Neurocomputing 471 (2022) 61-69

Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

Co-attention network with label embedding for text classification

Minqian Liu^{a,1}, Lizhao Liu^{b,1}, Junyi Cao^b, Qing Du^{b,*}

^a School of Computer Science and Engineering, South China University of Technology, China ^b School of Software Engineering, South China University of Technology, China

ARTICLE INFO

Article history: Received 1 June 2020 Revised 14 July 2021 Accepted 29 October 2021 Available online 4 November 2021 Communicated by Zidong Wang

Keywords: Text-label co-attention Label embedding Text classification Natural language processing Deep learning

ABSTRACT

Most existing methods for text classification focus on extracting a highly discriminative text representation, which, however, is typically computationally inefficient. To alleviate this issue, label embedding frameworks are proposed to adopt the label-to-text attention that directly uses label information to construct the text representation for more efficient text classification. Although these label embedding methods have achieved promising results, there is still much space for exploring how to use the label information more effectively. In this paper, we seek to exploit the label information by further constructing the text-attended label representation with text-to-label attention. To this end, we propose a Coattention Network with Label Embedding (CNLE) that jointly encodes the text and labels into their mutually attended representations. In this way, the model is able to attend to the relevant parts of both. Experiments show that our approach achieves competitive results compared with previous state-ofthe-art methods on 7 multi-class classification benchmarks and 2 multi-label classification benchmarks. © 2021 Elsevier B.V. All rights reserved.

1. Introduction

Text classification is one of the fundamental tasks in Natural Language Processing (NLP), which has been widely applied in sentiment analysis [1], question answering [2], and so on. Given an input text, the goal of text classification is to assign it with one or multiple labels from the predefined label set.

Existing methods for text classification mainly focus on encoding the input text and conducting classification with the encoded text representation. Traditionally, Brown et al. [3] and Wallach [4] apply feature engineering techniques and use hand-crafted features to represent the text. In recent years, deep neural networks have been applied to encode the text, including Convolutional Neural Networks (CNNs) [5,6] and Recurrent Neural Networks (RNNs) [7,8]. By employing deep models, previous methods are able to extract a more discriminative text representation to improve the performance of text classification. However, these deep models typically consume a large number of computational resources.

To alleviate this issue, Wang et al. [9] and Du et al. [10] adopt label embedding frameworks for more efficient text classification. They notice that the impact of label information on learning the text representation is indirect [9]. Hence, they propose to embed all the labels in the predefined label set into the same space as the word embedding and then adopt label-to-text attention to produce a label-attended text representation. By incorporating the label information in a more direct way, these label embedding approaches achieve promising results with fewer parameters and less computation. Thus, a natural question to ask is how to use the information from labels more effectively.

When exploiting the labels, we argue that it is helpful to leverage the information from the text. Intuitively, in a fine-grained sentiment classification task with five classes, humans can first easily exclude some obviously wrong answers (*e.g.*, very negative and negative) by briefly glancing at the text. Then, they can pay more attention to the remaining categories (*e.g.*, neutral, positive, and very positive). Motivated by this, we further incorporate the textto-label attention into text classification. In this way, we can focus on finding the label(s) that more closely match(es) the text during encoding the label embedding.

In this paper, we propose a text-label co-attention mechanism to obtain the text-attended label representation and the labelattended text representation. Inspired by the co-attention adopted in question answering [11–13], we introduce a Co-attention Network with Label Embedding (CNLE), where we jointly encode the text and label into their mutually attended representations for more effective text classification. Our model consists of two modules, namely the Text-Label Co-attentive Encoder (TLCE) and the Adaptive Label Decoder (ALD). In particular, the TLCE produces the mutually attended representations to focus on the relevant





Corresponding author.
 E-mail addresses: csmqliu@mail.scut.edu.cn (M. Liu), selizhaoliu@mail.scut.edu.cn (L. Liu), sejaycao@mail.scut.edu.cn (J. Cao), duqing@scut.edu.cn (Q. Du).

¹ Equal contribution.

parts of both, while the ALD leverages these representations to generate the prediction. Our method is more compatible with multi-label classification compared with other approaches. First, multi-label classification typically has more classes. We use multi-head co-attention in TLCE to let the model focus more on the labels that are more likely to be correct. Second, the labels in multi-label classification are often correlated. Previous work [14] shows that capturing the correlation among labels is critical for multi-label classification. Thus, we adopt a generative approach for ALD to capture the correlation among the labels, which not only models the correlation among labels but also serves as a unified framework for both multi-class classification and multi-label classification. The extensive experiments on several benchmark datasets demonstrate the effectiveness of the proposed method.

The main contributions of this paper are summarized as follows.

- For more effective text classification, we propose a text-label co-attention mechanism to obtain both the label-attended text representation and text-attended label representation. This allows the model to focus on the relevant parts of both the text and labels to benefit text classification.
- We devise a Co-attention Network with Label Embedding (CNLE) consisting of a Text-Label Co-attentive Encoder (TLCE) and an Adaptive Label Decoder (ALD). The TLCE aims to obtain the mutually attended representations of text and labels, and the ALD leverages these representations to tackle both the multi-class and the multi-label classification without modifying the model.
- We evaluate our method on 7 benchmarks of multi-class classification and 2 benchmarks of multi-label classification. The experiments show that our method achieves competitive results compared with previous state-of-the-art methods.

2. Related work

Text Classification. Traditional methods for text classification [15–17] use feature engineering techniques such as N-grams [3] and bag-of-words (BoW) [4] as the feature extractor, and then apply Support Vector Machine (SVM) [17] as the classifier. In recent years, models that are based on neural networks [18-26], such as Convolutional Neural Networks (CNNs) [27-29,5,6] and Recurrent Neural Networks (RNNs) [30,7,8,31], have been widely applied in text classification to extract more informative features from the text. Moreover, Yang et al. [14] propose a Sequence Generation Model (SGM) to apply a sequence-to-sequence generative approach for the multi-label classification. During the decoding stage, SGM is able to capture the correlations among the labels since it produces the next label given the labels predicted at previous steps. In this work, we also adopt the encoder-decoder architecture to generate the target label(s). In this way, our model is able to cope with both the multi-class classification and multilabel classification without modifying the model. Compared with traditional approaches (e.g., SVM), large pre-trained models such as BERT [21] and XLNet [22] recently achieved substantial improvement in terms of performance thanks to their powerful encoding ability. However, these methods typically rely on a highly discriminative text representation that may require abundant computation resources to obtain. In this work, in addition to exploiting the text representation, we also seek to leverage the label information for more effective and efficient text classification.

Label Embedding. Label embedding for text classification has been studied in multitask learning [32] to tackle the potential loss of label information. Wang et al. [9] and Du et al. [10] propose to view text classification as a word-label matching problem. Wang et al. [9] notice that the use of label information only occurs in learning the classifier on top of the model and its impact on learning the text representation is indirect. To address this issue, they directly incorporate label information from the bottom of the model by introducing the label embedding with a label-to-text attention mechanism. These methods show that the label embedding is informative for the downstream classification task. With a simpler architecture with fewer parameters, they still yield very promising results. In this work, we further incorporate the textto-label attention and produce the text-label co-attended representation for text classification. Compared with the previous label embedding methods, our method leverages the feedback from the text representation to encode more information into the labels.

Co-attention Mechanism. Co-attention [33–35] is widely applied in the multi-modal learning between images and language (e.g., visual question answering). Lu et al. [33] propose a co-attention model that jointly reasons about the image and question information. Recently. Lu et al. [35] incorporate a co-attention transformer architecture to learn the joint representations of the image content and natural language. Moreover, the co-attention has also been applied in question answering [11–13]. Seo et al. [11] adopt context-to-query attention and query-to-context attention to reduce the information loss. Xiong et al. [12] adopt the coattention to attend to important parts in the document and the question. Co-attention [36,33,12,13,37,35] in the previous methods typically fuse the information from two separated sources (e.g., image-language or question-text) to produce the attended source representations for each source respectively. Different from the classic co-attention mechanism, we modify the popular selfattention [38] into a co-attention between the source and target (i.e., the text sequence and label sequence) to jointly produce the mutually attended representations for both the text and labels. In this way, our model is able to focus on the relevant parts of both to benefit text classification.

3. Co-attention network with label embedding

We propose a Co-attention Network with Label Embedding (CNLE) to jointly encode the input text sequence and the label sequence, and then use their co-attended representations to generate the target label(s) for text classification. The overall scheme of CNLE is shown in Fig. 1.

Notations. Throughout the paper, we use the following notations. We denote the text sequence as $\mathbf{x} = [x_1, ..., x_m]$ and the label set as $\mathscr{C} = \{l_1, ..., l_c\}$, where *m* denotes the number of words in \mathbf{x} and *c* denotes the number of classes (*e.g.*, for binary classification, c = 2). In this work, we consider text classification as a generation problem, where we generate one label at each decoding timestep. Specifically, given the text \mathbf{x} , we aim to produce the probabilities $\hat{\mathbf{y}}_t$ for all classes to predict the target label at *t*-th decoding step. For the multi-class classification, only a single step of decoding is required, while the multi-label classification requires multiple decoding steps. The notations used in the paper are summarized in Table 1.

3.1. General scheme

In this paper, we focus on incorporating text-to-label attention based on previous label embedding methods [9,10] to further exploit the information from labels. To this end, we propose a text-label co-attentive architecture to obtain the text-attended label representation and label-attended text representation during the encoding process and then leverage both representations to generate a target label sequence. The intuition behind our proposed co-attention mechanism is that the information of text and labels can be mutually fused with each other. In this way, the



Fig. 1. The overall architecture of the Co-attention Network with Label Embedding (CNLE). The illustration shows the encoding process and one timestep of the decoding process on the multi-label classification task. The Text-Label Co-attentive Encoder (TLCE) takes a text sequence and a label sequence as inputs to produce their mutually attended representations. Then, the Adaptive Label Decoder (ALD) uses both the text and label representation to compute the probability for each class in the current timestep. Best viewed in color.

Summary of the notations used in this paper.

GeneralmThe number of words in the textScalar c The number of classesScalar d Hidden dimension (other than transformer layer)Scalar d Label set c κ Text sequence $m \times d_{emb}$ \mathbf{I} Label set expresentation $m \times d_{emb}$ \mathbf{Z}_{ql} Label-attended text representation $m \times d$ \mathbf{I}_{qc} Decoder function $ f_{enc}$ Encoder function $ f_{dec}$ Decoder function $ f_{dec}$ Decoder function of transformer layerScalar k_{mb} Hidden dimension of transformer layerScalar k_{emb} Label embeddings $c \times d_{emb}$ k_{mc} Text encoding from contextual encoding layer $m \times d$ k_{mc} Label encoding from contextual encoding layer $c \times d_{m}$ k_{arc} Label-attended text encoding $m \times d_h$ k_{arc} Fuse tencoding from contextual encoding layer $c \times d_{m}$ k_{arc} Label-attended label encoding $m \times d_h$ k_{arc} Label-attended label encoding $m \times d_h$	Symbol Description	Shape
cThe number of classesScalardHidden dimension (other than transformer layer)Scalar \mathscr{C} Label set C \mathscr{K} Text sequence $m \times d_{emb}$ ILabel sequence $c \times d_{emb}$ Z_{nl} Label-attended text representation $m \times d$ Z_{nlx} Text-attended label representation $c \times d$ f_{enc} Decoder function $ f_{dec}$ Decoder function $ f_{dec}$ Decoder function function headsScalar p The number of attention headsScalar q_{mb} Embeddings $m \times d_{emb}$ L_{emb} Label embeddings $c \times d$ L_{emc} Label encoding from contextual encoding layer $c \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d_h$ L_{att} Text-attended label encoding $c \times d_h$ L_{att} Text-attended label encoding $c \times d_h$	<i>m</i> The number of words in the text	Scalar
dHidden dimension (other than transformer layer)Scalar \mathscr{C} Label setc \mathbf{x} Text sequence $m \times d_{emb}$ \mathbf{z}_{nl} Label sequence $c \times d_{emb}$ \mathbf{z}_{nl} Label-attended text representation $m \times d$ \mathbf{z}_{px} Text-attended label representation $c \times d$ f_{enc} Encoder function $ f_{dec}$ Decoder function $ f_{dec}$ Decoder function fransformer layerScalar d_{h} Hidden dimension of transformer layerScalar \mathbf{z}_{emb} Word embeddings $m \times d_{emb}$ \mathbf{z}_{emb} Label embeddings $m \times d_{emb}$ \mathbf{z}_{proj} Projected word embeddings $m \times d$ \mathbf{k}_{proj} Projected label embeddings $m \times d$ \mathbf{k}_{mc} Text cooling from contextual encoding layer $m \times d$ \mathbf{k}_{mc} Label encoding from contextual encoding layer $c \times d$ \mathbf{k}_{mc} Ext encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{mc} Ext encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{mc} Ext encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{mc} Ext encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{mc} Ext encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{mc} Ext encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{mc} Ext encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{mc} Ext encoding from contextual encoding layer </td <td>c The number of classes</td> <td>Scalar</td>	c The number of classes	Scalar
\mathscr{C} Label set c \mathbf{x} Text sequence $m \times d_{emb}$ \mathbf{l} Label sequence $c \times d_{emb}$ $\mathbf{z}_{s(l)}$ Label-attended text representation $m \times d$ $\mathbf{z}_{s(l)}$ Text-attended label representation $c \times d$ \mathbf{f}_{enc} Encoder function- f_{dec} Decoder function- f_{dec} Decoder function- f_{dec} The number of attention headsScalar d_h Hidden dimension of transformer layerScalar χ_{emb} Word embeddings $m \times d_{emb}$ \mathbf{L}_{emb} Label embeddings $c \times d$ \mathbf{k}_{proj} Projected word embeddings $c \times d$ \mathbf{k}_{rnc} Text encoding from contextual encoding layer $m \times d$ \mathbf{k}_{att} Label encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{att} Label-attended text encoding $m \times d_h$ \mathbf{k}_{arco} Fext-attended label encoding layer $m \times d_h$ \mathbf{k}_{arco} Fext-attended label encoding layer $m \times d_h$ \mathbf{k}_{arco} Fuenc det text encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{arco} Fuenc text encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{arco} Fuenc text encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{arco} Fuenc text encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{arco} Fuenc text encoding from contextual encoding layer $m \times d_h$ \mathbf{k}_{arco} Fuenc text encoding from contextual encoding	d Hidden dimension (other than tran	Scalar
xText sequence $m \times d_{emb}$ ILabel sequence $C \times d_{emb}$ $z_{x l}$ Label-attended text representation $C \times d_{emb}$ $z_{x l}$ Text-attended label representation $c \times d$ f_{enc} Encoder function $ f_{dec}$ Decoder function $-$ TLCE d_{emb} Embedding dimensionScalar d_h Hidden dimension of transformer layerScalar p The number of attention headsScalar χ_{emb} Word embeddings $m \times d_{emb}$ L_{emb} Label embeddings $m \times d$ χ_{proj} Projected word embeddings $m \times d$ χ_{proj} Projected label embeddings $m \times d$ χ_{enc} Text encoding from contextual encoding layer $m \times d$ χ_{att} Label-attended text encoding $m \times d_h$ χ_{att} Label-attended label encoding layer $c \times d$ χ_{att} Label-attended laxt encoding from co-attention layer $m \times d_h$ χ_{att} Label-attended laxt encoding from co-attention layer $m \times d_h$ χ_{att} Label-attended laxt encoding from co-attention layer $m \times d_h$ χ_{att} Label-attended laxt encoding from co-attention layer $m \times d_h$ χ_{att} Label-attended laxt encoding from co-attention layer $m \times d_h$ χ_{att} Label-attended laxt encoding from co-attention layer $m \times d_h$ χ_{att} Label encoding from co-attention layer $m \times d_h$	C Label set	С
ILabel sequence $c \times d_{emb}$ $\mathbf{z}_{k l}$ Label-attended text representation $m \times d$ $\mathbf{z}_{ x}$ Text-attended label representation $c \times d$ f_{enc} Encoder function- f_{dec} Decoder function-TLCE d_{emb} Embedding dimensionScalar d_h Hidden dimension of transformer layerScalar p The number of attention headsScalar χ_{emb} Word embeddings $m \times d_{emb}$ L_{emb} Label embeddings $m \times d$ \mathbf{x}_{proj} Projected word embeddings $m \times d$ \mathbf{x}_{enc} Text encoding from contextual encoding layer $m \times d$ \mathbf{x}_{att} Label encoding from contextual encoding layer $c \times d$ \mathbf{x}_{att} Label-attended text encoding from contextual encoding layer $m \times d_h$ \mathbf{x}_{att} Text-attended label encoding from contextual encoding layer $m \times d_h$ \mathbf{x}_{att} Text-attended taket necoding layer $m \times d_h$ \mathbf{x}_{att} Text-attended taket necoding layer $m \times d_h$ \mathbf{x}_{att} Label-attended text encoding layer $m \times d_h$ \mathbf{x}_{att} Text-attended label encoding from contextual encoding layer $m \times d_h$ \mathbf{x}_{att} Text-attended label encoding from contextual encoding layer $m \times d_h$ \mathbf{x}_{att} Text-attended label encoding from contextual encoding layer $m \times d_h$ \mathbf{x}_{att} Text-attended label encoding from contextual encoding layer $m \times d_h$ \mathbf{x}_{att} Text-attende	x Text sequence	$m imes d_{emb}$
$z_{k l}$ Label-attended text representation $m \times d$ $z_{l x}$ Text-attended label representation $c \times d$ f_{enc} Encoder function $ f_{dc}$ Decoder function $ f_{dc}$ Decoder function $-$ TLCE d_{emb} Embedding dimensionScalar d_h Hidden dimension of transformer layerScalar p The number of attention headsScalar k_{emb} Word embeddings $m \times d_{emb}$ L_{emb} Label embeddings $m \times d_{emb}$ L_{emb} Label embeddings $m \times d$ L_{proj} Projected word embeddings $m \times d$ K_{enc} Text encoding from contextual encoding layer $m \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ L_{att} Text-attended tabel encoding $m \times d_h$ L_{att} Text-attended tabel encoding from co-attention layer $m \times d_h$	Label sequence	$c imes d_{emb}$
$z_{l x}$ Text-attended label representation $c \times d$ f_{enc} Encoder function- f_{dec} Decoder function-TLCE d_{emb} Embedding dimensionScalar d_h Hidden dimension of transformer layerScalar p The number of attention headsScalar k_{emb} Word embeddings $m \times d_{emb}$ L_{emb} Label embeddings $c \times d_{emb}$ L_{proj} Projected word embeddings $m \times d$ L_{proj} Projected label embeddings $c \times d$ L_{enc} Label encoding from contextual encoding layer $m \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ L_{att} Text-attended tabel encoding $m \times d_h$ L_{att} Text-attended label encoding from co-attention layer $m \times d_h$	$\mathbf{z}_{x l}$ Label-attended text representation	m imes d
f_{enc} f_{dec} Encoder function Decoder function-TLCE d_{emb} Embedding dimensionScalar d_h Hidden dimension of transformer layerScalar p The number of attention headsScalar p The number of attention headsScalar k_{emb} Word embeddings $m \times d_{emb}$ L_{emb} Label embeddings $c \times d_{emb}$ L_{proj} Projected word embeddings $m \times d$ L_{proj} Projected label embeddings $c \times d$ L_{enc} Label encoding from contextual encoding layer $m \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ L_{att} Text-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding $m \times d_h$ L_{att} Text-attended label encoding from co-attention layer $m \times d$	$\mathbf{z}_{l x}$ Text-attended label representation	c imes d
f_{dec} Decoder function-TLCE d_{emb} Embedding dimensionScalar d_h Hidden dimension of transformer layerScalar p The number of attention headsScalar p The number of attention headsScalar χ_{emb} Word embeddings $m \times d_{emb}$ L_{emb} Label embeddings $c \times d_{emb}$ χ_{proj} Projected word embeddings $m \times d$ L_{proj} Projected label embeddings $c \times d$ χ_{enc} Text encoding from contextual encoding layer $m \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ χ_{att} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding from co-attention layer $m \times d$	<i>f</i> _{enc} Encoder function	-
TLCE d_{emb} Embedding dimensionScalar d_h Hidden dimension of transformer layerScalar p The number of attention headsScalar p The number of attention headsScalar X_{emb} Word embeddings $m \times d_{emb}$ L_{emb} Label embeddings $c \times d_{emb}$ X_{proj} Projected word embeddings $m \times d$ L_{proj} Projected label embeddings $c \times d$ L_{proj} Projected label embedding layer $m \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ L_{att} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding from co-attention layer $m \times d_h$ L_{att} Text-attended label encoding from co-attention layer $m \times d$	<i>f</i> _{dec} Decoder function	-
d_h Hidden dimension of transformer layerScalar p The number of attention headsScalar p The number of attention headsScalar X_{emb} Word embeddings $m \times d_{emb}$ L_{emb} Label embeddings $c \times d_{emb}$ X_{proj} Projected word embeddings $m \times d$ L_{proj} Projected label embeddings $c \times d$ L_{proj} Projected label embedding layer $m \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ L_{att} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding $c \times d_h$ L_{att} Text-attended label encoding from co-attention layer $m \times d_h$	<i>d</i> _{emb} Embedding dimension	Scalar
pThe number of attention headsScalar X_{emb} Word embeddings $m \times d_{emb}$ L_{emb} Label embeddings $c \times d_{emb}$ X_{proj} Projected word embeddings $m \times d$ L_{proj} Projected label embeddings $c \times d$ L_{proj} Projected label embeddings $c \times d$ L_{enc} Text encoding from contextual encoding layer $m \times d$ L_{enc} Label-attended text encoding layer $c \times d$ X_{att} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding $m \times d_h$ L_{att} Text-attended label encoding from co-attention layer $m \times d$	d _h Hidden dimension of transformer l	Scalar
X_{emb} Word embeddings $m \times d_{emb}$ L_{emb} Label embeddings $c \times d_{emb}$ X_{proj} Projected word embeddings $m \times d$ L_{proj} Projected label embeddings $c \times d$ L_{proj} Projected label embeddings $c \times d$ L_{proj} Projected label embeddings $c \times d$ L_{enc} Text encoding from contextual encoding layer $c \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ X_{att} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding $c \times d_h$ X_{enc} Fused text encoding from co-attention layer $m \times d$	p The number of attention heads	Scalar
L_{emb} Label embeddings $c \times d_{emb}$ X_{proj} Projected word embeddings $m \times d$ L_{proj} Projected label embeddings $c \times d$ X_{enc} Text encoding from contextual encoding layer $m \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ X_{att} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding $c \times d_h$ X_{arc} Eused text encoding from co-attention layer $m \times d_h$	X _{emb} Word embeddings	$m imes d_{emb}$
X_{proj} Projected word embeddings $m \times d$ L_{proj} Projected label embeddings $c \times d$ X_{enc} Text encoding from contextual encoding layer $m \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ X_{art} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding $c \times d_h$ X_{arc} Eused text encoding from co-attention layer $m \times d_h$	Label embeddings	$c imes d_{emb}$
L_{proj} Projected label embeddings $c \times d$ X_{enc} Text encoding from contextual encoding layer $m \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ X_{att} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding $c \times d_h$ X_{att} Even encoding from co-attention layer $m \times d_h$	X _{proj} Projected word embeddings	m imes d
X_{enc} Text encoding from contextual encoding layer $m \times d$ L_{enc} Label encoding from contextual encoding layer $c \times d$ X_{att} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding $c \times d_h$ X_{enc} Fused text encoding from co-attention layer $m \times d$	L _{proj} Projected label embeddings	c imes d
L_{enc} Label encoding from contextual encoding layer $c \times d$ X_{att} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding $c \times d_h$ X_{enco} Fused text encoding from co-attention layer $m \times d$	X _{enc} Text encoding from contextual enc	m imes d
X_{att} Label-attended text encoding $m \times d_h$ L_{att} Text-attended label encoding $c \times d_h$ X_{barrow} Fused text encoding from co-attention layer $m \times d_h$	Label encoding from contextual en-	c imes d
L_{att} Text-attended label encoding $c \times d_h$ X_{6irea} Fused text encoding from co-attention layer $m \times d$	X _{att} Label-attended text encoding	$m imes d_h$
X _{funce} Fused text encoding from co-attention layer $m \times d$	L _{att} Text-attended label encoding	$c imes d_h$
$-\gamma_{\rm usc}$ rated the choosing non-co-dimensionary $n_{\rm eff}$	X _{fuse} Fused text encoding from co-attent	m imes d
\mathbf{L}_{fuse} Fused label encoding from co-attention layer $c imes d$	L _{fuse} Fused label encoding from co-atter	c imes d
\mathbf{X}_{fin} Final text encoding from final encoding layer $m \times d$	X _{fin} Final text encoding from final enco	m imes d
\mathbf{L}_{fin} Final label encoding from final encoding layer $c \times d$	L _{fin} Final label encoding from final enco	c imes d
ALD T The total number of decoding steps Scalar	<i>T</i> The total number of decoding step:	Scalar
$\hat{\mathbf{y}}$ Predicted probability sequence $T \times c$	ŷ Predicted probability sequence	T imes c
\mathbf{e}_t The embedding of the ground true label at the <i>t</i> -th decoding step $d \times 1$	\mathbf{e}_t The embedding of the ground true	coding step $d \times 1$
\mathbf{r}_t Recurrent context state at the <i>t</i> -th decoding step $d \times 1$	r _t Recurrent context state at the <i>t</i> -th	$d \times 1$
\mathbf{h}_t Hidden state at the <i>t</i> -th decoding step $d \times 1$	h _t Hidden state at the <i>t</i> -th decoding s	d imes 1
\mathbf{c}_t Cell state at the <i>t</i> -th decoding step $d \times 1$	c _t Cell state at the t-th decoding step	d imes 1
\mathbf{a}_t Attention weights at the <i>t</i> -th decoding step $m \times 1$	a _t Attention weights at the <i>t</i> -th decor	$m \times 1$
$\hat{\mathbf{y}}_t$ Predicted probability at the <i>t</i> -th decoding step $1 \times c$	$\hat{\mathbf{y}}_t$ Predicted probability at the <i>t</i> -th de	1 imes c
1[·] Indicator function –	1[·] Indicator function	-
\mathscr{L} Objective function –	\mathscr{L} Objective function	-

model is able to focus on the relevant parts in the text sequence and label sequence.

Formally, we consider a text sequence \mathbf{x} containing m word tokens in the document and a label sequence \mathbf{I} containing c label tokens in the predefined label set \mathscr{C} . It is worth noting that the

label sequence is converted from the label set \mathscr{C} (see Section 4.1.3 for details) and it is identical for each sample during both training and inference. We aim to incorporate the label information and text information to obtain a label-attended text representation \mathbf{z}_{xl} and a text-attended label representation \mathbf{z}_{tlx} :

M. Liu, L. Liu, J. Cao et al.

$$\mathbf{z}_{\mathbf{x}|l}, \mathbf{z}_{l|\mathbf{x}} = f_{enc}(\mathbf{x}, \mathbf{l}), \tag{1}$$

where f_{enc} is any mapping function to be learned. After acquiring these two representations, we generate a sequence of probabilities $\hat{\mathbf{y}}$ for each class with a decoder:

$$\hat{\mathbf{y}} = f_{dec}(\mathbf{z}_{x|l}, \mathbf{z}_{l|x}), \tag{2}$$

where f_{dec} is any mapping function to be learned. In this way, our decoder is able to leverage the mutually attended representations to facilitate the label prediction.

In the following, we will present the architecture design of our model. Particularly, we detail the Text-Label Co-attentive Encoder (TLCE) in Section 3.2 and introduce the Adaptive Label Decoder (ALD) in Section 3.3.

3.2. Text-label co-attentive encoder

In this subsection, we introduce our proposed Text-Label Coattentive Encoder (TLCE) in detail. The TLCE aims to jointly encode the text sequence and the label sequence into mutually attended text and label representations. Specifically, the TLCE consists of a contextual encoding layer, a co-attention layer, and a final encoding layer.

3.2.1. Contextual encoding layer

To encode the word tokens and label tokens into meaningful encodings, we adopt a Bidirectional Long Short Term Memory (BiLSTM) [39] to capture the contextual cues from the text and the correlation (which we consider as a special case of contextual information) among the labels.

Given an input word sequence $\mathbf{x} \in \mathbb{R}^m$ and a label sequence $\mathbf{l} \in \mathbb{R}^c$, we first map them into word embeddings $\mathbf{X}_{emb} \in \mathbb{R}^{m \times d_{emb}}$ and label embeddings $\mathbf{L}_{emb} \in \mathbb{R}^{c \times d_{emb}}$ respectively. The word embeddings are initialized with the pre-trained word vectors [40,41] while the label embeddings are randomly initialized. For computation efficiency, we apply two independent linear projection layer to project \mathbf{X}_{emb} and \mathbf{L}_{emb} into the more compact embeddings with a smaller dimension, *i.e.*, $\mathbf{X}_{proj} \in \mathbb{R}^{m \times d}$ and $\mathbf{L}_{proj} \in \mathbb{R}^{c \times d}$, respectively. Here, $d < d_{emb}$.

To capture the contextual information in the text sequence and the correlation in the label sequence, we apply a BiLSTM on the projected word embeddings $\mathbf{X}_{proj} \in \mathbb{R}^{m \times d}$ and the projected label embeddings $\mathbf{L}_{proj} \in \mathbb{R}^{c \times d}$ respectively:

$$\begin{aligned} \mathbf{X}_{enc} &= \text{BiLSTM}(\mathbf{X}_{proj}), \\ \mathbf{L}_{enc} &= \text{BiLSTM}(\mathbf{L}_{proj}), \end{aligned} \tag{3}$$

where $\mathbf{X}_{enc} \in \mathbb{R}^{m \times d}$ is the text encoding and $\mathbf{L}_{enc} \in \mathbb{R}^{c \times d}$ is the label encoding. To reduce the number of parameters, we share the weights of this BiLSTM.

3.2.2. Co-attention layer

To acquire the text-attended label representation and the labelattended text representation, we modify the widely applied selfattention module in the Transformer [38] into a co-attention layer. Instead of applying the self-attention, we employ a co-attention between the text encoding and the label encoding to attend to each other. Intuitively, the label-attended text representation can help the model to focus more on the relevant words for the classification task, while the text-attended label representation is able to emphasize the labels that match the text better.

For convenience, we first recap the scaled dot product attention and multi-head attention in Transformer [38]. The scaled dot product attention is defined as follows:

Attention(
$$\mathbf{Q}, \mathbf{K}, \mathbf{V}$$
) = Softmax($\frac{\mathbf{Q}\mathbf{K}^{l}}{\sqrt{d_{h}}}$) \mathbf{V} , (4)

where $\mathbf{Q} \in \mathbb{R}^{q \times d_k}$, $\mathbf{K} \in \mathbb{R}^{k \times d_k}$, and $\mathbf{V} \in \mathbb{R}^{k \times d_v}$. And the multi-head attention is defined as:

$$MultiHead(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = Softmax([\mathbf{H}_1; \dots; \mathbf{H}_p])\mathbf{W}^0,$$

where $\mathbf{H}_i = Attention(\mathbf{Q}\mathbf{W}_i^Q, \mathbf{K}\mathbf{W}_i^K, \mathbf{V}\mathbf{W}_i^V),$ (5)

and the projecting parameters are $\mathbf{W}_{i}^{Q} \in \mathbb{R}^{d_{k} \times d_{p}}, \mathbf{W}_{i}^{K} \in \mathbb{R}^{d_{k} \times d_{p}},$ $\mathbf{W}_{i}^{V} \in \mathbb{R}^{d_{v} \times d_{p}},$ and $\mathbf{W}_{i}^{O} \in \mathbb{R}^{pd_{p} \times d_{h}}$. We use $d_{k} = d_{v} = d_{p}$ and $d_{p} = d_{h}/p$ is the dimension for each head, where d is the interval model dimension and p is the number of heads. We use the [·] to denote the concatenate operation throughout the paper. The original multi-head self-attention [38] only considers the text modal, *i.e.*, the matrices \mathbf{Q}, \mathbf{K} and \mathbf{V} are text encoding. To mutually refine the text and label, we convert the self-attention into a co-attention by feeding both the text encoding \mathbf{X}_{enc} and the label encoding \mathbf{L}_{enc} into the multi-head attention module:

$$\begin{aligned} \mathbf{X}_{att} &= \text{MultiHead}_X(\mathbf{X}_{enc}, \mathbf{L}_{enc}, \mathbf{L}_{enc}), \\ \mathbf{L}_{att} &= \text{MultiHead}_L(\mathbf{L}_{enc}, \mathbf{X}_{enc}, \mathbf{X}_{enc}), \end{aligned} \tag{6}$$

where $\mathbf{X}_{att} \in \mathbb{R}^{m \times d_h}$ and $\mathbf{L}_{att} \in \mathbb{R}^{c \times d_h}$ are the label-attended text representation and the text-attended label representation respectively.

Similar to the Transformer [38] architecture, we also apply residual connections and two independent Feed Forward Networks (FFN) following Layer Normalization (LN) [42] to obtain their fused encodings $\mathbf{X}_{fuse} \in \mathbb{R}^{m \times d}$ and $\mathbf{L}_{fuse} \in \mathbb{R}^{c \times d}$:

$$\begin{aligned} \mathbf{X}_{fuse} &= \mathrm{LN}_{X}(\mathrm{FFN}_{X}(\mathbf{X}_{att}) + \mathbf{X}_{enc}), \\ \mathbf{L}_{fuse} &= \mathrm{LN}_{L}(\mathrm{FFN}_{L}(\mathbf{L}_{att}) + \mathbf{L}_{enc}). \end{aligned} \tag{7}$$

The FFN projects the input into dimension d and further propagates the mutually attended information, while the residual connection fuses the attended representations with their original encodings.

3.2.3. Final encoding layer

To further leverage the contextual information of the labelattended text encoding and the correlation of the text-attended label encoding, we apply a final encoding layer on top of the coattention layer.

In the final encoding process, two independent BiLSTMs are applied to propagate the mutually attended information in the fused text encoding and the fused label encoding, respectively. One BiLSTM encodes the fused text encoding X_{fuse} into the text final representation X_{fin} :

$$\mathbf{X}_{fin} = \text{BiLSTM}_{X}(\mathbf{X}_{fuse}),\tag{8}$$

where $\mathbf{X}_{fin} \in \mathbb{R}^{m \times d}$. Note that the hidden state $\mathbf{h} \in \mathbb{R}^{d \times 1}$ and the cell state $\mathbf{c} \in \mathbb{R}^{d \times 1}$ inside the BiLSTM_{*X*} are preserved for the subsequent decoding process. They are used to initialize the hidden state and cell state in LSTM Decoder respectively (see Section 3.3.1 for details).

Another BiLSTM encodes the fused label encoding to yield the final representation of the label sequence $\mathbf{L}_{fin} \in \mathbb{R}^{c \times d}$:

$$\mathbf{L}_{fin} = \text{BiLSTM}_L(\mathbf{L}_{fuse}). \tag{9}$$

3.3. Adaptive label decoder

In this subsection, we elaborate on our Adaptive Label Decoder (ALD) in detail. The ALD leverages both the text representation and the label representation to generate the label(s). For each timestep,

the decoding process of ALD is separated into two steps: 1) obtain the hidden state, the cell state, and the recurrent context state with an LSTM Decoder; 2) compute a probability for each class via an adaptive classifier.

3.3.1. LSTM decoder

To cope with both the multi-class and the multi-label classification without any modification to the model, we adopt an LSTM Decoder to generate the label(s), which is able to further model the correlations between labels [14].

Specifically, we adopt a standard LSTMCell with attention [43] to implement the LSTM Decoder. During training, we first look up the label embedding $\mathbf{e}_{t-1} \in \mathbb{R}^{d \times 1}$ for the ground true label in the (t-1)-th decoding step. Note that during inference, we look up the embedding for the predicted label instead. Then, the LSTMCell takes the label embedding \mathbf{e}_{t-1} , the recurrent context state $\mathbf{r}_{t-1} \in \mathbb{R}^{d \times 1}$, the hidden state $\mathbf{h}_{t-1} \in \mathbb{R}^{d \times 1}$, and the cell state $\mathbf{c}_{t-1} \in \mathbb{R}^{d \times 1}$ of previous timestep as inputs. It outputs the hidden state \mathbf{h}_t and the cell state \mathbf{c}_t for the current timestep t as follows:

$$\mathbf{h}_{t}, \mathbf{c}_{t} = \text{LSTMCell}([\mathbf{e}_{t-1}; \mathbf{r}_{t-1}], \mathbf{h}_{t-1}, \mathbf{c}_{t-1}). \tag{10}$$

Here, we initialize the \mathbf{h}_0 and \mathbf{c}_0 with \mathbf{h} and \mathbf{c} from the encoding process respectively. Both the \mathbf{e}_0 and \mathbf{r}_0 are initialized with zero vectors.

After we obtain the hidden state \mathbf{h}_t , we compute the attention weights $\mathbf{a}_t \in \mathbb{R}^{m \times 1}$ to help the decoder to focus on the text information relevant to the timestep *t*:

$$\mathbf{a}_t = \text{Softmax}(\mathbf{X}_{fin}\mathbf{W}_1\mathbf{h}_t), \tag{11}$$

where $\mathbf{W}_1 \in \mathbb{R}^{d \times d}$ is a trainable matrix. Then, we obtain the recurrent context state \mathbf{r}_t for timestep *t* as follows:

$$\mathbf{r}_t = \operatorname{Tanh}(\mathbf{W}_2[\mathbf{X}_{fin}^T \mathbf{a}_t; \mathbf{h}_t]), \tag{12}$$

where $\mathbf{W}_2 \in \mathbb{R}^{d \times 2d}$ is a trainable matrix. The recurrent context state \mathbf{r}_t will be used in the next timestep to compute the hidden state \mathbf{h}_{t+1} as in Eq. 10.

Then, the hidden state \mathbf{h}_t will be passed to the adaptive classifier for subsequent procedures.

3.3.2. Adaptive classifier

To leverage the informative text-attended label representation, we propose an adaptive classifier. Comparing with the linear projection layer used in most of the existing methods [5,30,44,31,45], the adaptive classifier is able to attentively leverage the text-attended label representation to directly output a probability for each class.

Given the hidden state \mathbf{h}_t of timestep t, the adaptive classifier takes the final label encoding \mathbf{L}_{fin} and the hidden state \mathbf{h}_t as inputs and produce the probabilities $\hat{\mathbf{y}}_t$ at timestep t:

$$\hat{\mathbf{y}}_t = \text{Sotfmax}(\mathbf{L}_{fin}\mathbf{W}_3\mathbf{h}_t), \tag{13}$$

where $\mathbf{W}_3 \in \mathbb{R}^{d \times d}$ is a trainable matrix and $\hat{\mathbf{y}}_t \in \mathbb{R}^{1 \times c}$ is a vector that contains the probability for each class at timestep *t*. In this way, we directly incorporate the text-attended label representation to facilitate the classification process.

After generating the probabilities for all the *T* timesteps, we compute the objective function \mathscr{L} to optimize our model as follows:

$$\mathscr{L} = -\sum_{t=1}^{T} \sum_{i \in \mathscr{C}} \mathbb{1}[\mathbf{y}_{i} = i] \log(\hat{\mathbf{y}}_{i}),$$
(14)

where \mathbf{y}_t is the ground true label at timestep t. $\mathbb{1}[\cdot]$ is the indicator function, where $\mathbb{1}[\mathbb{A}] = \mathbb{1}$ if A is true and $\mathbb{1}[\mathbb{A}] = \mathbb{0}$ if A is false. Note

that Eq. 14 can be applied in both the multi-label classification and multi-class classification task.

4. Experiments

4.1. Experimental settings

4.1.1. Datasets

We evaluate the models on 7 multi-class and 2 multi-label datasets. The statistics summary of these datasets is presented in Table 2.

For multi-class classification, we follow [28] and use 7 standard benchmark datasets with contents specified as follows. Note that we consider the binary classification as a special case of the multi-class classification.

- Yelp Full Review (Yelp F.): The Yelp Review dataset is obtained from the Yelp Dataset Challenge in 2015. This task aims to predict the sentiment polarity labels, *i.e.*, very negative, negative, neutral, positive, and very positive.
- Yelp Polarity Review (Yelp P.): This dataset contains the same set of the review texts in the Yelp Full Review, but only considers two sentiment polarities, *i.e.*, negative and positive.
- Amazon Full Review (Amz. F.): The Amazon dataset is obtained from the Stanford Network Analysis Project (SNAP) [46]. Similar to the Yelp Review dataset, the Amazon Full Review dataset also has five sentiment stars ranging from 1 to 5.
- Amazon Polarity Review (Amz. P.): Similar to the Yelp Polarity Review, this dataset also considers the star 1 and 2 as negative, and the star 4 and 5 as positive.
- AG News (AG): The AG News dataset is obtained from the Internet news articles [47]. It has four categories for topic classification: world, entertainment, sports, and business.
- DBPedia (DBP.): The DBPedia is constructed from Wikipedia [48]. It has 14 non-overlapping ontology classes from DBpedia 2014.
- Yahoo! Answers (Yah. A.): The Yahoo! Answers dataset is constructed from Yahoo! Answers Comprehensive Questions and Answers version 1.0 dataset in the Yahoo! Webscope program.

For multi-label classification, we evaluate our method on 2 popular datasets, namely AAPD [14] and Reuters-21578 [49].

- Arxiv Academic Paper Dataset (AAPD): The AAPD dataset is built by collecting the abstract and the corresponding subjects of 55,840 papers from the computer science academic website.¹
- Reuters-21578: Reuters is a dataset collected from Reuters news articles with 90 categories.

4.1.2. Evaluation metrics

We use accuracy as the metric for multi-class classification. For multi-label classification, we adopt the micro-averaged F1 score [50], which comes from the class-weighted harmonic mean of precision and recall. Denote the true positive, false positive, and false negative counts of the *i*-th class as TP_i , FP_i , and FN_i respectively. The micro-F1 score is computed as follows:

micro - F₁ =
$$\frac{\sum_{i=1}^{c} 2TP_i}{\sum_{i=1}^{c} 2TP_i + FP_i + FN_i}$$
. (15)

4.1.3. Implementation details

To convert the predefined label set into our input label sequence, we randomly choose the order of labels (which is fixed

¹ https://arxiv.org/

Statistics summary of multi-class and multi-label classification datasets. #Training and #Test refer to the number of samples in the training set and test set respectively.

Туре	Dataset	#Classes	#Training	#Test	Task
Multi-class	Yelp Full	5	650 k	50 k	sentiment
	Yelp Polarity	2	560 k	38 k	sentiment
	Amazon Full	5	3,000 k	650 k	sentiment
	Amazon Polarity	2	3,600 k	400 k	sentiment
	AG News	4	120 k	76 k	topic
	DBPedia	14	560 k	70 k	ontology
	Yahoo! Answers	10	1,400 k	60 k	topic
Multi-label	AAPD	54	53,840	1,000	topic
	Reuters-21578	90	9,598	3,299	topic

once chosen) and concatenate all the labels to obtain the label sequence. The label sequence contains all the labels in the label set and it is identical for every sample in the dataset. For word embeddings in the text sequence, we concatenate the 300dimension Glove [40] and the 100-dimension character n-gram embedding [41] to produce 400-dimension embeddings, *i.e.*, $d_{emb} = 400$, which are frozen during training. We initialize Out-Of-Vocabulary (OOV) words and label embeddings from a uniform distribution with the range [-0.01, 0.01] (also with 400 dimensions). The internal model dimension d is set to 256. For the BiLSTM, we set the output dimension to d/2 and concatenate the outputs of two directions to obtain the final outputs. The hidden dimension in the co-attention layer, *i.e.*, d_h is set to 150 and the number of attention heads p is set to 3. We use a greedy search decoding strategy for the decoder since labels are usually short in length. Our model is trained with an Adam [51] optimizer along with the warmup schedule as in [38]. β_1 , β_2 and ϵ in the Adam optimizer are set to 0.9, 0.98 and 10^{-9} respectively. The number of steps for warmup is set to 800 and the batch size is 32. During training, we apply gradient clipping that re-scales the gradient's norm within the interval [0.0, 1.0) to avoid unstable optimization. For the datasets that contain less than 20 M samples, we apply a sub-sample strategy to augment the dataset to obtain 20 M samples in total. We stop the training process if the validation metric does not change for 5 consecutive epochs. All experiments are conducted on a single Titan-XP GPU.

4.2. Quantitative results

In this part, we compare the performance of different methods on 9 benchmarks in total. We present the results of multi-class classification in Table 3 and the results of multi-label classification in Table 4. Due to space constraints, we use abbreviations for the names of the datasets that are introduced in Section 4.1 in Table 3.

Multi-class Classification. As shown in Table 3, our CNLE generally achieves higher accuracy compared with the previous baseline methods on the 6 multi-class datasets, which validates the effectiveness of our method. Label embedding methods (*i.e.*, LEAM [9], EXAM [10]) achieve promising performance compared with the methods with more complex architecture, such as VDCNN [5] and Encoder1-Encoder2 [45] (abbreviated as E1-E2 in Table 3). Different from them, our model uses both the text-attended label presentation and label-attended text presentation. The results show that our model achieves considerable improvement in performance comparing with previous state-of-the-art methods. This agrees well with our intuition that further exploiting the information from labels improves the text classification task more effectively.

Multi-label Classification. We also compare our model with previous state-of-the-arts methods in multi-label classification. From Table 4, our model outperforms the prior works, often by a large margin. Significantly, our model surpasses previous methods by 2.9% on the Reuters dataset. The improvement indicates that our model is able to better cope with multi-label classification.

4.3. Qualitative results

We visualize both the label-to-text attention and the text-tolabel attention in the co-attention layer in Fig. 2. As depicted in the figure, the important parts for classification, such as *the new york sirloin cooked well, but it was pretty tough*, have higher attention scores on both sides. It shows that our co-attention layer captures the classification-relevant parts in both the text sequence and the label sequence, which verifies the effectiveness of the proposed text-label co-attention mechanism. Note that some of the attention heads in the multi-head attention module turn out to be redundant and this issue has been studied in some literature [52,53]. Thus, the corresponding attention scores on the redundant heads are not discriminative, making them look similar. Nevertheless, some attention heads (as shown in Fig. 2 with colors in orange, light blue, and blue) are discriminative enough to facilitate the classification.

5. Further experiments

5.1. Ablation studies

To further validate the effectiveness of our CNLE, we conduct two ablation studies on both multi-class and multi-label benchmarks. The results are shown in Table 5 and Table 6 respectively. In the first ablation study (CNLE w/o Label Embedding), we do not use the label embedding (*i.e.*, label sequence) at all. In another word, only the text sequence is fed into the model as input, which yields a model similar to most existing methods. In the second ablation (CNLE w/o Text-to-label Attention), we retain the label embedding and label-to-text attention. However, we remove the text-to-label attention such that the label embedding is not attended by the text. Ablation results indicate that with the textattended label representation, our method generally yields a considerable improvement in terms of performance. This shows the effectiveness of our text-label co-attentive architecture.

5.2. Comparison with large pre-trained models

Recently, several large pre-trained models such as BERT [21] and XLNet [22] use large scale unsupervised corpus and obtain a powerful language representation ability. These models indeed achieve better performance than ours, but we argue that they are computationally inefficient and might not be practical in some real-life settings. In Table 7, we present the comparison between our CNLE and two large pre-trained models in terms of performance and the number of parameters. We clearly see that they have more than 100 times parameters comparing with CNLE, but only with less than 0.03% performance gain. This is a big disadvan-

Test accuracy on 7 multi-class benchmarks. The best results are in bold.

Model	Yelp P.	Yelp F.	Amz. P.	Amz. F.	AG	DBP.	Yah. A.	
Bags-of-words [28]	92.24	57.99	90.40	54.64	88.81	96.61	68.89	
LSTM [28]	94.74	58.17	93.90	59.43	86.06	98.55	70.84	
VDCNN [5]	95.72	64.26	95.69	63.00	91.27	98.71	73.43	
D. LSTM [31]	92.60	59.60	-	-	92.10	98.70	73.70	
Bi-BloSAN [19]	94.56	62.13	-	-	92.45	98.77	76.28	
LEAM (Linear) [9]	93.43	61.03	-	-	91.75	98.32	75.22	
LEAM [9]	95.31	64.09	-	-	92.45	99.02	77.42	
EXAM [10]	-	-	95.50	61.90	93.00	99.00	74.80	
E1-E2 [45]	96.70	67.00	96.00	63.10	93.20	99.00	75.00	
Transformer [38]	96.13	65.34	-	-	93.89	98.98	-	
CNLE (Ours)	97.13	68.15	96.23	64.18	94.00	99.17	75.78	

Table 4

Test micro-F1 on 2 multi-label benchmarks. The best results are in bold.

Model	AAPD	Reuters
SVM [44]	69.1	86.1
TextCNN [27]	51.4	80.8
HAN [8]	68.0	85.2
XML-CNN [6]	68.7	86.2
SGM [14]	71.0	78.8
LSTM _{base} [44]	69.6	84.9
LSTM _{reg} [44]	70.5	87.0
CNLE (Ours)	71.7	89.9

tage in the real-life scenarios with strict time and space constraints.

5.3. Analysis on the weights sharing strategy

We conduct a comprehensive experiment to analyze the effect of different weights sharing strategies for TLCE. The results are shown in Table 8. From Table 8, whether sharing the weights or not does not affect the performance significantly. This result should not be surprising since both the "text" and "label" are just

Table 6 Ablation study of CNLE on 2 multi-label bonchmark

blation	study of	CNLE on 2	multi-label	benchmarks. I	Best resul	lts are i	n bold.
---------	----------	-----------	-------------	---------------	------------	-----------	---------

Model	AAPD	Reuters
CNLE w/o Label Embedding	71.6	88.0
CNLE w/o Text-to-label Attention	70.6	89.4
CNLE	71.7	89.9

tokens containing some semantic information. Moreover, this is a piece of convincing evidence that we can model both the text sequence and the label sequence in one modality (even within one module). Due to the high similarity of text and labels in the semantic space, sharing the parameters in the encoding module does not impact the performance significantly. The experimental results further validate that modeling co-attention between text and label is a reasonable choice.

6. Conclusion

In this paper, we introduce a text-label co-attention mechanism for more effective text classification. We seek to obtain the mutu-



Fig. 2. Visualization of the attention scores in the multi-head co-attention layer. We visualize them from a sample with label *Neutral* in the test set of the Yelp Full Review dataset. The top three rows show the label-to-text multi-head attention scores, while the three rows at the bottom show the text-to-label attention scores. Best viewed in color.

Table 5

Ablation study of CNLE on 4 multi-class benchmarks. Best results are in bold.

Model	Yelp P.	Yelp F.	AG	DBP.
CNLE w/o Label Embedding	96.96	68.07	93.74	99.15
CNLE w/o Text-to-label Attention	96.66	68.10	93.93	99.17
CNLE	97.13	68.15	94.00	99.17

Comparison with large pre-trained models in terms of the test accuracy, number of parameters and the average performance gain (Δ). Best results are in bold.

Model	Yelp P.	Yelp F.	Amz. P.	Amz. F.	DBP.	Yah. A.	Params. (M)	Δ (%)
CNLE (Ours) BERT [21]	97.13 98.11	68.15 70.68	96.23 97.37	64.18 65.83	99.17 99.36	75.78 77.62	3 340	- 0.015
XLNet [22]	98.45	72.20	97.60	67.74	99.38	-	340	0.025

Table 8

Analysis of the effect of sharing the weights in different modules of TLCE in multi-class and multi-label classification. CE, CA, and FE refer to the contextual encoding layer, coattention layer, and final encoding layer, respectively. The """ indicates that the weights of this module are shared between text and label sequences. Note that the metrics used in multi-class and multi-label classification are accuracy and micro-F1 respectively.

CE	CA	FE	Yelp F.	Amz. F.	AAPD	Reuters
			68.15	64.18	71.7	89.8
-	100		68.15	64.21	71.1	89.8
		1	68.25	64.18	72.1	89.7

ally attended representations for both the text sequence and the label sequence. In this way, we are able to focus on the relevant parts of the text and labels to facilitate text classification. Moreover, with the generative decoder, we are able to cope with both multi-class classification and multi-label classification without modifying the model. The extensive experiments on 7 multi-class classification benchmarks and 2 multi-label benchmarks demonstrate the superiority of our method over the considered baseline methods. In future work, the idea of the co-attention mechanism can be further applied to other NLP tasks, such as few-shot text classification and machine reading comprehension.

CRediT authorship contribution statement

Minqian Liu: Conceptualization, Methodology, Software, Writing - original draft, Validation. Lizhao Liu: Conceptualization, Methodology, Software, Writing - original draft. Junyi Cao: Methodology, Software, Writing - original draft, Investigation. Qing Du: Writing - review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] T. Sahni, C. Chandak, N.R. Chedeti, M. Singh, Efficient twitter sentiment classification using subjective distant supervision, in: 2017 9th International Conference on Communication Systems and Networks (COMSNETS), 2017, pp. 548–553.
- [2] D. Zhang, W.S. Lee, Question classification using support vector machines, in: Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval, 2003, pp. 26–32.
- [3] P.F. Brown, P.V. Desouza, R.L. Mercer, V.J.D. Pietra, J.C. Lai, Class-based n-gram models of natural language, Comput. Linguistics 18 (4) (1992) 467–479.
- [4] H.M. Wallach, Topic modeling: beyond bag-of-words, in: Proceedings of the 23rd international conference on Machine learning, 2006, pp. 977–984.
- [5] A. Conneau, H. Schwenk, L. Barrault, Y. Lecun, Very deep convolutional networks for text classification, in: Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers), 2017, pp. 1107–1116.
- [6] J. Liu, W.-C. Chang, Y. Wu, Y. Yang, Deep learning for extreme multi-label text classification, in: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2017, pp. 115–124.
- [7] P. Liu, X. Qiu, X. Huang, Recurrent neural network for text classification with multi-task learning, in: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, 2016, pp. 2873–2879.
- [8] Z. Yang, D. Yang, C. Dyer, X. He, A. Smola, E. Hovy, Hierarchical attention networks for document classification, in: Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies, 2016, pp. 1480–1489.

- [9] G. Wang, C. Li, W. Wang, Y. Zhang, D. Shen, X. Zhang, R. Henao, L. Carin, Joint embedding of words and labels for text classification, in: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2018, pp. 2321–2331.
- [10] C. Du, Z. Chen, F. Feng, L. Zhu, T. Gan, L. Nie, Explicit interaction model towards text classification, in: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, 2019, pp. 6359–6366..
- [11] M. Seo, A. Kembhavi, A. Farhadi, H. Hajishirzi, Bidirectional attention flow for machine comprehension, in: International Conference on Learning Representations, 2017..
- [12] C. Xiong, V. Zhong, R. Socher, Dynamic coattention networks for question answering, in: International Conference on Learning Representations, 2017..
- [13] B. McCann, N.S. Keskar, C. Xiong, R. Socher, The natural language decathlon: Multitask learning as question answering, arXiv preprint arXiv:1806.08730..
- [14] P. Yang, X. Sun, W. Li, S. Ma, W. Wu, H. Wang, Sgm: sequence generation model for multi-label classification, in: Proceedings of the 27th International Conference on Computational Linguistics, 2018, pp. 3915–3926.
- [15] A. McCallum, K. Nigam, et al., A comparison of event models for naive bayes text classification, in: AAAI-98 workshop on learning for text categorization, vol. 752, Citeseer, 1998, pp. 41–48..
- [16] H.-J. Kim, J.-U. Kim, Y.-G. Ra, Boosting nai ve bayes text classification using uncertainty-based selective sampling, Neurocomputing 67 (2005) 403–410.
- [17] S.R. Gunn, et al., Support vector machines for classification and regression, 1998..
- [18] S. Zhou, Q. Chen, X. Wang, Active deep learning method for semi-supervised sentiment classification, Neurocomputing 120 (2013) 536–546.
- [19] T. Shen, T. Zhou, G. Long, J. Jiang, C. Zhang, Bi-directional block self-attention for fast and memory-efficient sequence modeling, in: International Conference on Learning Representations, 2018..
- [20] G. Liu, J. Guo, Bidirectional lstm with attention mechanism and convolutional layer for text classification, Neurocomputing 337 (2019) 325–338.
- [21] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), 2019, pp. 4171–4186.
- [22] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R.R. Salakhutdinov, Q.V. Le, Xlnet: Generalized autoregressive pretraining for language understanding, in: Advances in neural information processing systems, 2019, pp. 5754–5764..
- [23] J. Kim, S. Jang, E. Park, S. Choi, Text classification using capsules, Neurocomputing 376 (2020) 214–221.
- [24] X. Jing, F. Wu, Z. Li, R. Hu, D. Zhang, Multi-label dictionary learning for image annotation, IEEE Trans. Image Process. 25 (6) (2016) 2712–2725.
- [25] X. Jing, X. Zhu, F. Wu, R. Hu, X. You, Y. Wang, H. Feng, J. Yang, Super-resolution person re-identification with semi-coupled low-rank discriminant dictionary learning, IEEE Trans. Image Process. 26 (3) (2017) 1363–1378.
- [26] X. Zhu, X. Jing, X. You, W. Zuo, S. Shan, W. Zheng, Image to video person reidentification by learning heterogeneous dictionary pair with feature projection matrix, IEEE Trans. Inf. Forensics Secur. 13 (3) (2018) 717–732.
- [27] Y. Kim, Convolutional neural networks for sentence classification, in: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2014, pp. 1746–1751.
- [28] X. Zhang, J. Zhao, Y. LeCun, Character-level convolutional networks for text classification, in: Advances in neural information processing systems, 2015, pp. 649–657.
- [29] P. Wang, B. Xu, J. Xu, G. Tian, C.-L. Liu, H. Hao, Semantic expansion using word embedding clustering and convolutional neural network for improving short text classification, Neurocomputing 174 (2016) 806–814.
- [30] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Comput. 9 (8) (1997) 1735–1780.
- [31] D. Yogatama, C. Dyer, W. Ling, P. Blunsom, Generative and discriminative text classification with recurrent neural networks, arXiv preprint arXiv:1703.01898..

M. Liu, L. Liu, J. Cao et al.

- [32] H. Zhang, L. Xiao, W. Chen, Y. Wang, Y. Jin, Multi-task label embedding for text classification, in: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2018, pp. 4545–4553.
- [33] J. Lu, J. Yang, D. Batra, D. Parikh, Hierarchical question-image co-attention for visual question answering, in: Advances in neural information processing systems, 2016, pp. 289–297..
- [34] Z. Yu, J. Yu, J. Fan, D. Tao, Multi-modal factorized bilinear pooling with coattention learning for visual question answering, in: Proceedings of the IEEE international conference on computer vision, 2017, pp. 1821–1830.
- [35] J. Lu, D. Batra, D. Parikh, S. Lee, Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks, Adv. Neural Inform. Process. Syst. (2019) 13–23.
- [**36**] Y. Li, T. Liu, J. Hu, J. Jiang, Topical co-attention networks for hashtag recommendation on microblogs, Neurocomputing 331 (2019) 356–365.
- [37] L. Zhang, Z. Guan, A. Hauptmann, The co-attention model for tiny activity analysis, Neurocomputing 105 (2013) 51–60.
- [38] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, in: Advances in neural information processing systems, 2017, pp. 5998–6008.
- [39] A. Graves, J. Schmidhuber, Framewise phoneme classification with bidirectional lstm and other neural network architectures, Neural Networks 18 (5–6) (2005) 602–610.
- [40] J. Pennington, R. Socher, C.D. Manning, Glove: Global vectors for word representation, in: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP), 2014, pp. 1532–1543.
- [41] K. Hashimoto, C. Xiong, Y. Tsuruoka, R. Socher, A joint many-task model: growing a neural network for multiple nlp tasks, in: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2017, pp. 1923–1933.
- [42] J.L. Ba, J.R. Kiros, G.E. Hinton, Layer normalization, arXiv preprint arXiv:1607.06450..
- [43] T. Sahni, C. Chandak, N.R. Chedeti, M. Singh, Efficient twitter sentiment classification using subjective distant supervision, in: 2017 9th International Conference on Communication Systems and Networks (COMSNETS), 2017, pp. 548–553.
- [44] A. Adhikari, A. Ram, R. Tang, J. Lin, Rethinking complex neural network architectures for document classification, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), 2019, pp. 4046–4051.
- [45] G. Niu, H. Xu, B. He, X. Xiao, H. Wu, G. Sheng, Enhancing local feature extraction with global representation for neural text classification, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019, pp. 496–506.
- [46] J. McAuley, J. Leskovec, Hidden factors and hidden topics: understanding rating dimensions with review text, in: Proceedings of the 7th ACM conference on Recommender systems, 2013, pp. 165–172.
- [47] G.M. Del Corso, A. Gulli, F. Romani, Ranking a stream of news, in: Proceedings of the 14th international conference on World Wide Web, 2005, pp. 97–106.
- [48] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P.N. Mendes, S. Hellmann, M. Morsey, P. Van Kleef, S. Auer, et al., Dbpedia-a large-scale, multilingual knowledge base extracted from wikipedia, Semantic Web 6 (2) (2015) 167–195.
- [49] C. Apté, F. Damerau, S.M. Weiss, Automated learning of decision rules for text categorization, ACM Trans. Inform. Syst. (TOIS) 12 (3) (1994) 233–251.
- [50] C.D. Manning, P. Raghavan, H. Schütze, Introduction to information retrieval, Cambridge University Press, 2008.
- [51] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: International Conference on Learning Representations, 2015.
- [52] E. Voita, D. Talbot, F. Moiseev, R. Sennrich, I. Titov, Analyzing multi-head selfattention: Specialized heads do the heavy lifting, the rest can be pruned, in: Proceedings of the 57th Conference of the Association for Computational Linguistics (Volume 1: Long Papers), 2019, pp. 5797–5808.
- [53] M. Behnke, K. Heafield, Losing heads in the lottery: pruning transformer attention in neural machine translation, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020, pp. 2664–2674.



Minqian Liu received the B.E. degree from the School of Computer Science and Engineering, South China University of Technology. He is currently studying for the Ph.D. degree in computer science and application at the Department of Computer Science, Virginia Tech. His main research interests include natural language processing, information retrieval, computer vision, and deep learning.



Lizhao Liu received the B.E. degree from the School of Software Engineering, South China University of Technology, where he is currently studying for the master's degree in electrical and information engineering. His main research interests cover machine learning, visual recognition and deep learning.



Junyi Cao received the B.E. degree at the School of Software Engineering, South China University of Technology. He is currently studying for the master's degree in computer science and technology at the Department of Computer Science and Engineering, Shanghai Jiao Tong University. His main research interests cover neural networks, visual recognition and deep learning.



Qing Du received the B.S. degree in computer science and technology, the master's degree in computer application, and the Ph.D. degree in computer application from the South China University of Technology respectively, where she is currently an Associate Professor with the School of Software Engineering. Her research interests include information retrieval, recommendation systems, natural language processing, and deep learning.

Neurocomputing 471 (2022) 61-69