSTM (Qian, Huang, and Zhu 2017)	50.6
Word Embedding with additive attention	47.47
Word Embedding with s2t self-attention	48.87
Multi-head with s2t self-attention	49.14
Bi-LSTM with s2t self-attention	49.95
DiSAN without directions	49.41
DiSAN	51.72

Table 2: Test accuracy of fine-grained sentiment analysis on Stanford Sentiment Treebank (SST) dataset.

Sentiment analysis aims to analyze the sentiment of a sentence or a paragraph, e.g., a movie or a product review. We use Stanford Sentiment Treebank (SST)⁴ (Socher et al. 2013) for the experiments, and only focus on the fine-grained movie review sentiment classification over five classes, i.e., very negative, negative, neutral, positive and very positive. We use the standard train/dev/test sets split with 8,544/1,101/2,210 samples. Similar to Section 5.1, we employ a single sentence encoding model to obtain a sentence representation s of a movie review, then pass it into a 300D fully connected layer. Finally, a 5-unit output layer with softmax is used to calculate a probability distribution over the five classes.

In Table 2, we compare previous works with DiSAN on test accuracy. To the best of our knowledge, DiSAN improves the last best accuracy (given by CNN-Tensor) by 0.52%. Compared to tree-based models with heavy use of the prior structure, e.g., MV-RNN, RNTN and Tree-LSTM, DiSAN outperforms them by 7.32%, 6.02% and 0.72%, respectively. Additionally, DiSAN achieves better performance than CNN-based models. More recent works tend to focus on lexicon-based sentiment analysis, by exploring sentiment lexicons, negation words and intensity words.

⁴<https://nlp.stanford.edu/sentiment/>

Nonetheless, DiSAN still outperforms these fancy models, such as NCSL (+0.62%) and LR-Bi-LSTM (+1.12%).

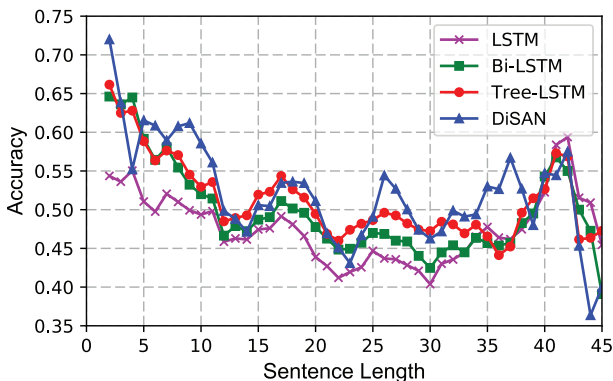


Figure 5: Fine-grained sentiment analysis accuracy vs. sentence length. The results of LSTM, Bi-LSTM and Tree-LSTM are from Tai, Socher, and Manning (2015) and the result of DiSAN is the average over five random trials.

It is also interesting to see the performance of different models on the sentences with different lengths. In Figure 5, we compare LSTM, Bi-LSTM, Tree-LSTM and DiSAN on different sentence lengths. In the range of (5, 12), the length range for most movie review sentences, DiSAN significantly outperforms others. Meanwhile, DiSAN also shows impressive performance for slightly longer sentences or paragraphs in the range of (25, 38). DiSAN performs poorly when the sentence length ≥ 38 , in which however only 3.21% of total movie review sentences lie.

5.3 Experiments on Other NLP Tasks

Multi-Genre Natural Language Inference Multi-Genre Natural Language Inference (MultiNLI)⁵ (Williams, Nangia, and Bowman 2017) dataset consists of 433k sentence pairs annotated with textual entailment information. This dataset is similar to SNLI, but it covers more genres of spoken and written text, and supports a distinctive cross-genre generalization evaluation. However, MultiNLI is a quite new dataset, and its leaderboard does not include a session for the sentence-encoding only model. Hence, we only compare DiSAN with the baselines provided at the official website. The results of DiSAN and two sentence-encoding models on the leaderboard are shown in Table 3. Note that the prediction accuracies of Matched and Mismatched test datasets are obtained by submitting our test results to Kaggle open evaluation platforms⁶: *MultiNLI Matched Open Evaluation* and *MultiNLI Mismatched Open Evaluation*.

Semantic Relatedness The task of semantic relatedness aims to predict a similarity degree of a given pair of sentences. We show an experimental comparison of different

⁵<https://www.nyu.edu/projects/bowman/multinli/>

⁶<https://inclass.kaggle.com/c/multinli-matched-open-evaluation> and <https://inclass.kaggle.com/c/multinli-mismatched-open-evaluation>

Method	Matched	Mismatched
cBoW	0.65200	0.64759
Bi-LSTM	0.67507	0.67248
DiSAN	0.70977	0.71402

Table 3: Experimental results of prediction accuracy for different methods on MultiNLI.

methods on Sentences Involving Compositional Knowledge (SICK)⁷ dataset (Marelli et al. 2014). SICK is composed of 9,927 sentence pairs with 4,500/500/4,927 instances for train/dev/test. The regression module on the top of DiSAN is introduced by Tai, Socher, and Manning (2015). The results in Table 4 show that DiSAN outperforms the models from previous works in terms of Pearson’s r and Spearman’s ρ indexes.

Model	Pearson’s r	Spearman’s ρ	MSE
Meaning Factory ^a	.8268	.7721	.3224
ECNU ^b	.8414	/	/
DT-RNN ^c	.7923 (.0070)	.7319 (.0071)	.3822 (.0137)
SDT-RNN ^c	.7900 (.0042)	.7304 (.0042)	.3848 (.0042)
Cons. Tree-LSTM ^d	.8582 (.0038)	.7966 (.0053)	.2734 (.0108)
Dep. Tree-LSTM ^d	.8676 (.0030)	.8083 (.0042)	.2532 (.0052)
DiSAN	.8695 (.0012)	.8139 (.0012)	.2879 (.0036)

Table 4: Experimental results for different methods on SICK sentence relatedness dataset. The reported accuracies are the mean of five runs (standard deviations in parentheses). Cons. and Dep. represent Constituency and Dependency, respectively. ^a(Bjerva et al. 2014), ^b(Zhao, Zhu, and Lan 2014), ^c(Socher et al. 2014), ^d(Tai, Socher, and Manning 2015).

Sentence Classifications The goal of sentence classification is to correctly predict the class label of a given sentence in various scenarios. We evaluate the models on four sentence classification benchmarks of various NLP tasks, such as sentiment analysis and question-type classification. They are listed as follows. 1) **CR**: Customer review (Hu and Liu 2004) of various products (cameras, etc.), which is to predict whether the review is positive or negative; 2) **MPQA**: Opinion polarity detection subtask of the MPQA dataset (Wiebe, Wilson, and Cardie 2005); 3) **SUBJ**: Subjectivity dataset (Pang and Lee 2004) whose labels indicate whether each sentence is subjective or objective; 4) **TREC**: TREC question-type classification dataset (Li and Roth 2002). The experimental results of DiSAN and existing methods are shown in Table 5.

5.4 Case Study

To gain a closer view of what dependencies in a sentence can be captured by DiSAN, we visualize the attention probability $p(z = i | \mathbf{x}, x_j)$ or alignment score by heatmaps. In

⁷<http://clic.cimec.unitn.it/composes/sick.html>

Model	CR	MPQA	SUBJ	TREC
cBoW ^a	79.9	86.4	91.3	87.3
Skip-thought ^b	81.3	87.5	93.6	92.2
DCNN ^c	/	/	/	93.0
AdaSent ^d	83.6 (1.6)	90.4 (0.7)	92.2 (1.2)	91.1 (1.0)
SRU ^e	84.8 (1.3)	89.7 (1.1)	93.4 (0.8)	93.9 (0.6)
Wide CNNs ^e	82.2 (2.2)	88.8 (1.2)	92.9 (0.7)	93.2 (0.5)
DiSAN	84.8 (2.0)	90.1 (0.4)	94.2 (0.6)	94.2 (0.1)

Table 5: Experimental results for different methods on various sentence classification benchmarks. The reported accuracies on CR, MPQA and SUBJ are the mean of 10-fold cross validation, the accuracies on TREC are the mean of dev accuracies of five runs. All standard deviations are in parentheses. ^a(Mikolov et al. 2013a), ^b(Kiros et al. 2015), ^c(Kalchbrenner, Grefenstette, and Blunsom 2014), ^d(Zhao, Lu, and Poupart 2015), ^e(Lei and Zhang 2017).

particular, we will focus primarily on the probability in forward/backward DiSA blocks (Figure 6), forward/backward fusion gates F in Eq.(19) (Figure 7), and the probability in multi-dimensional source2token self-attention block (Figure 8). For the first two, we desire to demonstrate the dependency at token level, but attention probability in DiSAN is defined on each feature, so we average the probabilities along the feature dimension.

We select two sentences from SNLI test set as examples for this case study. Sentence 1 is *Families have some dogs in front of a carousel and sentence 2 is volleyball match is in progress between ladies.*

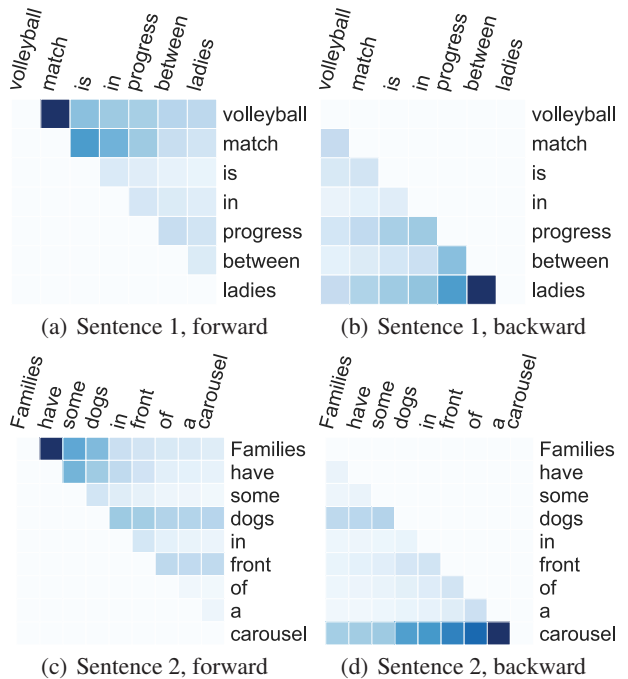


Figure 6: Attention probability in forward/backward DiSA blocks for the two example sentences.

Figure 6 shows that 1) semantically important words such as nouns and verbs usually get large attention, but stop words (*am, is, are, etc.*) do not; 2) globally important words, e.g., *volleyball, match, ladies* in sentence 1 and *dog, front, carousel* in sentence 2, get large attention from all other words; 3) if a word is important to only some of the other words (e.g. to constitute a phrase or sense-group), it gets large attention only from these words, e.g., attention between *progress, between* in sentence 1, and attention between *families, have* in sentence 2.

This also shows that directional information can help to generate context-aware word representation with temporal order encoded. For instance, for word *match* in sentence 1, its forward DiSA focuses more on word *volleyball*, while its backward attention focuses more on *progress* and *ladies*, so the representation of word *match* contains the essential information of the entire sentence, and simultaneously includes the positional order information.

In addition, forward and backward DiSAs can focus on different parts of a sentence. For example, the forward one in sentence 2 pays attention to the word *families*, whereas the backward one focuses on the word *carousel*. Since forward and backward attentions are computed separately, it avoids normalization over multiple significant words to weaken their weights. Note that this is a weakness of traditional attention compared to RNN, especially for long sentences.

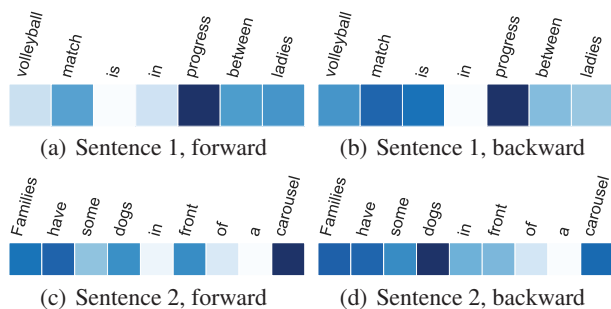


Figure 7: Fusion Gate F in forward/backward DiSA blocks.

In Figure 7, we show that the gate value F in Eq.(19). The gate combines the input and output of masked self-attention. It tends to select the input representation h instead of the output s if the corresponding weight in F is large. This shows that the gate values for meaningless words, especially stop words is small. The stop words themselves cannot contribute important information, so only their semantic relations to other words might help to understand the sentence. Hence, the gate tends to use their context features given by masked self-attention.

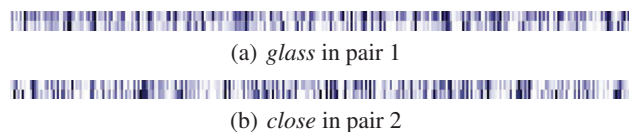


Figure 8: Two pairs of attention probability comparison of same word in difference sentence contexts.

In Figure 8, we show the two multi-dimensional source2token self-attention score vectors of the same word in the two sentences, by their heatmaps. The first pair has two sentences: one is *The glass bottle is big*, and another is *A man is pouring a glass of tea*. They share the same word is *glass* with different meanings. The second pair has two sentences: one is *The restaurant is about to close* and another is *A biker is close to the fountain*. It can be seen that the two attention vectors for the same words are very different due to their different meanings in different contexts. This indicates that the multi-dimensional attention vector is not redundant because it can encode more information than one single score used in traditional attention and it is able to capture subtle difference of the same word in different contexts or sentences. Additionally, it can also alleviate the weakness of the attention over long sequence, which can avoid normalization over entire sequence in traditional attention only once.

6 Conclusion

In this paper, we propose two novel attention mechanisms, multi-dimensional attention and directional self-attention. The multi-dimensional attention performs a feature-wise selection over the input sequence for a specific task, and the directional self-attention uses the positional masks to produce the context-aware representations with temporal information encoded. Based on these attentions, Directional Self-Attention Network (DiSAN) is proposed for sentence-encoding without any recurrent or convolutional structure. The experiment results show that DiSAN can achieve state-of-the-art inference quality and outperform existing works (LSTM, etc.) on a wide range of NLP tasks with fewer parameters and higher time efficiency.

In future work, we will explore the approaches to using the proposed attention mechanisms on more sophisticated tasks, e.g. question answering and reading comprehension, to achieve better performance on various benchmarks.

7 Acknowledgments

This research was funded by the Australian Government through the Australian Research Council (ARC) under grant 1) LP160100630 partnership with Australia Government Department of Health, and 2) LP150100671 partnership with Australia Research Alliance for Children and Youth (ARACY) and Global Business College Australia (GBCA).

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Bjerva, J.; Bos, J.; Van der Goot, R.; and Nissim, M. 2014. The meaning factory: Formal semantics for recognizing textual entailment and determining semantic similarity. In *SemEval@ COLING*, 642–646.
- Bowman, S. R.; Angeli, G.; Potts, C.; and Manning, C. D. 2015. A large annotated corpus for learning natural language inference. In *EMNLP*.
- Bowman, S. R.; Gauthier, J.; Rastogi, A.; Gupta, R.; Manning, C. D.; and Potts, C. 2016. A fast unified model for parsing and sentence understanding. In *ACL*.

- Chung, J.; Gulcehre, C.; Cho, K.; and Bengio, Y. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *NIPS*.
- Clevert, D.-A.; Unterthiner, T.; and Hochreiter, S. 2016. Fast and accurate deep network learning by exponential linear units (elus). In *ICLR*.
- Glorot, X., and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 249–256.
- Graves, A.; Jaitly, N.; and Mohamed, A.-r. 2013. Hybrid speech recognition with deep bidirectional lstm. In *Automatic Speech Recognition and Understanding (ASRU), 2013 IEEE Workshop on*, 273–278. IEEE.
- Hermann, K. M.; Kocisky, T.; Grefenstette, E.; Espeholt, L.; Kay, W.; Suleyman, M.; and Blunsom, P. 2015. Teaching machines to read and comprehend. In *NIPS*.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- Hu, M., and Liu, B. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, 168–177. ACM.
- Hu, M.; Peng, Y.; and Qiu, X. 2017. Reinforced mnemonic reader for machine comprehension. *arXiv preprint arXiv:1705.02798*.
- Kalchbrenner, N.; Grefenstette, E.; and Blunsom, P. 2014. A convolutional neural network for modelling sentences. *arXiv preprint arXiv:1404.2188*.
- Kim, Y.; Jernite, Y.; Sontag, D.; and Rush, A. M. 2016. Character-aware neural language models. In *AAAI*.
- Kim, Y. 2014. Convolutional neural networks for sentence classification. In *EMNLP*.
- Kingma, D., and Ba, J. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Kiros, R.; Zhu, Y.; Salakhutdinov, R. R.; Zemel, R.; Urtasun, R.; Torralba, A.; and Fidler, S. 2015. Skip-thought vectors. In *NIPS*.
- Kokkinos, F., and Potamianos, A. 2017. Structural attention neural networks for improved sentiment analysis. *arXiv preprint arXiv:1701.01811*.
- Lei, T., and Zhang, Y. 2017. Training rnns as fast as cnns. *arXiv preprint arXiv:1709.02755*.
- Lei, T.; Barzilay, R.; and Jaakkola, T. 2015. Molding cnns for text: non-linear, non-consecutive convolutions. In *EMNLP*.
- Li, X., and Roth, D. 2002. Learning question classifiers. In *ACL*.
- Li, J.; Luong, M.-T.; Jurafsky, D.; and Hovy, E. 2015. When are tree structures necessary for deep learning of representations? *arXiv preprint arXiv:1503.00185*.
- Liu, Y.; Sun, C.; Lin, L.; and Wang, X. 2016. Learning natural language inference using bidirectional lstm model and inner-attention. *arXiv preprint arXiv:1605.09090*.
- Luong, M.-T.; Pham, H.; and Manning, C. D. 2015. Effective approaches to attention-based neural machine translation. In *EMNLP*.
- Marelli, M.; Menini, S.; Baroni, M.; Bentivogli, L.; Bernardi, R.; and Zamparelli, R. 2014. A sick cure for the evaluation of compositional distributional semantic models. In *LREC*.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013b. Distributed representations of words and phrases and their compositionality. In *NIPS*.
- Mou, L.; Men, R.; Li, G.; Xu, Y.; Zhang, L.; Yan, R.; and Jin, Z. 2016. Natural language inference by tree-based convolution and heuristic matching. In *ACL*.
- Munkhdalai, T., and Yu, H. 2017a. Neural semantic encoders. In *EACL*.
- Munkhdalai, T., and Yu, H. 2017b. Neural tree indexers for text understanding. In *EACL*.
- Pang, B., and Lee, L. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *ACL*.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- Qian, Q.; Huang, M.; and Zhu, X. 2017. Linguistically regularized lstms for sentiment classification. In *ACL*.
- Rush, A. M.; Chopra, S.; and Weston, J. 2015. A neural attention model for abstractive sentence summarization. In *EMNLP*.
- Seo, M.; Kembhavi, A.; Farhadi, A.; and Hajishirzi, H. 2017. Bidirectional attention flow for machine comprehension. In *ICLR*.
- Shang, L.; Lu, Z.; and Li, H. 2015. Neural responding machine for short-text conversation. In *ACL*.
- Socher, R.; Perelygin, A.; Wu, J. Y.; Chuang, J.; Manning, C. D.; Ng, A. Y.; Potts, C.; et al. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP*.
- Socher, R.; Karpathy, A.; Le, Q. V.; Manning, C. D.; and Ng, A. Y. 2014. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics* 2:207–218.
- Srivastava, N.; Hinton, G. E.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15(1):1929–1958.
- Sukhbaatar, S.; Weston, J.; Fergus, R.; et al. 2015. End-to-end memory networks. In *NIPS*.
- Tai, K. S.; Socher, R.; and Manning, C. D. 2015. Improved semantic representations from tree-structured long short-term memory networks. In *ACL*.
- Teng, Z.; Vo, D.-T.; and Zhang, Y. 2016. Context-sensitive lexicon features for neural sentiment analysis. In *EMNLP*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. In *NIPS*.
- Vendrov, I.; Kiros, R.; Fidler, S.; and Urtasun, R. 2016. Order-embeddings of images and language. In *ICLR*.
- Wiebe, J.; Wilson, T.; and Cardie, C. 2005. Annotating expressions of opinions and emotions in language. *Language resources and evaluation* 39(2):165–210.
- Williams, A.; Nangia, N.; and Bowman, S. R. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*.
- Zeiler, M. D. 2012. Adadelat: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Zhao, H.; Lu, Z.; and Poupard, P. 2015. Self-adaptive hierarchical sentence model. In *IJCAI*.
- Zhao, J.; Zhu, T.; and Lan, M. 2014. Ecnu: One stone two birds: Ensemble of heterogenous measures for semantic relatedness and textual entailment. In *SemEval@ COLING*, 271–277.