UvA-DARE (Digital Academic Repository)

Pre-train, Interact, Fine-tune: A Novel Interaction Representation for Text Classification

Zheng, J.; Cai, F.; Chen, H.; de Rijke, M.

DOI

[10.1016/j.ipm.2020.102215](https://doi.org/10.1016/j.ipm.2020.102215)

Publication date

2020

Document Version

Submitted manuscript

Published in

Information Processing & Management

[Link to publication](#)

Citation for published version (APA):

Zheng, J., Cai, F., Chen, H., & de Rijke, M. (2020). Pre-train, Interact, Fine-tune: A Novel Interaction Representation for Text Classification. *Information Processing & Management*, 57(6), [102215]. <https://doi.org/10.1016/j.ipm.2020.102215>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Pre-train, Interact, Fine-tune: A Novel Interaction Representation for Text Classification

Jianming Zheng^a, Fei Cai^{a,*}, Honghui Chen^a, Maarten de Rijke^b

^a*Science and Technology on Information Systems Engineering Laboratory,
National University of Defense Technology, Hunan, 410073, China*

^b*Informatics Institute, University of Amsterdam, Amsterdam, 1098 XH, The Netherlands*

Abstract

Text representation can aid machines in understanding text. Previous work on text representation often focuses on the so-called forward implication, i.e., preceding words are taken as the context of later words for creating representations, thus ignoring the fact that the semantics of a text segment is a product of the mutual implication of words in the text: later words contribute to the meaning of preceding words. We introduce the concept of interaction and propose a two-perspective interaction representation, that encapsulates a local and a global interaction representation. Here, a *local* interaction representation is one that interacts among words with parent-children relationships on the syntactic trees and a *global* interaction interpretation is one that interacts among all the words in a sentence. We combine the two interaction representations to develop a Hybrid Interaction Representation (HIR).

Inspired by existing feature-based and fine-tuning-based pretrain-finetuning approaches to language models, we integrate the advantages of feature-based and fine-tuning-based methods to propose the Pre-train, Interact, Fine-tune (PIF) architecture.

We evaluate our proposed models on five widely-used datasets for text classification tasks. Our ensemble method, HIR_P, outperforms state-of-the-art base-

*Corresponding Author

Email addresses: ADDRESS (Jianming Zheng), caifei@nudt.edu.cn (Fei Cai), ADDRESS (Honghui Chen), derijke@uva.nl (Maarten de Rijke)

lines with improvements ranging from 2.03% to 3.15% in terms of error rate. In addition, we find that, the improvements of PIF against most state-of-the-art methods is not affected by increasing of the length of the text.

Keywords: Interaction representation, Pre-training, Fine-tuning, Classification

1. Introduction

Text representations map text spans into real-valued vectors or matrices. They have come to play a crucial role in machine understanding of text. Applications include sentiment classification (Tang et al., 2015), question answering (Qin et al., 2017), summarization (Ren et al., 2017), and sentence inference (Parikh et al., 2016).

Previous work on text representation can be categorized into three main types (Xie et al., 2016), i.e., *statistics-based*, *neural-network-based* and *pre-training-based* embeddings. Statistics-based embedding models are estimated based on a statistical indicator, e.g., the frequency of co-occurring words (in bag-of-words models (Joachims, 1998)), the frequency of co-occurring word pairs (in n-gram models (Zhang et al., 2015)), and the weights of words in different documents (the TF-IDF model (Robertson, 2004)). Neural-network-based embedding models mainly rely on a neural network architecture to learn a text representation, based on a hidden layer (Joulin et al., 2017), convolutional neural networks (CNNs) (Kim, 2014) or recurrent neural networks (RNNs) (Liu et al., 2016). Additionally, this type of methods may also consider the syntactic structure to reflect the semantics of text, e.g., recursive neural networks (Socher et al., 2013) and tree-structured long short-term memory networks (Tree-LSTM) (Tai et al., 2015). Pretraining-based embedding models adopt a feature-based (Mikolov et al., 2013; Pennington et al., 2014; McCann et al., 2017; Peters et al., 2018) or fine-tuning strategy (Dai & Le, 2015; Howard & Ruder, 2018; Devlin et al., 2019; Yang et al., 2019) to capture the semantics and syntactic information from a large text corpora.

Figure 1: Overview pipeline of Pre-train Interact Fine-tune

In general, the aforementioned models work well for the task of text classification. (Joulin et al., 2017; Kim, 2014; Zhang et al., 2015; Howard & Ruder, 2018) However, in existing embedding models, the generated process of the vectorized representation of a text usually follows a so-called one-way action. That is to say, representations generated for the preceding text are taken as the context to determine the representations of later texts. Although a bidirectional LSTM considers bidirectional actions, it simply concatenates two one-way actions to get the embeddings. We argue that the semantics as defined in terms of a text representation should be a product of interactions of all source elements (e.g., words or sentences) in the text. Restrictions to one-way actions may result in a partial semantic loss (Saif et al., 2016), causing the poor performances in the downstream applications. We hypothesize that although these interaction relations may be learned by neural networks with enough samples, explicitly modeling such interaction relations can directly make text representation more informative and effective. Furthermore, recent unsupervised representation learning has proven to be effective and promising in the field of natural language processing (McCann et al., 2017; Peters et al., 2018; Howard & Ruder, 2018; Devlin et al., 2019; Yang et al., 2019). So far, these approaches are limited to a single strategy (either feature-based or fine-tuning strategy), which results in a so-called fine-tune error, which may be trapped in the local best.

Thus, as illustrated in Figure 1, we focus on the task of text classification and propose a novel pipeline with the following ingredients:

1. **pre-train** language model on a large text corpus to get the related word embeddings and neural networks parameters;
2. **interact** the word embeddings based on the pre-trained parameters to obtain the interaction representation; and
3. **fine-tune** the classifier with the interaction representation and pre-trained word embeddings as input.

More specifically, in the interaction representation layer, we propose a two-perspective interaction representation using a *Local Interaction Representation* (LIR) and a *Global Interaction Representation* (GIR). The LIR applies an attention mechanism (Bahdanau et al., 2015) inside the syntactic structure of a sentence, e.g., the dependency-based parse trees or constituency-based parse trees, to reflect the local interaction of adjacent words. The GIR employs an attention mechanism with an enumeration-based strategy to represent the interactions of all words in a sentence. After that, we combine LIR and GIR to into a *Hybrid Interaction Representation* (HIR) model to represent both local and global interactions of words in a sentence. For the pretrain-finetuning process, we combine the feature-based and the fine-tuning strategies and propose a *hybrid language model pretrain-finetuning* (HLMPf) approach. HLMPf first follows the fine-tuning strategy to employ the pre-trained embeddings and neural network parameters as the initialization of the interaction representation layer. Then, according to the feature-based strategy, HLMPf applies the pre-trained embeddings as additional features and concatenates the interaction representation in the classifier fine-tuning layer.

For evaluation, we conduct a comprehensive experiment on five publicly available benchmark datasets for the task of text classification. The experimental results show that our proposal with interaction representations and the hybrid pretrain-finetuning strategy outperforms the state-of-the-art baselines for

text classification, with improvements ranging from 2.03% to 3.15% in terms of accuracy.

The main contributions of our work are as follows:

1. We propose a novel pipeline for the task of text classification, i.e., *Pre-train, Interact, Fine-tune* (PIF).
2. To the best of our knowledge, ours is the first attempt to model word interactions for text representation. We introduce a two-perspective interaction representation for text classification, i.e., a Local Interaction Representation (LIR) and a Global Interaction Representation (GIR), which are then combined to generate a Hybrid Interaction Representation (HIR) model.
3. We combine the advantages of two popular language model pretrain-fine-tuning strategies (feature-based and fine-tuning) and propose the hybrid language model pretrain-finetuning (HLMPf).
4. We analyze the effectiveness of our proposal and find that it outperforms the state-of-the-art methods for text classification in terms of accuracy.

2. Related Work

In this section, we briefly summarize the general statistical approaches for text representation in Section 2.1 and the neural-networks-based methods in Section 2.2. We then describe the recent work on language model pre-training for downstream applications in Section 2.3.

2.1. Statistics-based representation

As a word is the most basic unit of semantics, the traditional one-hot representation model converts a word in a vocabulary into a sparse vector with a single high value (i.e., 1) in its position and the others with a low value (i.e., 0). The representation is employed in the Bag-of-Words (BoW) model (Joachims, 1998) to reflect the word frequency. However, the BoW model only symbolizes

the word and cannot reflect the semantic relationship between words. Consequently, the bag-of-means model (Zhang et al., 2015) was proposed to cluster the word embeddings learned by the word2vec model (Mikolov et al., 2013). Furthermore, the bag-of-n-grams (Zhang et al., 2015) was developed to take the n-grams (up to 5-grams) as the vocabulary in the BoW model. In addition, with some extra statistical information, e.g., TF-IDF, a better document representation can be produced (Robertson, 2004). Other text features, e.g., the noun phrases (Lewis, 1992) and the tree kernels (Post & Bergsma, 2013), were incorporated into the model construction.

Clearly, a progressive step has been made in statistical based representation (Bernauer et al., 2018). However, such traditional statistical representation approaches inevitably face the problems of data sparsity and dimensionality, leading to no applications on large-scale corpora. In addition, such approaches are simply built on shallow statistics, and a deeper semantic information of the text has not been well developed.

Instead, our proposal in this paper based on neural networks has the ability to learn a low-dimensional and distributed representation to overcome such problems.

2.2. Neural-based representation

Since Bengio et al. (2000) first employed the neural network architecture to train a language model, considerable attention has been devoted to proposing neural network-related models for text representation. For instance, the FastText model (Joulin et al., 2017) employs one hidden layer to integrate the subword information and obtains satisfactory results. However, this model simply averages all word embeddings and discards the word order. In view of that, Liu et al. (2016) employed the recurrent structure, i.e., RNNs, to consider the word order and to jointly learn text representation across multiple related tasks. Compared to RNNs, CNNs are easier to train and capture the local word-pair information (Kim, 2014; Zhang et al., 2015).

Furthermore, a combination of neural network models are integrated to develop the advantage of each single neural network. For example, Lai et al. (2015) proposed the recurrent convolutional neural networks (RCNN), which adopted the recurrent structure to grasp the context information and employed a max-pooling layer to identify the key components in text. Besides, other document features have been injected into the document modeling. For instance, Zheng et al. (2019) took the hierarchical structure of text into account. He et al. (2018) transformed the document-level knowledge to improve the performance of aspect-level sentiment classification.

Although these approaches have been proved effective in the downstream applications, they completely depend on the structure of network to implicitly represent a document, ignoring the interaction that exists among the source elements in a document, e.g., words or sentences. However, our proposal can model the interaction as the starting point to better reflect the semantic relationship between words in a sentence, which we argue can help improve the performance of downstream tasks, e.g, sentimental classification.

2.3. Language model pre-training-based representation

The language pre-training model has been shown effective for the natural language processing tasks, e.g., question answering (McCann et al., 2017), textual entailment (Peters et al., 2018), semantic role labeling (Devlin et al., 2019) sentimental analysis (Dai & Le, 2015), etc. These pre-training models can be mainly classified into two classes, i.e., feature-based models and fine-tuning models.

The feature-based models generate the pre-trained embeddings from other tasks, where the output can be regarded as the additional features for the current task architecture. For instance, word2vec (Mikolov et al., 2013) and GloVe (Pennington et al., 2014) focus on transforming words into the distributed representations and capturing the syntactics as well as the semantics by pre-training the neural language models on a large text corpora. In addition, McCann et al. (2017) concentrated on the machine translation task to get the contextualize

word vectors (CoVe). Since these word-level models suffer from the word-polysemy, Peters et al. (2018) developed the sequence-level model, i.e., ELMo, to capture the complex word features across different linguistic contexts and then use ELMo to generate the context-sensitive word embeddings.

Different from the feature-based strategy (Mehta & Majumder, 2018), the fine-tuning models first produce the contextual word presentations which have been pre-trained from unlabeled text and fine-tune for a supervised downstream task. For instance, Dai & Le (2015) trained a sequence auto-encoder model on unlabeled text as an initialization of another supervised network. However, this method suffers from overfitting and requires some in-domain knowledge to improve the performance. Consequently, Universal Language Model Fine-tuning (ULMFit) (Howard & Ruder, 2018) was developed, which leveraged the general-domain pre-training and the novel fine-tuning techniques to prevent overfitting. In addition, Devlin et al. (2019) proposed two unsupervised tasks, i.e., masked language model and next sentence prediction, to further improve fine-tuning process. In addition, XLNet (Yang et al., 2019) was proposed to employ the permutation language model to capture the bidirectional context and avoid the pretrain-finetune discrepancy.

Although the language pre-training model based representations have been proposed and proved promising in the NLP tasks, these methods are limited to either feature-based or fine-tune-based strategy. Our proposal combine their respective characteristics to improve the performance of downstream applications.

3. Proposed Models

In this section, we first formally describe how to compute the interaction representation in Section 3.1, which can be divided into three parts, i.e., LIR (see Section 3.1.1), GIR (see Section 3.1.2) and HIR (see Section 3.1.3). And then, we introduce the HLMPf approach in detail (see Section 3.2), which is the combination of the feature-based and fine-tuning strategies.

3.1. Interaction representation

We describe the Local Interaction Representation (LIR) of adjacent words and introduce the Global Interaction Representation (GIR) of all words in a sentence. After that, a Hybrid Interaction Representation (HIR) model is proposed.

3.1.1. Local interaction representation

We introduce an attentive tree LSTM that computes a local representation of words. The idea of an *action* of a word on another word is that the former assigns a semantic weight to the latter.

The experiments we conduct related to LIR are based on constituency-based trees, but we explain the core concepts for both dependency-based and constituency-based trees. Given a dependency-based parse tree, let $C(p)$ denote the set of child words of a parent word x_p . To define the attentive tree LSTM, we introduce hidden states and memory cells h_k and c_k ($k \in \{1, 2, \dots, |C(p)|\}$) for every child word, respectively. As shown in Fig. 2, unlike the Tree-LSTM model in (Tai et al., 2015) that only performs the one-way action (child words $\mapsto$ parent word), LIR also considers an action in the opposite direction, i.e., parent word $\mapsto$ child words.

Let us explain this in detail. In an action parent word $\mapsto$ child words, we regard the parent word x_p as a controller that assigns semantic weights based on the attention mechanism to its child words in a sentence Saraiva et al. (2016). Thus, we first convert the parent word x_p into a hidden representation $\overline{h_p}$ as follows:

$$\overline{h_p} = \tanh(W^{(h)}P_{x_p} + b^{(h)}), \quad (1)$$

where P_{x_p} is the pre-trained word embedding for parent word x_p ; $W^{(h)}$ and $b^{(h)}$ are the weight matrix and the bias term, respectively. Then, we employ a general content-based function (Luong et al., 2015) to connect the parent word and the child words as follows:

$$\alpha_k = \overline{h_p}W_\alpha h_k, \quad (2)$$

Figure 2: Structure of the local interaction representation model. (For simplicity, we write U_p for $(U^{(i)}P_{x_p}, U^{(o)}P_{x_p}, U^{(u)}P_{x_p}, U^{(f)}x_p)$.)

where α_k is the connective representation of $\bar{h}_p$ and the hidden state h_k , and W_α is the connective matrix to be learned. After that, we apply a softmax function on a sequence of connective representations $\{\alpha_1, \alpha_2, \dots, \alpha_{|C(p)|}\}$ to get the weight λ_k as follows:

$$\lambda_k = \frac{\exp(\alpha_k)}{\sum_{i=1}^{|C(p)|} \exp(\alpha_i)}. \quad (3)$$

Finally, we represent the hidden interaction state $\tilde{h}_p$ that relates to all child states of the parent word x_p , i.e.,

$$\tilde{h}_p = \sum_{i \in C(j)} \lambda_i h_i. \quad (4)$$

In the action child words $\mapsto$ parent word in Fig. 2, we use the hidden interaction

state $\tilde{h}_p$ and the parent word x_p as input to the LSTM cell and obtain

$$i_p = \sigma(U^{(i)}x_p + W^{(i)}\tilde{h}_p + b^{(i)}), \quad (5)$$

$$o_p = \sigma(U^{(o)}x_p + W^{(o)}\tilde{h}_p + b^{(o)}), \quad (6)$$

$$u_p = \tanh(U^{(u)}x_p + W^{(u)}\tilde{h}_p + b^{(u)}), \quad (7)$$

$$f_{kp} = \sigma(U^{(f)}x_p + W^{(f)}h_k + b^{(f)}), \quad (8)$$

where i_p , o_p and f_{kp} are the input gate, the output gate and the forget gate, respectively; u_p is the candidate hidden state of x_p . For i_p , o_p , u_p and f_{kp} , we have a corresponding weight matrix of x_p (i.e., $U^{(i)}$, $U^{(o)}$, $U^{(u)}$ and $U^{(f)}$), a weight matrix of $\tilde{h}_p$ (or h_k) (i.e., $W^{(i)}$, $W^{(o)}$, $W^{(u)}$ and $W^{(f)}$), and a bias term (i.e., $b^{(i)}$, $b^{(o)}$, $b^{(u)}$ and $b^{(f)}$). Finally, we can get the memory cell c_p and the hidden state h_p of the parent word x_p as follows:

$$c_p = i_p \odot u_p + \sum_{k=1}^{|C(p)|} f_{kp} \odot c_k, \quad (9)$$

$$h_p = o_p \odot \tanh(c_p), \quad (10)$$

where $\odot$ is element-wise multiplication and c_k is the memory cell of a child word.

Similarly, given a constituency-based tree, let x_l and x_r denote the left child word and the right child word of a parent word x_p . Since the parent word x_p is a non-terminal node (i.e., x_p is a zero vector), we use x_l and x_r as the controller instead of P_{x_p} , respectively. Therefore, following Eq. (2)–(4), we obtain the hidden interaction states $\tilde{h}_l$ and $\tilde{h}_r$ related to x_l and x_r , respectively. We concatenate $\tilde{h}_l$ and $\tilde{h}_r$ to represent the hidden interaction states of the parent word, i.e., $\tilde{h}_p = [\tilde{h}_l; \tilde{h}_r]$. Again, following Eq. (5)–(10), we can get the memory cell c_p and the hidden state h_p for parent word x_p .

At this stage we have represented the local interaction, and each word has been updated by the interaction representation.

3.1.2. Global interaction representation

Unlike LIR, which captures the syntactic relation between words, GIR adopts an enumeration-based strategy to employ an attention mechanism on all words in a sentence.

In detail, after implementing Tree-LSTM on all n words in a sentence, we can have the hidden representations $\{h_1, h_2, \dots, h_n\}$ corresponding to the words $\{x_1, x_2, \dots, x_n\}$. In order to represent the interaction between a word x_g and the other words in a sentence, we regard the word x_g as a controller that can assign semantic weights to other words in $\{x_1, x_2, \dots, x_n\}$ excluding x_g itself. Similarly, we employ a general content-based function to connect the word x_g with other words as follows:

$$\alpha_{gk} = h_g W_\alpha h_k, \quad (11)$$

where α_{gk} is the connective representation of h_g and h_k ($g, k \in (1, 2, \dots, n)$). After that, we can get all connective representations $\{\alpha_{g1}, \alpha_{g2}, \dots, \alpha_{gn}\}$ between the word x_g and other words. Then, we can apply a softmax function on the connective representation sequence to calculate the weight as follows:

$$\lambda_{gk} = \frac{\exp(\alpha_{gk})}{\sum_{i=1}^n \exp(\alpha_{gi})}, \quad (12)$$

where λ_{gk} is the weight of word x_k in $\{x_1, x_2, \dots, x_n\}$ that interacts with word x_g . Finally, we obtain the interaction representation r_g as follows:

$$r_g = \sum_{i=1}^n \lambda_{gi} h_i. \quad (13)$$

By doing so, we enumerate all words in a sentence and can return a sequence of interaction representations as $\{r_1, r_2, \dots, r_n\}$. We then adopt a max-pooling on this sequence to produce the sentence embeddings s by

$$s = \max\{r_1, r_2, \dots, r_n\}. \quad (14)$$

This completes the definition of the global interaction representation. We can train the sentence representation s to update the pre-trained embeddings.

3.1.3. Hybrid interaction representation

In order to capture both local and global interactions between words, we combine LIR and GIR to form a hybrid interaction representation model (HIR) for text representation. HIR first follows the procedure of LIR to produce the hidden state representations $\{h_1, h_2, \dots, h_n\}$ for the corresponding word $\{x_1, x_2, \dots, x_n\}$. Then, HIR employs the process of GIR on these hidden state representations to get the final sentence embeddings s .

Eventually, in the process of class prediction, we apply a softmax classifier on the sentence embeddings s to get a predicted label $\hat{s}$, where $\hat{s} \in \mathcal{Y}$ and $\mathcal{Y}$ is the class label set, i.e.,

$$\hat{s} = \arg \max p(\mathcal{Y} | s), \quad (15)$$

where

$$p(\mathcal{Y} | s) = \text{softmax}(W^{(s)}s + b^{(s)}). \quad (16)$$

Here, $W^{(s)}$ and $b^{(s)}$ are the reshape matrix and the bias term, respectively. For formulating the loss function in HIR, we combine the corresponding loss in LIR and GIR as

$$L = \frac{\gamma}{n} \sum_{i=1}^n \log p(\tilde{w}_i | h_i) - (1 - \gamma) \log p(\tilde{s} | s), \quad (17)$$

where the former loss comes from LIR and the latter from GIR, γ is the trade-off parameter. In addition, h_i is the hidden state and $\tilde{w}_i$ is the true class label of word x_i in LIR; $\tilde{s}$ is the true class label of sentence embeddings s in GIR. In addition, $\tilde{w}_i$ and $\tilde{s}$ can be trained using the dataset.

We have now introduced the main process of our HIR model. Clearly, as shown in Algorithm 1, we first employed bi-lstm process the pre-trained word sequence to build their semantics relations from step 1 to 2. Then, with the help of syntactic parse tool, we can get the parent-child set T . Following the bottom-up traversal algorithm, we show how to model the local interaction

Algorithm 1 Hybrid Interaction Representation

Input: The pre-trained embeddings for each word in a sentence s_t , i.e., $\{P_{x_1}, P_{x_2}, \dots, P_{x_n}\}$; the pre-trained parameters ψ of the neural networks from the language pre-training layer.

Output: The interaction representation for the word sequence, i.e., $\{I_{x_1}, I_{x_2}, \dots, I_{x_n}\}$

- 1: $\vec{h}_t \leftarrow LSTM(x_t, h_{t-1})$, $\overleftarrow{h}_t \leftarrow LSTM(x_t, h_{t-1})$
- 2: $h_t = (\vec{h}_t; \overleftarrow{h}_t)$, $t = 1, 2, \dots, n$
- 3: Get the set of parent word p_i and its child words by syntactic parsing:
 $T = \{p_i, C(p_i)\} \leftarrow \text{syntactic parse}(s)$, $i = 1, \dots, |T|$
- 4: **for** each parent word $\{p_i, C(p_i)\} \in T$ **do**
- 5: $\overline{h}_{p_i} = \tanh(W^{(h)} P_{x_{p_i}} + b^{(h)})$
- 6: **for** each child word in $C(p_i)$: **do**
- 7: Get the connective representation: $\alpha_k = \overline{h}_{p_i} W_\alpha h_k$
- 8: **end for**
- 9: $\tilde{h}_p = \sum_{i \in C(j)} \lambda_i h_i$, where $\lambda_k = \frac{\exp(\alpha_k)}{\sum_{i=1}^{|C(p)|} \exp(\alpha_i)}$
 %% parent word $\mapsto$ child words
- 10: $h_p \leftarrow LSTM(x_p, \tilde{h}_p)$ %% child words $\mapsto$ parent word
- 11: **end for**
 %% This loop for LIR that follows the bottom-up algorithm to traverse the
 syntactic parsing tree.
- 12: **for** each word w_g in sentence s_t **do**
- 13: Word w_g is regarded as the parent word
- 14: **for** each word in sentence s_t excluding word w_g **do**
- 15: $\alpha_{gk} = h_g W_\alpha h_k$
- 16: **end for**
- 17: $r_g = \sum_{i=1}^n \lambda_{gi} h_i$, where $\lambda_{gk} = \frac{\exp(\alpha_{gk})}{\sum_{i=1}^n \exp(\alpha_{gi})}$
- 18: $s = \max\{r_1, r_2, \dots, r_n\}$
- 19: **end for** %% This loop for GIR.
- 20: Optimize the loss function: $L = \frac{\gamma}{n} \sum_{i=1}^n \log p(\tilde{w}_i | h_i) - (1 - \gamma) \log p(\tilde{s} | s)$
- 21: Update $I_{x_i} \leftarrow h_{x_i}$
- 22: **return** $\{I_{x_1}, I_{x_2}, \dots, I_{x_n}\}$

representation between parent word and child words from step 4 to 11. While from step 12 to 19, we show how to compute the global interaction representation between all words. At last, we optimize the loss function to jointly training the process of LIR and GIR. And update the pre-trained embeddings with the interaction representations.

3.2. Hybrid language model pretrain-finetuning

Unsupervised representation learning, as a fundamental tool, has been shown effective in many language processing tasks (McCann et al., 2017; Peters et al., 2018; Howard & Ruder, 2018; Devlin et al., 2019; Yang et al., 2019). Here, we propose the *hybrid language model pretrain-finetuning* (HLMPf) method, which integrates the respective advantages in the PIF pipeline shown in Figure 1. The details of HLMPf are shown in Algorithm 2.

We first follow the BERT (Devlin et al., 2019) model to train the language model pre-training layer. From step 2 to 3, we employ the fine-tuning strategy to fine-tune the interaction representation layer and the language model pre-training layer. After that, we follow the ELMo approach (Peters et al., 2018) to obtain the context-aware word embeddings. From step 5 to 6, we show how to further fine-tune all neural layers following the feature-based strategy.

Specially, since fine-tuning all layers at once will result in catastrophic forgetting, we adopt the gradual unfreezing strategy (Howard & Ruder, 2018) to fine-tune all neural layers.

Algorithm 2 Hybrid Language Model Pretrain-finetuning

Input: The text need to be trained.

Output: The trained parameters ψ of all neural networks; the trained word embeddings $\{W_{x_1}, W_{x_2}, \dots, W_{x_n}\}$.

- 1: Pre-train the masked language model and next sentence prediction tasks to get the pre-trained neural networks and related parameters.
- 2: Add the interaction representation layer to the pre-training layer.
- 3: Following algorithm 1, optimize the loss function to fine-tune the related parameters.

%% the fine-tuning strategy

- 4: Pre-train some supervised tasks to get the context-aware word embeddings, i.e., $\{C_{x_1}, C_{x_2}, \dots, C_{x_n}\}$.
- 5: Add the classifier fine-tuning layer to the former combination layer.
- 6: Use the $\{I_{x_i}; C_{x_i}\}$ as the input of the classifier fine-tuning layer to further fine-tune the related parameters.

%% the feature-based strategy

- 7: **return** ψ and $\{W_{x_1}, W_{x_2}, \dots, W_{x_n}\}$
-

4. Experiments

We start by providing an overview of the text representation model to be discussed in this paper and list the research questions that guide our experiments. Then we describe the task and datasets that we evaluate our proposals on. We conclude the section by specifying the settings of the parameters in our experiments.

4.1. Model summary and research questions

Table 1 list the models to be discussed. Among these models, LSTM, Char-level CNN, LIR, GIR and HIR models are neural based representation and don't experience the pretrain-finetuning process.

Baselines Four state-of-the-art baselines: two neural based representation model (i.e., LSTM (Liu et al., 2016), C-CNN (Zhang et al., 2015)), two language model pre-training based representation model (i.e., CoVe (McCann et al., 2017), ULMFiT (Howard & Ruder, 2018)).

Our proposals Nine flavors of approaches that we introduce in this paper: three interaction representation models (i.e., LIR, GIR and HIR), three interaction representation models in the BERT architecture (i.e., LIR_B, GIR_B, HIR_B) and the Pre-train, Interact, Fine-tune (PIF) architecture (i.e., LIR_P, GIR_P, HIR_P).

To assess the quality of our proposed interaction representation models and the PIF architecture, we consider a text classification task and seek to answer the following questions:

RQ1 Does the interaction representation incorporated in the text representation model help to improve the performance for text classification?

RQ2 Compared with the existing pretrain-finetuning approaches, does our proposed PIF architecture help to improve the model performance for text classification?

Table 1: An overview of models discussed in the paper.

Model	Description	Source	Finetuning
LSTM	A long and short-term memory network (LSTM) based representation model.	(Lai et al., 2015)	×
C-CNN	A CNN based representation model in the character level.	(Zhang et al., 2015)	×
CoVe	A text representation model transferred from the machine translation model.	(McCann et al., 2017)	✓
ULMFiT	A text representation model based on general-domain language model pre-train, target task language model and classifier fine-tune.	(Howard & Ruder, 2018)	✓
LIR	A text representation model based on the local interaction representation.	This paper	×
GIR	A text representation model based on the global interaction representation.	This paper	×
HIR	A text representation model based on the hybrid interaction representation.	This paper	×
LIR_B	A text representation model based on the local interaction representation model in the BERT fine-tuning architecture.	This paper	✓
GIR_B	A text representation model based on the global interaction representation model in the BERT fine-tuning architecture.	This paper	✓
HIR_B	A text representation model based on the hybrid interaction representation model in the BERT fine-tuning architecture.	This paper	✓
LIR_P	A text representation model based on the local interaction representation model in the Pre-train Interact Fine-tune architecture.	This paper	✓
GIR_P	A text representation model based on the global interaction representation model in the Pre-train Interact Fine-tune architecture.	This paper	✓
HIR_P	A text representation model based on the hybrid interaction representation model in the Pre-train Interact Fine-tune architecture.	This paper	✓

RQ3 How does the trade-off parameter between LIR and GIR (as encoded in γ) impact the performance of HIR related model in terms of classification accuracy?

RQ4 Is the performance of our proposal sensitive to the length L of text to be classified?

Table 2: Dataset statistics.

Dataset	IMDb	Yelp	TREC	AG	DBpedia
type	sentiment	sentiment	question	topic	topic
# training documents	25K	560K	5 K	120 K	560K
# text documents	2K	50K	0.5K	7.6K	70K
# classes	2	5	6	4	14

4.2. Datasets

We evaluate our proposal on five publicly available datasets used in different application domains, e.g., sentiment analysis, questions classification and topic classification, which are widely used by the state-of-the-art models for text classification, e.g., CoVe (McCann et al., 2017) and ULMFiT (Howard & Ruder, 2018). Table 2 details the statistics of the datasets. We use accuracy as the evaluation metric to compare the performance of discussed models.

Sentiment analysis Sentiment analysis mainly concentrates on the movie review and shopping review datasets. For example, IMDb dataset proposed by (Maas et al., 2011) is a movie review dataset with binary sentimental labels. While Yelp dataset compiled by (Zhang et al., 2015) is a shopping review dataset that has two versions, i.e., binary and five-class version. We concentrate on the five-class version (Johnson & Zhang, 2017).

Question classification For question classification, Voorhees & Tice (1999) collected open-domain fact-based questions and divided them into broad semantic categories, which has six-class and fifty-class versions. We mainly focus on the small six-class version and hold out 452 examples for validation and leave 5,000 for training, which is similar to (McCann et al., 2017).

Topic classification For topic classification, we evaluate our proposals on the task of news article and ontology classification. We use the AG news corpus collected by Zhang et al. (2015), which has four classes of news with only the titles and description fields. In addition, the DBpedia dataset,

collected by Zhang et al. (2015), is used, which contains the title and abstract of each Wikipedia article with 14 non-overlapping ontology classes. In general, the dataset division is the same as in (Zhang et al., 2015).

4.3. Model configuration and training

For data preprocessing, we split the text into sentences and tokenized each sentence using Stanford’s CoreNLP (Manning et al., 2014). In addition, we discard the words with single characters and other punctuation and convert the upper-case letters to the lower-cases letters. In order to fit in the BERT pre-training, we add a special token for each sentence, e.g., [CLS] and [SEP]. The other data preprocessing follow the same way as (Johnson & Zhang, 2017)

For model configuration, we use the same set of hyper-parameters across all datasets to evaluate the robustness of our proposal. In the process of pre-training, we directly employ the trained $BERT_{base}$ ¹ as our language model pre-training layer for simplicity. As for the feature-based process, we follow the ELMo model² and employ AWD-LSTM (Merity et al., 2018) on the trained BERT layer to get the context-aware word embeddings. For classifier fine-tuning layer, we adopt a softmax classifier and set the size of hidden layer to 100. In addition, we set the dimension of word embeddings and hidden representation in the interaction representation layer to 400 and 200, respectively. We also apply a dropout of 0.4 to layers and 0.05 to the embedding layers.

For the whole training process, we use a batch size of 64, a base learning rate of 0.004 and 0.01 for fine-tuning the interaction representation layer and the classifier fine-tuning layer, respectively. We employ a batch normalization mechanism (Ioffe & Szegedy, 2015) to accelerate the training of the neural networks. Gradient clipping is applied by scaling gradients when the norm may exceed a threshold of 5 (Pascanu et al., 2013). For the fine-tuning process, we adopt the gradual unfreezing strategy (Howard & Ruder, 2018) to fine-tune all neural layers.

¹<https://github.com/google-research/bert>

²<https://github.com/allenai/bilm-tf>

5. Results and Discussion

In Section 5.1, we examine the performance of our proposal incorporated with the interaction representation and the HLMPf on five public datasets, which aims at answering **RQ1** and **RQ2**. Then, in Section 5.2, we analyze the impact of the trade-off parameter γ in HIR related model to answer **RQ3**. Finally, to answer **RQ4**, section Section 5.3 focuses on investigating the impact on the text classification by varying the text length.

5.1. Performance comparison

5.1.1. Performances about the interaction representation

To answer **RQ1**, we first compare the performance of the basic interaction representation based models (i.e., LIR, GIR and HIR) with the baselines and present the results in Table 3

Table 3: Error rate (%) about the interaction representation on different datasets. (The results of the best baseline and the best performer in each column are underlined and boldfaced, respectively. Results marked with * are re-printed from (Zhang et al., 2015; McCann et al., 2017; Howard & Ruder, 2018; Zhou et al., 2016). The rest are obtained by our own implementation. Statistical significance of pairwise differences (the best proposed model vs. the best neural-network-based baseline) are determined by a t -test ($\blacktriangle/\blacktriangledown$ for $\alpha = .01$)

Datasets	IMDb	Yelp	TREC	AG	DBpedia
LSTM	8.72	41.83*	7.66	13.94*	1.45*
C-CNN	7.36	37.95*	6.48	9.51*	1.55*
CoVe	8.2*	—	4.2*	—	—
ULMFiT	<u>4.6*</u>	<u>29.98*</u>	<u>3.6*</u>	<u>5.01*</u>	<u>0.80*</u>
LIR	6.86	35.58	5.76	7.83	1.31
GIR	6.92	35.46	5.87	8.20	1.37
HIR	6.73 $\blacktriangle$	34.18 $\blacktriangle$	5.44 $\blacktriangle$	7.53 $\blacktriangle$	1.24 $\blacktriangle$

As to the baselines, we present two types of representation models, i.e., the neural-network based model (LSTM and C-CNN) and the pretrain-finetuning based model (CoVe and ULMFiT). For the neural-network based model, C-CNN achieves a better performance than LSTM. While in the pretrain-finetuning based model, ULMFiT is obviously the better one. Interestingly, comparing these two types of models, we can find that the representation models with

pretrain-finetuning process have super advantages in terms of reducing error rate. Specially, with regard to C-CNN, ULMFiT reduces the error dramatically by 37.5%, 26.6%, 44.4%, 89.8% and 48.4% on the corresponding datasets (IMDb, Yelp, TREC, AG and DBpedia in order, which is the same in the following text). This may be due to the fact that the pre-training on a large text corpora can capture the deep syntactic and semantic information, which cannot be realized by only training on the neural networks.

Similarly, our proposals only with the interaction representation, i.e., LIR, GIR and HIR, cannot beat the state-of-the-art pretrain-finetuning based model, i.e., ULMFiT. But for the neural-network based baselines, our proposals can achieve better performance in terms of error rate. In particular, HIR is the best performing model among our proposals, which shows an improvement against the best neural-network based baseline, i.e., C-CNN, resulting in 8.6%, 9.9%, 16.0%, 26.3% and 20% reduction in terms of error rate on the respective datasets. LIR and GIR, following HIR, can outperform C-CNN on all datasets. The aforementioned findings indicate that compared with the traditional neural-network based models, modeling the interaction process explicitly can better capture the semantics relation between source elements in the text and generate more meaningful text representation. Especially for HIR, by representing the local and global interaction between words, it is more effective to improve the performance of the downstream applications.

5.1.2. Performances about the pretrain-finetuning

In section 5.1.1, the effectiveness of the pretrain-finetuning based and the interaction-related models have been proven. However, the basic interaction representation based models cannot beat the state-of-the-art pretrain-finetuning based model, i.e., ULMFiT. Hence, we incorporate them with the popular pretrain-finetuning architecture (i.e., BERT) and our PIF architecture to get the corresponding models (i.e., LIR_B , GIR_B , HIR_B and LIR_P , GIR_P , HIR_P), respectively. To answer **RQ2**, we compare the performance of these proposed models with ULMFiT and present their experimental results in Table 4.

Table 4: Error rate (%) about the pretrain-finetuning process on different datasets. (The results of the best baseline and the best performer in each column are underlined and boldfaced, respectively. Results marked with * are re-printed from (Howard & Ruder, 2018). The rest are obtained by our own implementation. Statistical significance of pairwise differences (the best proposed model vs. the best neural-network-based baseline) are determined by a t -test ($\blacktriangle/\blacktriangledown$ for $\alpha = .01$)

Datasets	IMDb	Yelp	TREC	AG	DBpedia
ULMFiT	<u>4.6*</u>	<u>29.98*</u>	<u>3.6*</u>	<u>5.01*</u>	<u>0.80*</u>
LIR _B	4.58	28.42	3.54	4.93	0.81
GIR _B	4.69	28.84	3.55	5.03	0.84
HIR _B	4.24	28.31	3.37	4.88	0.78
LIR _P	4.25	27.33	3.40	4.92	0.80
GIR _P	4.31	27.66	3.48	4.94	0.81
HIR _P	4.04$\blacktriangle$	27.07$\blacktriangle$	3.33$\blacktriangle$	4.85$\blacktriangle$	0.77$\blacktriangle$

Clearly, as shown in Table 4, our basic interaction-related models incorporated with the pretrain-finetuning process generally outperform the state-of-the-art model, i.e., ULMFiT, except for some cases, e.g., the LIR_B on DBpedia, GIR_B on AG and DBpedia, GIR_P on DBpedia. This findings again prove that our basic interaction representation models have the promising perspectives under the pretrain-finetuning architecture. With regard to the BERT architecture, our interaction-related models present the similar accuracy distribution to the basic interaction representation models in Table 3. HIR_B is the best performer using the BERT architecture, followed by LIR_B and GIR_B. Specially, for each dataset, HIR_B shows an obvious improvement of 7.9%, 5.6%, 6.4%, 2.6% and 2.5% against ULMFiT, respectively. While LIR_B, except on the DBpedia, also gains a minor improvement of 0.4%, 5.2%, 1.7%, 1.6% against ULMFiT, respectively. GIR_B, a bit worser than LIR_B, beats the ULMFiT on 3 out of 5 datasets.

The similar findings can also be found in the PIF architecture. In particular, HIR_P achieves the best performance not only in the PIF architecture but among all discussed models. Compared with the baseline ULMFiT, HIR_P gains substantial improvements of 12.1%, 9.7%, 7.5%, 3.2%, 3.8% in terms of error rate on respective datasets. In addition, LIR_P wins the comparisons against

ULMFiT, resulting in 7.6%, 8.8%, 5.6%, 1.8% improvements on the respective datasets and an equal performance on DBpedia. While GIR_P defeats the ULMFiT model on 4 out of 5 datasets.

Furthermore, comparing the same type of interaction models with different architectures (e.g., type LIR: LIR, LIR_B , LIR_P), we can find that there exists a unchanged ranking order of performance on each dataset, i.e., $\text{LIR}_P > \text{LIR}_B > \text{LIR}$, $\text{GIR}_P > \text{GIR}_B > \text{GIR}$, $\text{HIR}_P > \text{HIR}_B > \text{HIR}$. This ranking order demonstrates that our proposed PIF architecture that combines the feature-based and fine-tuning based strategies is the most effective architecture, followed by the fine-tuning based strategy, BERT. While the neural-network based models are worse than the former kinds of models.

5.2. Parameters analysis

Next we turn to **RQ3** and conduct a parameter sensitivity analysis of our HIR related models, i.e., HIR, HIR_B and HIR_P . Clearly, as shown in Table 3 and Table 4, for different datasets, the same model has varied error rates on different orders of magnitude, e.g., HIR on IMDB and Yelp (6.73 *vs* 34.18). To better present the γ effect of the same model on different datasets, we introduce an evaluation metric, Relative Error Rate (RER), which is defined as, given a dataset, the relative improvement ratio of the lowest error rate with regard to the others with different γ s. In addition, we examine the performances of these three models in terms of RER by gradually changing the parameters γ from 0 to 1 with an interval 0.1. We plot the RER results of HIR, HIR_B and HIR_P in Figure 3a, Figure 3b and Figure 3c, respectively.

As shown in Figure 3a, HIR achieves the lowest error rate when $\gamma = 0.5$ on all datasets (except $\gamma = 0.6$ for Yelp dataset), which is 0 in the figure. In addition, the RER of HIR on each dataset decreases consistently when γ varies from 0 to 0.5 (0.6 for Yelp); after that, the RER metric goes up when γ changes from 0.5 (0.6 for Yelp) to 1. The similar phenomena can be found in Figure 3b and Figure 3c. HIR_B and HIR_P both achieve the lowest error rate when $\gamma = 0.5$. In addition, the RER of these two models on each dataset first keeps

a stable decrease to the lowest point 0 and then increases stably until $\gamma = 1$.

Interestingly, comparing the curve gradient on both sides of $\gamma = 0.5$, we can find that the gradient of the left side is steeper than that of the right side, which indicates that GIR can result in the increase of error rate more easily than LIR. Furthermore, comparing the same model on different datasets, we can find that the change ranges of RER on IMDB, DBpedia and TREC, is greater than that on Yelp and AG. The phenomena may be due to the differences of statistical characteristics among these datasets, which require further experiments to find potential reasons.

Curiously, we also want to find whether the relation $\text{HIR}_P \lessdot \text{HIR}_B \lessdot \text{HIR}$ can always keep unchanged when the trade-off parameter γ increases from 0 to 1. Due to the text space, we only select the dataset Yelp as the analytical object, which has the highest error rate among these datasets. We plot the experimental results in Figure 4. Clearly, as Figure 4 shows, we can find that the performance of HIR_P is the lowest in terms of error rate, followed by HIR_B , and the highest is HIR , when γ increases from 0 to 1. This result is consistent with the previous finding $\text{HIR}_P > \text{HIR}_B > \text{HIR}$, i.e., the effectiveness of our PIF architecture. On the other hand, it indicates that the effectiveness of our PIF architecture is not sensitive to the trade-off parameter γ .

5.3. Impact of the text length

To answer **RQ4**, we manually group the text according to the text length L , e.g., 0–100, 100–200, ..., 900–1000, >1000. We compare the performance of interaction representation related models, e.g., LIR, GIR, HIR, HIR_B and HIR_P , under different settings of text length. We plot this experimental results in Figure 5

Clearly, as shown in Figure 5, we can find the relation $\text{LIR} > \text{GIR} > \text{HIR} > \text{HIR}_B > \text{HIR}_P$ unchanged when text length increases. This phenomenon is consistent with the findings in Section 5.1.1, which indicates the effectiveness of interaction representation and PIF architecture is not affected by the text length.

(a) HIR performance on each dataset.

(b) HIR_B performance on each dataset.

(c) HIR_P performance on each dataset.

Figure 3: Effect on performance of the HIR related models in terms of RER by changing the trade-off parameter γ , tested on all datasets.

Figure 4: Effect on the performance of HIR related models in terms of error rate by changing the trade-off parameter γ , tested on the Yelp dataset.

Figure 5: Effect on the performance of interaction related models in terms of error rate with varied text length, tested on Yelp dataset.

Interestingly, as the text length increases, the performances of all discussed models decrease first to reach the lowest error rate at the point of group 100–200, and then keep a constant increase. This finding may be explained by the fact that the longer the text, the richer the information it provides, which results in targeting the class label of text more easily, i.e., the decrease of error rate in the earlier stage. But as the text length grows, the structure and semantics of text become more complex and variable, the proposed models find it harder to get the exact representation.

6. Conclusion and Future Work

In this paper, we focus on the task of text classification and propose a novel pipeline, the PIF architecture, which incorporates the respective advantages from feature based and fine-tuning based strategies in the language model pretrain-finetuning process. We also introduce the concept of interaction representation and propose a two-perspective interaction representation for sentence embeddings, i.e., a local interaction representation (LIR) and a global interaction representation (GIR). We combine these two representations to produce a hybrid interaction representation model, i.e., HIR.

We evaluate these models on five widely-used datasets for text classification. Our experimental results shows that: (1) compared with the traditional neural-network based models, our basic interaction-related models can help boost the performance for text classification in terms of error rate. (2) our proposed PIF architecture is more effective to help improve the text classification than the existing feature-based as well as the fine-tuning based strategies. Specially, HIR_P model present the best performance on each dataset. (3) the effectiveness of interaction representation and the PIF architecture is not affected by the text length.

As to future work, we plan to evaluate our models for other tasks so as to verify the robustness of the interaction representation models. In addition, the existing fine-tuning approach is too general. We want to investigate some task-sensitive fine-tuning methods to better improve the performance.

7. Acknowledgements

This work was partially supported by the National Natural Science Foundation of China under No. 61702526, the Defense Industrial Technology Development Program under No. JCKY2017204B064, the National Advanced Research Project under No. 6141B0801010b, Ahold Delhaize, the Association of Universities in the Netherlands (VSNU), and the Innovation Center for Artificial Intel-

ligence (ICAI). All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Bengio, Y., Ducharme, R., & Vincent, P. (2000). A neural probabilistic language model. In *NIPS* (pp. 932–938).
- Bernauer, L., Han, E. J., & Sohn, S. Y. (2018). Term discrimination for text search tasks derived from negative binomial distribution. *Inf. Process. Manage.*, 54, 370–379.
- Dai, A. M., & Le, Q. V. (2015). Semi-supervised sequence learning. In *NIPS* (pp. 3079–3087).
- Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT* (pp. 4171–4186).
- He, R., Lee, W. S., Ng, H. T., & Dahlmeier, D. (2018). Exploiting document knowledge for aspect-level sentiment classification. In *ACL* (pp. 579–585).
- Howard, J., & Ruder, S. (2018). Universal language model fine-tuning for text classification. In *ACL* (pp. 328–339).
- Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML* (pp. 448–456).
- Joachims, T. (1998). Text categorization with support vector machines: Learning with many relevant features. In *ECML* (pp. 137–142).
- Johnson, R., & Zhang, T. (2017). Deep pyramid convolutional neural networks for text categorization. In *ACL* (pp. 562–570).

- Joulin, A., Grave, E., Bojanowski, P., & Mikolov, T. (2017). Bag of tricks for efficient text classification. In *EACL* (pp. 427–431).
- Kenter, T., Borisov, A., & de Rijke, M. (2016). Siamese CBOW: Optimizing word embeddings for sentence representations. In *ACL* (pp. 941–951).
- Kim, Y. (2014). Convolutional neural networks for sentence classification. In *EMNLP* (pp. 1746–1751).
- Lai, S., Xu, L., Liu, K., & Jun, Z. (2015). Recurrent convolutional neural networks for text classification. In *AAAI* (pp. 2267–2273).
- Lewis, D. D. (1992). An evaluation of phrasal and clustered representations on a text categorization task. In *SIGIR* (pp. 37–50).
- Liu, P., Qiu, X., & Huang, X. (2016). Recurrent neural network for text classification with multi-task learning. In *IJCAI* (pp. 2873–2879).
- Luong, M. T., Pham, H., & Manning, C. D. (2015). Effective approaches to attention-based neural machine translation. In *EMNLP* (pp. 1412–1421).
- Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., & Potts, C. (2011). Learning word vectors for sentiment analysis. In *ACL* (pp. 142–150).
- Manning, C. D., Surdeanu, M., Bauer, J., Finkel, J. R., Bethard, S., & McClosky, D. (2014). The stanford corenlp natural language processing toolkit. In *ACL* (pp. 55–60).
- McCann, B., Bradbury, J., Xiong, C., & Socher, R. (2017). Learned in translation: Contextualized word vectors. In *NIPS* (pp. 6297–6308).
- Mehta, P., & Majumder, P. (2018). Effective aggregation of various summarization techniques. *Inf. Process. Manage.*, 54, 145–158.
- Merity, S., Keskar, N. S., & Socher, R. (2018). Regularizing and optimizing LSTM language models. In *ICLR*.

- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. In *ICLR*.
- Parikh, A. P., Täckström, O., Das, D., & Uszkoreit, J. (2016). A decomposable attention model for natural language inference. In *EMNLP* (pp. 2249–2255).
- Pascanu, R., Mikolov, T., & Bengio, Y. (2013). On the difficulty of training recurrent neural networks. In *ICML* (pp. 1310–1318).
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL* (pp. 1532–1543).
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. In *NAACL-HLT* (pp. 2227–2237).
- Post, M., & Bergsma, S. (2013). Explicit and implicit syntactic features for text classification. In *ACL* (pp. 866–872).
- Qin, C., Qinmin, H., Jimmy Xiangji, H., Liang, H., & Weijie, A. (2017). Enhancing recurrent neural networks with positional attention for question answering. In *SIGIR* (pp. 993–996).
- Ranzato, M. A., Boureau, Y. L., & Lecun, Y. (2007). Sparse feature learning for deep belief networks. In *NIPS* (pp. 1185–1192).
- Ren, P., Chen, Z., Ren, Z., & de Rijke, M. (2017). Leveraging contextual sentence relations for extractive summarization using a neural attention model. In *SIGIR* (pp. 95–104).
- Robertson, S. (2004). Understanding inverse document frequency: on theoretical arguments for IDF. *Journal of Documentation*, 60, 503–520.

- Saif, H., He, Y., Fernández, M., & Alani, H. (2016). Contextual semantics for sentiment analysis of twitter. *Inf. Process. Manage.*, 52, 5–19.
- Saraiva, P. C., Cavalcanti, J. M. B., de Moura, E. S., Gonçalves, M. A., & da Silva Torres, R. (2016). A multimodal query expansion based on genetic programming for visually-oriented e-commerce applications. *Inf. Process. Manage.*, 52, 783–800.
- Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C., Ng, A., & Potts, C. (2013). Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP* (pp. 1631–1642).
- Tai, K. S., Socher, R., & Manning, C. D. (2015). Improved semantic representations from tree-structured long short-term memory networks. In *ACL* (pp. 1556–1566).
- Tang, D., Qin, B., & Liu, T. (2015). Document modeling with gated recurrent neural network for sentiment classification. In *EMNLP* (pp. 1422–1432).
- Voorhees, E. M., & Tice, D. M. (1999). The TREC-8 question answering track evaluation. In *TREC* (p. 82).
- Xie, H., Li, X., Wang, T., Lau, R. Y. K., Wong, T., Chen, L., Wang, F. L., & Li, Q. (2016). Incorporating sentiment into tag-based user profiles and resource profiles for personalized search in folksonomy. *Inf. Process. Manage.*, 52, 61–72.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J. G., Salakhutdinov, R., & Le, Q. V. (2019). Xlnet: Generalized autoregressive pretraining for language understanding. *CoRR*, *abs/1906.08237*.
- Zhang, X., Zhao, J., & Lecun, Y. (2015). Character-level convolutional networks for text classification. In *NIPS* (pp. 649–657).
- Zheng, J., Cai, F., Chen, W., Feng, C., & Chen, H. (2019). Hierarchical neural representation for document classification. *Cognitive Computation*, 11, 317–327.

Zhou, P., Qi, Z., Zheng, S., Xu, J., Bao, H., & Xu, B. (2016). Text classification improved by integrating bidirectional LSTM with two-dimensional max pooling. In *COLING* (pp. 3485–3495).