



A Survey on Text Classification: From Traditional to Deep Learning

QIAN LI, HAO PENG, and JIANXIN LI, Beihang University, China

CONGYING XIA, University of Illinois at Chicago, USA

RENYU YANG, University of Leeds, UK

LICHAO SUN, Lehigh University, USA

PHILIP S. YU, University of Illinois at Chicago, USA

LIFANG HE, Lehigh University, USA

Text classification is the most fundamental and essential task in natural language processing. The last decade has seen a surge of research in this area due to the unprecedented success of deep learning. Numerous methods, datasets, and evaluation metrics have been proposed in the literature, raising the need for a comprehensive and updated survey. This paper fills the gap by reviewing the state-of-the-art approaches from 1961 to 2021, focusing on models from traditional models to deep learning. We create a taxonomy for text classification according to the text involved and the models used for feature extraction and classification. We then discuss each of these categories in detail, dealing with both the technical developments and benchmark datasets that support tests of predictions. A comprehensive comparison between different techniques, as well as identifying the pros and cons of various evaluation metrics are also provided in this survey. Finally, we conclude by summarizing key implications, future research directions, and the challenges facing the research area.

CCS Concepts: • General and reference → Document types; Surveys and overviews

Additional Key Words and Phrases: Deep learning, traditional models, text classification, evaluation metrics, challenges

ACM Reference format:

Qian Li, Hao Peng, Jianxin Li, Congying Xia, Renyu Yang, Lichao Sun, Philip S. Yu, and Lifang He. 2022. A Survey on Text Classification: From Traditional to Deep Learning. *ACM Trans. Intell. Syst. Technol.* 13, 2, Article 31 (April 2022), 41 pages.

<https://doi.org/10.1145/3495162>

The authors of this paper were supported by the National Key R&D Program of China through grant 2021YFB1714800, NSFC through grants (No. U20B2053 and 61872022), State Key Laboratory of Software Development Environment (SKLSD-E-2020ZX-12). Philip S. Yu was supported by NSF under grants III-1763325, III-1909323, III-2106758, and SaTC-1930941. Lifang He was supported by NSF ONR N00014-18-1-2009 and Lehigh's accelerator grant S00010293. This work was also sponsored by CAAI-Huawei MindSpore Open Fund.

Authors' addresses: Q. Li, H. Peng, and J. Li (corresponding author), Beihang University, Haidian district, Beijing, China, 100191; emails: {liqian, penghao, lijx}@act.buaa.edu.cn; C. Xia and P. S. Yu, University of Illinois at Chicago, 601 S Morgan, Chicago, IL, USA, 61820; emails: {cxia8, psyu}@uic.edu; R. Yang, University of Leeds, Leeds LS2 9JT, Leeds, England, UK; email: r.yang@leeds.ac.uk; L. Sun and L. He, Lehigh University, 27 Memorial Drive West, Bethlehem, PA, USA, 18015; emails: james.lichao.sun@gmail.com, lih319@lehigh.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

2157-6904/2022/04-ART31 \$15.00

<https://doi.org/10.1145/3495162>

1 INTRODUCTION

Text classification – the procedure of designating pre-defined labels for text – is an essential and significant task in many **Natural Language Processing (NLP)** applications, such as sentiment analysis [1, 2], topic labeling [3, 4], question answering [5, 6] and dialog act classification [7]. In the era of information explosion, it is time-consuming and challenging to process and classify large amounts of text data manually. Besides, the accuracy of manual text classification can be easily influenced by human factors, such as fatigue and expertise. It is desirable to use machine learning methods to automate the text classification procedure to yield more reliable and less subjective results. Moreover, this can also help enhance information retrieval efficiency and alleviate the problem of information overload by locating the required information.

Figure 1 illustrates a flowchart of the procedures involved in the text classification, under the light of traditional and deep analysis. Text data is different from numerical, image, or signal data. It requires NLP techniques to be processed carefully. The first important step is to preprocess text data for the model. Traditional models usually need to obtain good sample features by artificial methods and then classify them with classic machine learning algorithms. Therefore, the effectiveness of the method is largely restricted by feature extraction. However, different from traditional models, deep learning integrates feature engineering into the model fitting process by learning a set of nonlinear transformations that serve to map features directly to outputs.

From the 1960s until the 2010s, traditional text classification models dominated. Traditional methods mean statistics-based models, such as **Naïve Bayes (NB)** [8], **K-Nearest Neighbor (KNN)** [9], and **Support Vector Machine (SVM)** [10]. Comparing with the earlier rule-based methods, this method has obvious advantages in accuracy and stability. However, these approaches still need to do feature engineering, which is time-consuming and costly. Besides, they usually disregard the natural sequential structure or contextual information in textual data, making it challenging to learn the semantic information of the words. Since the 2010s, text classification has gradually changed from traditional models to deep learning models. Compared with the methods based on traditional, deep learning methods avoid designing rules and features by humans and automatically provide semantically meaningful representations for text mining. Therefore, most of the text classification research works are based on **Deep Neural Networks (DNNs)** [11], which are data-driven approaches with high computational complexity. Few works focus on traditional models to settle the limitations of computation and data.

1.1 Major Differences and Contributions

There have been several works reviewing text classification and its subproblems recently. Two of them are reviews of text classification. Kowsari et al. [12] surveyed different text feature extraction, dimensionality reduction methods, basic model structure for text classification, and evaluation methods. Minaee et al. [13] reviewed recent deep learning based text classification methods, benchmark datasets, and evaluation metrics. Unlike existing text classification reviews, we conclude existing models from traditional models to deep learning with works of recent years. Traditional models emphasize the feature extraction and classifier design. Once the text has well-designed characteristics, it can be quickly converged by training the classifier. DNNs can perform feature extraction automatically and learn well without domain knowledge. We then give the datasets and evaluation metrics for single-label and multi-label tasks and summarize future research challenges from data, models, and performance perspectives. Moreover, we summarize various information in three tables, including the necessary information of classic deep learning models, primary information of main datasets, and a general benchmark of state-of-the-art methods under different applications. In summary, this study's main contributions are as follows:

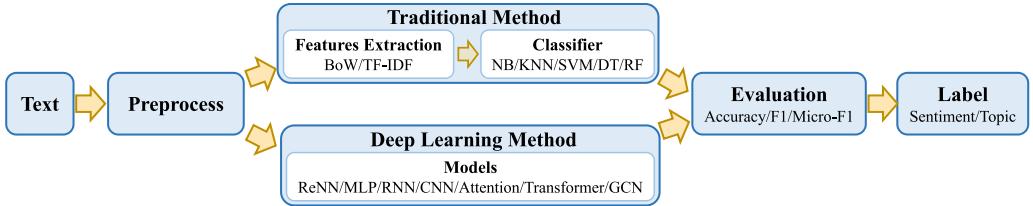


Fig. 1. Flowchart of the text classification with classic methods in each module. It is crucial to extract essential features for traditional methods, but features can be extracted automatically by deep learning methods.

- We introduce the process and development of text classification and present comprehensive analysis and research on primary models – from traditional to deep learning models – according to their model structures. We summarize the necessary information of deep learning models in terms of basic model structures in Table 1, including publishing years, methods, venues, applications, evaluation metrics, datasets, and code links.
- We introduce the present datasets and give the formulation of main evaluation metrics with the comparison of metrics, including single-label and multi-label text classification tasks. We summarize the necessary information of primary datasets in Table 2, including the number of categories, average sentence length, the size of each dataset, related papers, and data addresses.
- We summarize classification accuracy scores of models given in their articles, on benchmark datasets in Table 4 and conclude the survey by discussing the main challenges facing the text classification and key implications stemming from this study.

1.2 Organization of the Survey

The rest of the survey is organized as follows. Section 2 summarizes the existing models related to text classification, including traditional and deep learning models, including a summary table. Section 3 introduces the primary datasets with a summary table and evaluation metrics on single-label and multi-label tasks. We then give quantitative results of the leading models in classic text classification datasets in Section 4. Finally, we summarize the main challenges for deep learning text classification in Section 5 before concluding the article in Section 6.

2 TEXT CLASSIFICATION METHODS

Text classification is referred to as extracting features from raw text data and predicting the categories of text data based on such features. Numerous models have been proposed in the past few decades for text classification. For traditional models, NB [8] is the first model used for the text classification task. Whereafter, generic classification models are proposed, such as KNN [9], SVM [10], and **Random Forest (RF)** [14], which are called classifiers, and are widely used for text classification. Recently, the **eXtreme Gradient Boosting (XGBoost)** [15] and the **Light Gradient Boosting Machine (LightGBM)** [16] have arguably the potential to provide excellent performance. For deep learning models, TextCNN [17] has the highest number of references in these models, wherein a **Convolutional Neural Network (CNN)** [18] model has been introduced to solve the text classification problem for the first time. While not specifically designed for handling text classification tasks, the **Bidirectional Encoder Representation from Transformers (BERT)** [19] has been widely employed when designing text classification models, considering its effectiveness on numerous text classification datasets.

2.1 Traditional Models

Traditional models accelerate text classification with improved accuracy and make the application scope of traditional expand. The first thing is to preprocess the raw input text for training traditional models, which generally consists of word segmentation, data cleaning, and statistics. Then, text representation aims to express preprocessed text in a form that is much easier for computers and minimizes information loss, such as **Bag-Of-Words (BOW)** [20], N-gram [21], **Term Frequency-Inverse Document Frequency (TF-IDF)** [22], word2vec [23], and **Global Vectors** for word representation (**GloVe**) [24]. BOW means that all the words in the corpus are formed into a mapping array. According to the mapping array, a sentence can be represented as a vector. The i -th element in the vector represents the frequency of the i -th word in the mapping array of the sentence. The vector is the BOW of the sentence. At the core of the BOW is representing each text with a dictionary-sized vector. The individual value of the vector denotes the word frequency corresponding to its inherent position in the text. Compared to BOW, N-gram considers the information of adjacent words and builds a dictionary by considering the adjacent words. It is used to calculate the probability model of a sentence. The probability of a sentence is expressed as the joint probability of each word in the sentence. The probability of a sentence can be calculated by predicting the probability of the N -th word, given the sequence of the $(N-1)$ -th words. To simplify the calculation, the N-gram model adopts the Markov hypothesis [21]. A word appears only concerning the words that preceded it. Therefore, the N-gram model performs a sliding window with size N . By counting and recording the occurrence frequency of all fragments, the probability of a sentence can be calculated using the frequency of relevant fragments in the record. TF-IDF [22] uses the word frequency and inverses the document frequency to model the text. TF is the word frequency of a word in a specific article, and IDF is the reciprocal of the proportion of the articles containing this word to the total number of articles in the corpus. TF-IDF is the multiplication of the two. TF-IDF assesses the importance of a word to one document in a set of files or a corpus. The importance of a word increases proportionally with the number of times it appears in a document. However, it decreases inversely with its frequency in the corpus as a whole. The word2vec [23] employs local context information to obtain word vectors, as shown in Figure 2. Word vector refers to a fixed-length real value vector specified as the word vector for any word in the corpus. The word2vec uses two essential models: CBOW and Skip-gram. The former is to predict the current word on the premise that the context of the current word is known. The latter is to predict the context when the current word is known. The GloVe [24] – with both the local context and global statistical features – trains on the nonzero elements in a word-word co-occurrence matrix, as shown in Figure 2. It enables word vectors to contain as much semantic and grammatical information as possible. The construction method of word vector is: firstly, the co-occurrence matrix of words is constructed based on the corpus, and then the word vector is learned based on the co-occurrence matrix and GloVe model. Finally, the represented text is fed into the classifier according to selected features. Here, we discuss some representative classifiers in detail.

2.1.1 PGM-based Methods. **Probabilistic Graphical Models (PGMs)** express the conditional dependencies among features in graphs, such as the Bayesian network [25]. It is a combination of probability theory and graph theory.

Naïve Bayes (NB) [8] is the simplest and most broadly used model based on applying Bayes' theorem. The NB algorithm has an independent assumption: when the target value has been given, the conditions between text $T = [T_1, T_2, \dots, T_n]$ are independent (see Figure 3). The NB algorithm primarily uses the prior probability to calculate the posterior probability $P(y | T_1, T_2, \dots, T_n) = \frac{p(y) \prod_{j=1}^n p(T_j | y)}{\prod_{j=1}^n p(T_j)}$. Due to its simple structure, NB is broadly used for text classification tasks.

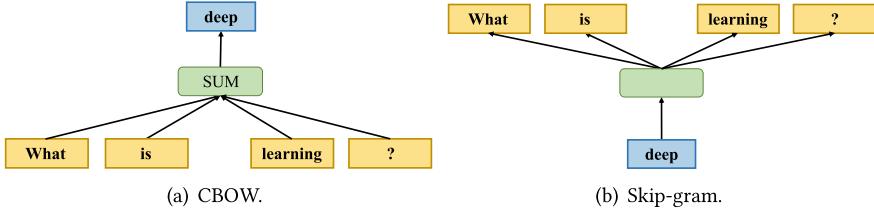


Fig. 2. The structure of word2vec, including CBOW and Skip-gram.

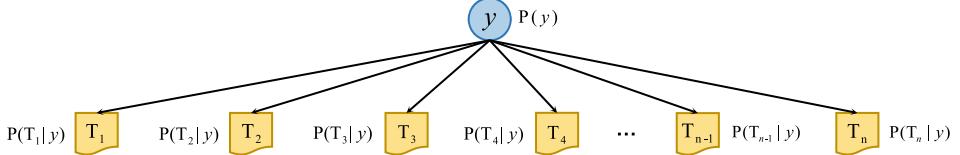
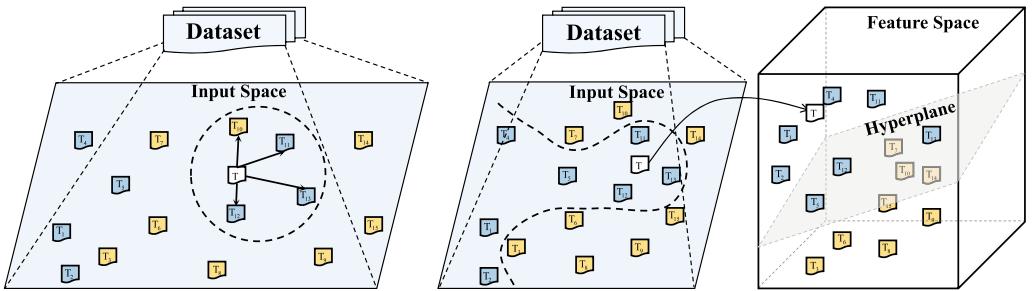


Fig. 3. The structure of Naïve Bayes.

Fig. 4. The structure of KNN where $k = 4$ (left) and the structure of SVM (right). Each node represents a text and nodes with different contours represent different categories.

Although the assumption that the features are independent is sometimes not actual, it substantially simplifies the calculation process and performs better. To improve the performance on smaller categories, Schneider [26] proposes a feature selection score method through calculating KL-divergence [27] between the training set and corresponding categories for multinomial NB text classification. Dai et al. [28] propose a transfer learning method named **Naïve Bayes Transfer Classification (NBTC)** to settle the different distribution between the training set and the target set. It uses the EM algorithm [29] to obtain a locally optimal posterior hypothesis on the target set. NB classifier is also used for fake news detection [30], and sentiment analysis [31], which can be seen as a text classification task. Bernoulli NB, Gaussian NB and Multinomial NB are three popular approaches of NB text classification [32]. Multinomial NB performs slightly better than Bernoulli NB on few labeled datasets [33]. Bayesian NB classifier with Gaussian event model [32] has been proven to be superior to NB with multinomial event model on **20 Newsgroups (20NG)** [34] and WebKB [35] datasets.

2.1.2 KNN-based Methods. At the core of the **K-Nearest Neighbors (KNN)** algorithm [9] is to classify an unlabeled sample by finding the category with the most samples on the k nearest samples. It is a simple classifier without building the model and can decrease complexity through the fasting process of getting k nearest neighbors. Figure 4 showcases the structure of KNN.

We can find k training texts approaching a specific text to be classified through estimating the in-between distance. Hence, the text can be divided into the most common categories found in k training set texts. The improvement of KNN algorithm mainly includes feature similarity [36], K value [37] and index optimization [38]. However, due to the positive correlation between model time/space complexity and the amount of data, the KNN algorithm takes an unusually long time on the large-scale datasets [39]. To decrease the number of selected features, Soucy et al. [40] propose a KNN algorithm without feature weighting. It manages to find relevant features, building the inter-dependencies of words by using a feature selection. When the data is extremely unevenly distributed, KNN tends to classify samples with more data. The **Neighbor-Weighted K-Nearest Neighbor (NWKNN)** [41] is proposed to improve classification performance on the unbalanced corpora. It casts a significant weight for neighbors in a small category and a small weight for neighbors in a broad class.

2.1.3 SVM-based Methods. Cortes and Vapnik [42] propose **Support Vector Machine (SVM)** to tackle the binary classification of pattern recognition. Joachims [10], for the first time, uses the SVM method for text classification representing each text as a vector. As illustrated in Figure 4, SVM-based approaches turn text classification tasks into multiple binary classification tasks. In this context, SVM constructs an optimal hyperplane in the one-dimensional input space or feature space, maximizing the distance between the hyperplane and the two categories of training sets, thereby achieving the best generalization ability. The goal is to make the distance of the category boundary along the direction perpendicular to the hyperplane to be the largest. Equivalently, this will result in the lowest error rate of classification. Constructing an optimal hyperplane can be transformed into a quadratic programming problem to obtain a globally optimal solution. Choosing the appropriate kernel function [43] and feature selection [44] are of the utmost importance to ensure SVM can deal with nonlinear problems and become a robust nonlinear classifier. Furthermore, active learning [45] and adaptive learning [46] method are used for text classification to reduce the labeling effort based on the supervised learning algorithm SVM. To analyze what the SVM algorithms learn and what tasks are suitable, Joachims [47] proposes a theoretical learning model combining the statistical traits with the generalization performance of an SVM, analyzing the features and benefits using a quantitative approach. **Transductive Support Vector Machine (TSVM)** [48] is proposed to lessen misclassifications of the particular test collections with a general decision function considering a specific test set. It uses prior knowledge to establish a more suitable structure and study faster.

2.1.4 DT-based Methods. Decision Trees (DT) [49] is a supervised tree structure learning method – reflective of the idea of divide-and-conquer – and is constructed recursively. It learns disjunctive expressions and has robustness for the text with noise. As shown in Figure 5, decision trees can be generally divided into two distinct stages: tree construction and tree pruning. It starts at the root node and tests the data samples (composed of instance sets, which have several attributes), and divides the dataset into diverse subsets according to different results. A subset of datasets constitutes a child node, and every leaf node in the decision tree represents a category. Constructing the decision tree is to determine the correlation between classes and attributes, further exploited to predict the record categories of unknown forthcoming types. The classification rules generated by the decision tree algorithm are straight-forward, and the pruning strategy [50] can also help reduce the influence of noise. Its limitation, however, mainly derives from inefficiency in coping with explosively increasing data size. More specifically, the **Iterative Dichotomiser 3 (ID3)** [51] algorithm uses information gain as the attribute selection criterion in the selection of each node – It is used to select the attribute of each branch node, and then select the attribute having the maximum information gain value to become the discriminant attribute of the current

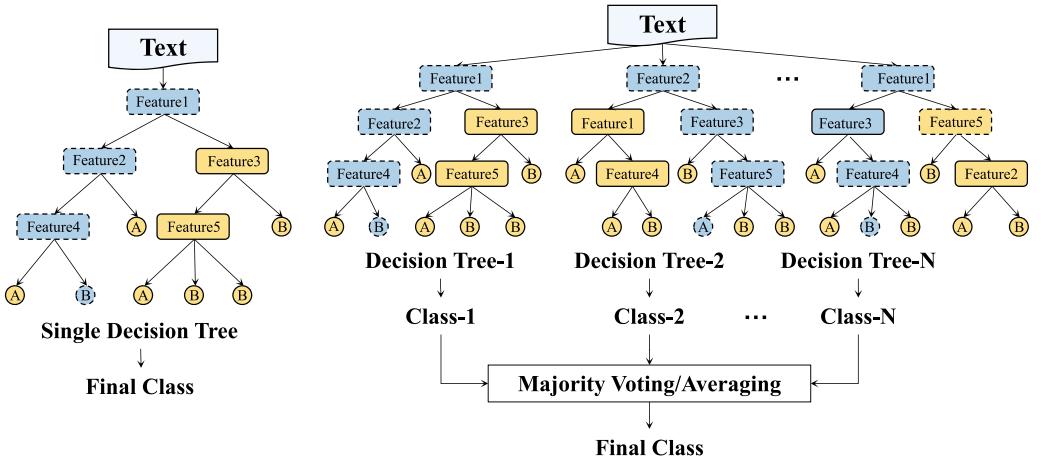


Fig. 5. An example of DT (left) and the structure of RF (right). The nodes with the dotted outline represent the nodes of the decision route. It has five features to predict whether each text belongs to category A or B.

node. Based on ID3, C4.5 [52] learns to obtain a map from attributes to classes, which effectively classifies entities unknown to new categories. DT based algorithms usually need to train for each dataset, which is low efficiency [53]. Thus, Johnson et al. [54] propose a DT-based symbolic rule system. The method represents each text as a vector calculated by the frequency of each word in the text, and induces rules from the training data. The learning rules are used for classifying the other data, being similar to the training data. Furthermore, to reduce the computational costs of DT algorithms, **Fast Decision-Tree (FDT)** [55] uses a two-pronged strategy: pre-selecting a feature set and training multiple DTs on different data subsets. Results from multiple DTs are combined through a data-fusion technique to resolve the cases of imbalanced classes.

2.1.5 Integration-based Methods. Integrated algorithms aim to aggregate the results of multiple algorithms for better performance and interpretation. Conventional integrated algorithms are bootstrap aggregation, such as RF [14], boosting such as the **Adaptive Boosting (AdaBoost)** [56] and XGBoost [15], and stacking. The bootstrap aggregation method trains multiple classifiers without strong dependencies and then aggregates their results. For instance, RF [14] consists of multiple tree classifiers wherein all trees depend on the value of the random vector sampled independently (depicted in Figure 5). It is worth noting that each tree within the RF shares the same distribution. The generalization error of an RF relies on the strength of each tree and the relationship among trees, and will converge to a limit with the increment of tree number in the forest. In boosting based algorithms, all labeled data are trained with the same weight to initially obtain a weaker classifier [57]. The weights of the data will then be adjusted according to the former result of the classifier. The training procedure will continue by repeating such steps until the termination condition is reached. Unlike bootstrap and boosting algorithms, stacking based algorithms break down the data into n parts and use n classifiers to calculate the input data in a cascade manner – Results from the upstream classifier will feed into the downstream classifier as input. The training will terminate once a pre-defined iteration number is targeted. The integrated method can capture more features from multiple trees. However, it helps little for short text. Motivated by this, Bouaziz et al. [58] combine data enrichment – with semantics in RFs for short text classification – to overcome the deficiency of sparseness and insufficiency of contextual information. In integrated algorithms, not all classifiers learn well. It is necessary to give different weights for each classifier.

To differentiate contributions of trees in a forest, Islam et al. [59] exploit the **Semantics Aware Random Forest (SARF)** classifier, choosing features similar to the features of the same class, for extracting features and producing the prediction values.

Summary. The parameters of NB are more diminutive, less sensitive to missing data, and the algorithm is simple. However, it assumes that features are independent of each other. When the number of features is large, or the correlation between features is significant, the performance of NB decreases. SVM can solve high-dimensional and nonlinear problems. It has a high generalization ability, but it is sensitive to missing data. KNN mainly depends on the surrounding finite adjacent samples, rather than discriminating class domain to determine the category. Thus, for the dataset to be divided with more crossover or overlap of the class domain, it is more suitable than other methods. DT is easy to understand and interpret. Given an observed model, it is easy to deduce the corresponding logical expression according to the generated decision tree. The traditional method is a type of machine learning. It learns from data, which are pre-defined features that are important to the performance of prediction values. However, feature engineering is tough work. Before training the classifier, we need to collect knowledge or experience to extract features from the original text. The traditional methods train the initial classifier based on various textual features extracted from the raw text. Toward small datasets, traditional models usually present better performance than deep learning models under the limitation of computational complexity. Therefore, some researchers have studied the design of traditional models for specific domains with less data.

2.2 Deep Learning Models

The DNNs consist of artificial neural networks that simulate the human brain to automatically learn high-level features from data, getting better results than traditional models in speech recognition, image processing, and text understanding. Input datasets should be analyzed to classify the data, such as a single-label, multi-label, unsupervised, unbalanced dataset. According to the trait of the dataset, the input word vectors are sent into the DNN for training until the termination condition is reached. The performance of the training model is verified by the downstream task, such as sentiment classification, question answering, and event prediction. We show some DNNs over the years in Table 1, including designs that are different from the corresponding basic models, evaluation metrics, and experimental datasets.

Numerous deep learning models have been proposed in the past few decades for text classification, as shown in Table 1. We tabulate primary information – including publication years, venues, applications, code links, evaluation metrics, and experiment datasets – of main deep learning models for text classification. The applications in this table include **Sentiment Analysis (SA)**, **Topic Labeling (TL)**, **News Classification (NC)**, **Question Answering (QA)**, **Dialog Act Classification (DAC)**, **Natural Language Inference (NLI)**, and **Relation Classification (RC)**. The multilayer perceptron [172] and the recursive neural network [173] are the first two deep learning approaches used for the text classification task, which improve performance compared with traditional models. Then, CNNs, **Recurrent Neural Networks (RNNs)**, and attention mechanisms are used for text classification [101, 174, 175]. Many researchers advance text classification performance for different tasks by improving CNN, RNN, and attention, or model fusion and multi-task methods. The appearance of BERT [19], which can generate contextualized word vectors, is a significant turning point in the development of text classification and other NLP technologies. Many researchers [142, 176] have studied text classification models based on BERT, which achieves better performance than the above models in multiple NLP tasks, including text classification. Besides, some researchers study text classification technology based on **Graph Neural Network (GNN)**

Table 1. Basic Information Based on Different Models

Model	Year	Method	Venue	Applications	Code Link	Metrics	Datasets
ReNN	2011	RAE [60]	EMNLP	SA, QA	[61]	Accuracy	MPQA, MR, EP
	2012	MV-RNN [62]	EMNLP	SA	[63]	Accuracy, F1	MR
	2013	RNTN [64]	EMNLP	SA	[65]	Accuracy	SST
MLP	2014	DeepRNN [66]	NIPS	SA, QA	-	Accuracy	SST-1;SST-2
	2014	Paragraph-Vec [67]	ICML	SA, QA	[68]	Error Rate	SST, IMDB
RNN	2015	DAN [69]	ACL	SA, QA	[70]	Accuracy, Time	RT, SST, IMDB
	2015	Tree-LSTM [1]	ACL	SA	[71]	Accuracy	SST-1, SST-2
	2015	S-LSTM [2]	ICML	SA	-	Accuracy	SST
	2015	TextRCNN [72]	AAAI	SA, TL	[73]	Macro-F1, etc.	20NG, Fudan, ACL, SST-2
	2015	MT-LSTM [6]	EMNLP	SA, QA	[74]	Accuracy	SST-1, SST-2, QC, IMDB
	2016	oh-2LSTMp [75]	ICML	SA, TL	[76]	Error Rate	IMDB, Elec, RCV1, 20NG
	2016	BLSTM-2DCNN [77]	COLING	SA, QA, TL	[78]	Accuracy	SST-1, Subj, TREC, etc.
	2016	Multi-Task [79]	IJCAI	SA	[80]	Accuracy	SST-1, SST-2, Subj, IMDB
	2017	DeepMoji [81]	EMNLP	SA	[82]	Accuracy	SS-Twitter, SE1604, etc.
	2017	TopicRNN [83]	ICML	SA	[84]	Error Rate	IMDB
CNN	2017	Miyato et al. [85]	ICLR	SA	[86]	Error Rate	IMDB, DBpedia, etc.
	2018	RNN-Capsule [87]	TheWebConf	SA	[88]	Accuracy	MR, SST-1, etc.
	2019	HM-DenseRNNs [89]	IJCAI	SA, TL	[90]	Accuracy	IMDB, SST-5, AG
	2014	TextCNN [17]	EMNLP	SA, QA	[91]	Accuracy	MR, SST-2, Subj, etc.
	2014	DCNN [5]	ACL	SA, QA	[92]	Accuracy	MR, TREC, Twitter
	2015	CharCNN [93]	NeurIPS	SA, QA, TL	[94]	Error Rate	AG, Yelp P, DBpedia, etc.
	2016	SeqTextRCNN [7]	NAACL	Dialog act	[95]	Accuracy	DSTC 4, MRDA, SwDA
	2017	XML-CNN [96]	SIGIR	NC, TL, SA	[97]	NDCG@K, etc.	EUR-Lex, Wiki-30K, etc.
	2017	DPCNN [98]	ACL	SA, TL	[99]	Error Rate	AG, DBpedia, YelpP, etc.
	2017	KPCNN [100]	IJCAI	SA, QA, TL	-	Accuracy	Twitter, AG, Bing, etc.
Attention	2018	TextCapsule [101]	EMNLP	SA, QA, TL	[102]	Accuracy	Subj, TREC, Reuters, etc.
	2018	HFT-CNN [103]	EMNLP	TL	[104]	Micro-F1, etc.	RCV1, Amazon670K
	2019	CCRCNN [105]	AAAI	TL	-	Accuracy	TREC, MR, AG
	2020	Bao et al. [106]	ICLR	TL	[107]	Accuracy	20NG, Reuters-2157, etc.
	2016	HAN [108]	NAACL	SA, TL	[109]	Accuracy	YelpF, YahooA, etc.
	2016	BI-Attention [110]	NAACL	SA	-	Accuracy	NLP&CC 2013 [111]
	2016	LSTMN [112]	EMNLP	SA	[113]	Accuracy	SST-1
	2017	Lin et al. [114]	ICLR	SA	[115]	Accuracy	Yelp, SNLI Age
	2018	SGM [116]	COLING	TL	[117]	HL, Micro-F1	RCV1-V2, AAPD
	2018	ELMo [118]	NAACL	SA, QA, NLI	[119]	Accuracy	SQuAD, SNLI, SST-5
Trans	2018	BiBlaSA [120]	ICLR	SA	[121]	Accuracy, Time	CR, MPQA, SUBJ, etc.
	2019	AttentionXML [122]	NeurIPS	TL	[123]	P@k, N@k, etc.	EUR-Lex, etc.
	2019	HAPN [124]	EMNLP	RC	-	Accuracy	FewRel, CSID
	2019	Proto-HATT [125]	AAAI	RC	[126]	Accuracy	FewRel
	2019	STCKA [127]	AAAI	SA, TL	[128]	Accuracy	Weibo, Product Review, etc.
	2020	HyperGAT [129]	EMNLP	TL, NC	[130]	Accuracy	20NG, Ohsumed, MR, etc.
	2020	MSMSA [131]	AAAI	ST, QA, NLI	-	Accuracy, F1	IMDB, MR, SST, SNLI, etc.
	2020	Choi [132]	EMNLP	SA, TL	-	Accuracy	SST2, IMDB, 20NG
	2019	BERT [19]	NAACL	SA, QA	[133]	Accuracy	SST-2, QQP, QNLI, CoLA
	2019	BERT-BASE [134]	ACL	TL	[135]	P@K, R@K, etc.	EUR-LEX
GNN	2019	Sun et al. [136]	CCL	SA, QA, TL	[137]	Error Rate	TREC, DBpedia, etc.
	2019	XLNet [138]	NeurIPS	SA, QA, NC	[139]	EM, F1, etc.	Yelp-2, AG, MNLI, etc.
	2019	RoBERTa [140]	arXiv	SA, QA	[141]	F1, Accuracy	SQuAD, MNLI-m, SST-2
	2020	GAN-BERT [142]	ACL	SA, NLI	[143]	F1, Accuracy	SST-5, MNLI
	2020	BAE [144]	EMNLP	SA, QA	[145]	Accuracy	Amazon, Yelp, MR, MPQA
	2020	ALBERT [146]	ICLR	SA, QA	[147]	F1, Accuracy	SST, MNLI, SQuAD
	2020	TG-Transformer [148]	EMNLP	SA, TL	-	Accuracy, Time	R8, R52, Ohsumed, etc.
	2020	X-Transformer [149]	KDD	SA, TL	[150]	P@K, R@K	Eurlex-4K, Wiki10-31K, etc.
	2021	LightXML [151]	arXiv	TL, ML, NLI	[152]	P@K, Time	AmazonCat-13K, etc.
	2018	DGCNN [153]	TheWebConf	TL	[154]	Macro-F1, etc.	RCV1, NYTimes
Others	2019	TextGCN [155]	AAAI	SA, TL	[156]	Accuracy	20NG, Ohsumed, R52, etc.
	2019	SGC [157]	ICML	NC, TL, SA	[158]	Accuracy, Time	20NG, R8, Ohsumed, etc.
	2019	Huang et al. [159]	EMNLP	NC, TL	[160]	Accuracy	R8, R52, Ohsumed
	2019	Peng et al. [161]	arXiv	NC, TL	-	Micro-F1, etc.	RCV1, EUR-Lex, etc.
	2020	TextTNG [162]	ACL	SA, NC, TL	[163]	Accuracy	MR, R8, R52, Ohsumed
	2020	TensorGCN [164]	AAAI	SA, NC, TL	[165]	Accuracy	20NG, R8, R52, Ohsumed, MR
	2020	MAGNET [166]	ICAART	TL	[167]	Micro-F1, HL	Reuters, RCV1-V2, etc.
	2017	Miyato et al. [85]	ICLR	SA, NC	[168]	Error Rate	IMDB, RCV1, et al.
	2018	TMN [169]	EMNLP	TL	-	Accuracy, F1	Snippets, Twitter, et al.
	2019	Zhang et al. [170]	NAACL	TL, NC	[171]	Accuracy	DBpedia, 20NG

Trans: Transformer. Time: training time.

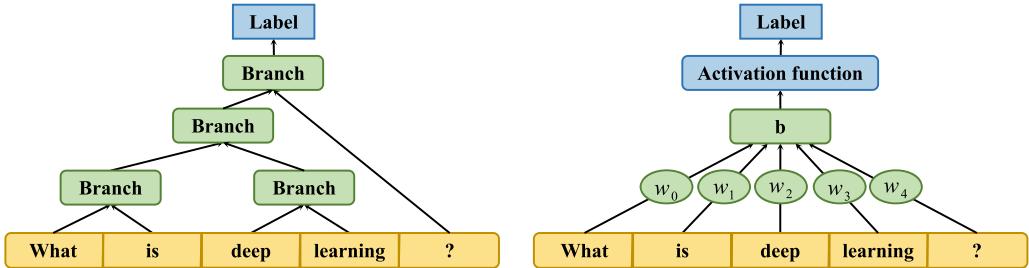


Fig. 6. The architecture of ReNN (left) and the architecture of MLP (right).

[155, 177] to capture structural information in the text, which cannot be replaced by other methods. Here, we classify DNNs by structure and discuss some representative models in detail.

2.2.1 ReNN-based Methods. Traditional models cost lots of time on design features for each task. Furthermore, in the case of deep learning, the meaning of “word vectors” is different: each input word is associated with a fixed-length vector whose values are either drawn at random or derived from a previous traditional process, thus forming a matrix L called word embedding matrix which represents the vocabulary words in a small latent semantic space, of generally 50 to 300 dimensions. The **Recursive Neural Network (ReNN)** [173] can automatically learn the semantics of text recursively and the syntax tree structure without feature design, as shown in Figure 6. We give an example of ReNN based models. First, each word of input text is taken as the leaf node of the model structure. Then all nodes are combined into parent nodes using a weight matrix. The weight matrix is shared across the whole model. Each parent node has the same dimension with all leaf nodes. Finally, all nodes are recursively aggregated into a parent node to represent the input text to predict the label.

ReNN-based models improve performance compared with traditional models and save on labor costs due to excluding feature designs used for different text classification tasks. The **Recursive AutoEncoder (RAE)** [60] is used to predict the distribution of sentiment labels for each input sentence and learn the representations of multi-word phrases. To learn compositional vector representations for each input text, the **Matrix-Vector Recursive Neural Network (MV-RNN)** [62] introduces a ReNN model to learn the representation of phrases and sentences. It allows that the length and type of input texts are inconsistent. MV-RNN allocates a matrix and a vector for each node on the constructed parse tree. Furthermore, the **Recursive Neural Tensor Network (RNTN)** [64] is proposed with a tree structure to capture the semantics of sentences. It inputs phrases with different lengths and represents the phrases by parse trees and word vectors. The vectors of higher nodes on the parse tree are estimated by the equal tensor-based composition function. For RNTN, the time complexity of building the textual tree is high, and expressing the relationship between documents is complicated within a tree structure. The performance is usually improved, with the depth being increased for DNNs. Therefore, Irsøy et al. [66] propose a **Deep Recursive Neural Network (DeepReNN)**, which stacks multiple recursive layers. It is built by binary parse trees and learns distinct perspectives of compositionality in language.

2.2.2 MLP-based Methods. A **MultiLayer Perceptron (MLP)** [172], sometimes colloquially called “vanilla” neural network, is a simple neural network structure that is used for capturing features automatically. As shown in Figure 6, we show a three-layer MLP model. It contains an input layer, a hidden layer with an activation function in all nodes, and an output layer. Each node

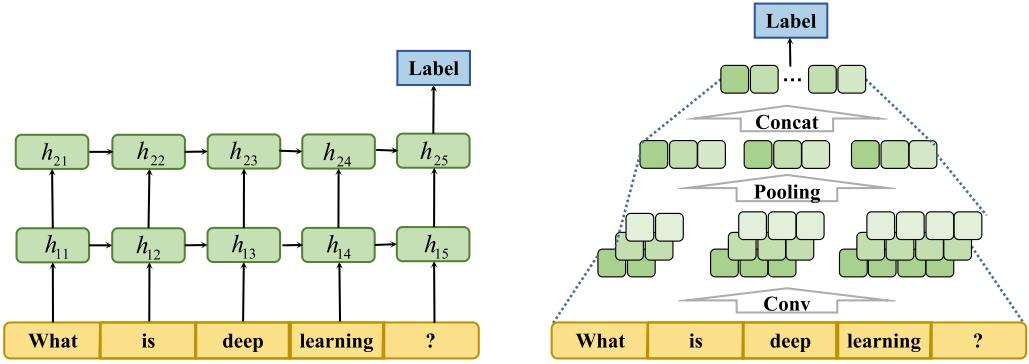


Fig. 7. The RNN based model (left) and the CNN based model (right).

connects with a certain weight w_i . It treats each input text as a bag of words and achieves high performance on many text classification benchmarks comparing with traditional models.

There are some MLP-based methods proposed by some research groups for text classification tasks. The **Paragraph Vector (Paragraph-Vec)** [67] is the most popular and widely used method, which is similar to the **Continuous Bag-Of-Words (CBOW)** [23]. It gets fixed-length feature representations of texts with various input lengths by employing unsupervised algorithms. Comparing with CBOW, it adds a paragraph token mapped to the paragraph vector by a matrix. The model predicts the fourth word by the connection or average of this vector to the three contexts of the word. Paragraph vectors can be used as a memory for paragraph themes and are used as a paragraph function and inserted into the prediction classifier.

2.2.3 RNN-based Methods. The **Recurrent Neural Network (RNN)** [173] is broadly used for capturing long-range dependency through recurrent computation. The RNN language model learns historical information, considering the location information among all words suitable for text classification tasks. We show an RNN model for text classification with a simple sample, as shown in Figure 7. Firstly, each input word is represented by a specific vector using a word embedding technology. Then, the embedding word vectors are fed into RNN cells one by one. The output of RNN cells are the same dimension with the input vector and are fed into the next hidden layer. The RNN shares parameters across different parts of the model and has the same weights of each input word. Finally, the label of input text can be predicted by the last output of the hidden layer.

To diminish the time complexity of the model and capture contextual information, Liu et al. [79] introduce a model for catching the semantics of long texts. It is a biased model that parsed the text one by one, making the following inputs profit over the former and decreasing the semantic efficiency of capturing the whole text. For modeling topic labeling tasks with long input sequences, TopicRNN [83] is proposed. It captures the dependencies of words in a document via latent topics and uses RNNs to capture local dependencies and latent topic models for capturing global semantic dependencies. **Virtual Adversarial Training (VAT)** [178] is a useful regularization method applicable to semi-supervised learning tasks. Miyato et al. [85] apply adversarial and virtual adversarial training text and employ the perturbation into word embedding rather than the original input text. The model improves the quality of the word embedding and is not easy to overfit during training. Capsule network [179] captures the relationships between features using dynamic routing between capsules comprised of a group of neurons in a layer. Wang et al. [87] propose an RNN-Capsule model with a simple capsule structure for the sentiment classification task.

In the backpropagation process of RNN, the weights are adjusted by gradients, calculated by continuous multiplications of derivatives. If the derivatives are extremely small, it may cause a gradient vanishing problem by continuous multiplications. **Long Short-Term Memory (LSTM)** [180], the improvement of RNN, effectively alleviates the gradient vanishing problem. It is composed of a cell to remember values on arbitrary time intervals and three gate structures to control information flow. The gate structures include input gates, forget gates, and output gates. The LSTM classification method can better capture the connection among context feature words, and use the forgotten gate structure to filter useless information, which is conducive to improving the total capturing ability of the classifier. Tree-LSTM [1] extends the sequence of LSTM models to the tree structure. The whole subtree with little influence on the result can be forgotten through the LSTM forgetting gate mechanism for the Tree-LSTM model.

Natural Language Inference (NLI) [181] predicts whether one text's meaning can be deduced from another by measuring the semantic similarity between each pair of sentences. To consider other granular matchings and matchings in the reverse direction, Wang et al. [182] propose a model for the NLI task named **Bilateral Multi-Perspective Matching (BiMPM)**. It encodes input sentences by the BiLSTM encoder. Then, the encoded sentences are matched in two directions. The results are aggregated in a fixed-length matching vector by another BiLSTM layer. Finally, the result is evaluated by a fully connected layer.

2.2.4 CNN-based Methods. **Convolutional Neural Networks (CNNs)** [18] are proposed for image classification with convolving filters that can extract features of pictures. Unlike RNN, CNN can simultaneously apply convolutions defined by different kernels to multiple chunks of a sequence. Therefore, CNNs are used for many NLP tasks, including text classification. For text classification, the text requires being represented as a vector similar to the image representation, and text features can be filtered from multiple angles, as shown in Figure 7. Firstly, the word vectors of the input text are spliced into a matrix. The matrix is then fed into the convolutional layer, which contains several filters with different dimensions. Finally, the result of the convolutional layer goes through the pooling layer and concatenates the pooling result to obtain the final vector representation of the text. The category is predicted by the final vector.

To try using CNN for the text classification task, an unbiased model of convolutional neural networks is introduced by Kim, called TextCNN [17]. It can better determine discriminative phrases in the max-pooling layer with one layer of convolution and learn hyperparameters except for word vectors by keeping word vectors static. Training only on labeled data is not enough for data-driven deep models. Therefore, some researchers consider utilizing unlabeled data. Johnson et al. [183] propose a CNN model based on two-view semi-supervised learning for text classification, which first uses unlabeled data to train the embedding of text regions and then labeled data. DNNs usually have better performance, but it increases the computational complexity. Motivated by this, a **Deep Pyramid Convolutional Neural Network (DPCNN)** [98] is proposed, with a little more computational accuracy, increasing by raising the network depth. The DPCNN is more specific than **Residual Network (ResNet)** [184], as all the shortcuts are exactly simple identity mappings without any complication for dimension matching.

According to the minimum embedding unit of text, embedding methods are divided into character-level, word-level, and sentence-level embedding. Character-level embeddings can settle **Out-Of-Vocabulary (OOV)** [185] words. Word-level embeddings learn the syntax and semantics of the words. Moreover, sentence-level embedding can capture relationships among sentences. Motivated by these, Nguyen et al. [186] propose a deep learning method based on a dictionary, increasing information for word-level embeddings through constructing semantic rules and deep CNN for character-level embeddings. Adams et al. [187] propose a character-level CNN model,

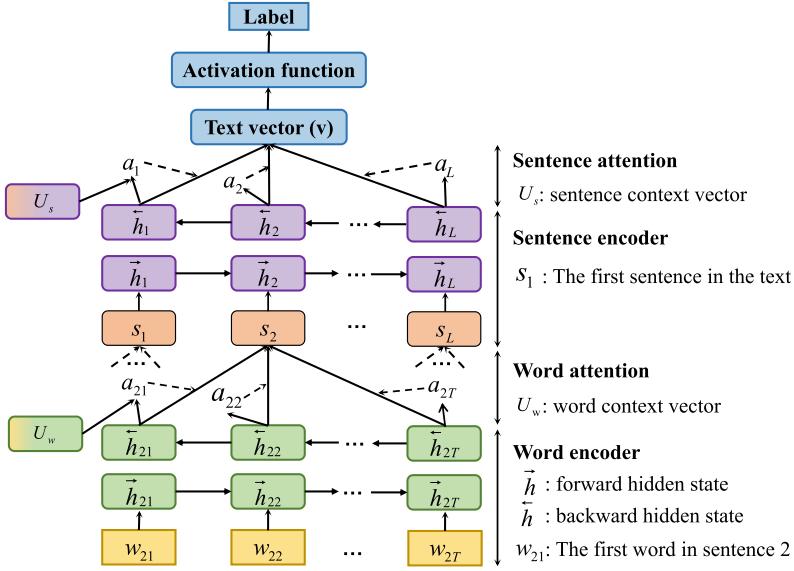


Fig. 8. The architecture of hierarchical attention network (HAN) [108].

called MGTC, to classify multi-lingual texts written. TransCap [188] is proposed to encapsulate the sentence-level semantic representations into semantic capsules and transfer document-level knowledge.

RNN based models capture the sequential information to learn the dependency among input words, and CNN based models extract the relevant features from the convolution kernels. Thus some works study the fusion of the two methods. BLSTM-2DCNN [77] integrates a **Bidirectional LSTM (BiLSTM)** with two-dimensional max pooling. It uses a 2D convolution to sample more meaningful information of the matrix and understands the context better through BiLSTM. Moreover, Xue et al. [189] propose MTNA, a combination of BiLSTM and CNN layers, to solve aspect category classification and aspect term extraction tasks.

2.2.5 Attention-based Methods. CNN and RNN provide excellent results on tasks related to text classification. However, these models are not intuitive enough for poor interpretability, especially in classification errors, which cannot be explained due to the non-readability of hidden data. The attention-based methods are successfully used in the text classification. Bahdanau et al. [190] first propose an attention mechanism that can be used in machine translation. Motivated by this, Yang et al. [108] introduce the **Hierarchical Attention Network (HAN)** to gain better visualization by employing the extremely informational components of a text, as shown in Figure 8. HAN includes two encoders and two levels of attention layers. The attention mechanism lets the model pay different attention to specific inputs. It aggregates essential words into sentence vectors firstly and then aggregates vital sentence vectors into text vectors. The attention-based models can learn how much contribution of each word and sentence for the classification judgment, which is beneficial for applications and analysis through the two levels of attention.

The attention mechanism can improve the performance with interpretability for text classification, which makes it popular. There are some other works based on attention. LSTMN [112] is proposed to process text step by step from left to right and does superficial reasoning through memory and attention. BI-Attention [110] is designed for cross-lingual text classification to catch bilingual long-distance dependencies. Hu et al. [191] propose an attention mechanism based on

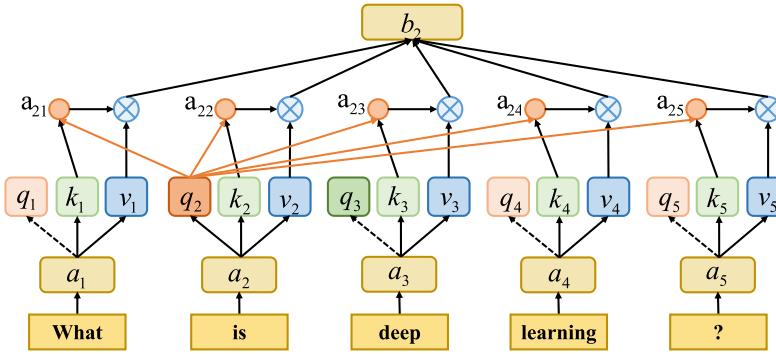


Fig. 9. An example of self-attention for calculating output vector b_2 .

category attributes for solving the imbalance of the number of various charges which contain few-shot charges. HAPN [124] is presented for few-shot text classification.

Self-attention [192] captures the weight distribution of words in sentences by constructing K, Q, and V matrices among sentences that can capture long-range dependencies on text classification. We give an example for self-attention, as shown in Figure 9. Each input word vector a_i can be represented as three n-dimensional vectors, including q_i , k_i , and v_i . After self-attention, the output vector b_i can be represented as $\sum_j softmax(a_{ij})v_j$ and $a_{ij} = q_i \cdot k_j / \sqrt{n}$. All output vectors can be parallelly computed. Lin et al. [114] used source token self-attention to explore the weight of every token to the entire sentence in the sentence representation task. To capture long-range dependencies, **Bi-directional Block Self-Attention Network (Bi-BLoSAN)** [120] uses an intra-block **Self-Attention Network (SAN)** to every block split by sequence and an inter-block SAN to the outputs.

Aspect-Based Sentiment Analysis (ABSA) [31, 193] breaks down a text into multiple aspects and allocates each aspect a sentiment polarity. The sentiment polarity can be divided into three types: positive, neutral, and negative. Some attention-based models are proposed to identify the fine-grained opinion polarity towards a specific aspect for aspect-based sentiment tasks. ATAE-LSTM [194] can concentrate on different parts of each sentence according to the input through the attention mechanisms. MGAN [195] presents a fine-grained attention mechanism with a coarse-grained attention mechanism to learn the word-level interaction between context and aspect.

To catch the complicated semantic relationship among each question and candidate answers for the QA task, Tan et al. [196] introduce CNN and RNN and generate answer embeddings by using a simple one-way attention mechanism affected through the question context. The attention captures the dependence among the embeddings of questions and answers. Extractive QA can be seen as the text classification task. It inputs a question and multiple candidates answers and classifies every candidate answer to recognize the correct answer. Furthermore, AP-BiLSTM [197] with a two-way attention mechanism can learn the weights between the question and each candidate answer to obtain the importance of each candidate answer to the question.

2.2.6 Pre-trained Methods. Pre-trained language models [198] effectively learn global semantic representation and significantly boost NLP tasks, including text classification. It generally uses unsupervised methods to mine semantic knowledge automatically and then constructs pre-training targets so that machines can learn to understand semantics.

As shown in Figure 10, we give differences in the model architectures among the **Embedding from Language Model (ELMo)** [118], OpenAI GPT [199], and BERT [19]. ELMo [118] is a deep

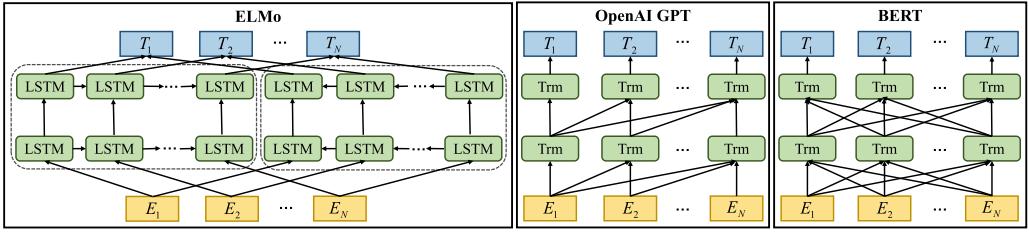


Fig. 10. Differences in pre-trained model architectures [19], including BERT, OpenAI GPT, and ELMo. E_i represents embedding of i th input. Trm represents the transformer block. T_i represents predicted tag of i th input.

contextualized word representation model, which is readily integrated into models. It can model complicated characteristics of words and learn different representations for various linguistic contexts. It learns each word embedding according to the context words with the bi-directional LSTM. GPT [199] employs supervised fine-tuning and unsupervised pre-training to learn general representations that transfer with limited adaptation to many NLP tasks. Furthermore, the domain of the target dataset does not need to be similar to the domain of unlabeled datasets. The training procedure of the GPT algorithm usually includes two stages. Firstly, the initial parameters of a neural network model are learned by a modeling objective on the unlabeled dataset. We can then employ the corresponding supervised objective to accommodate these parameters for the target task. To pre-train deep bidirectional representations from the unlabeled text through joint conditioning on both left and right context in every layer, the BERT model [19], proposed by Google, significantly improves performance on NLP tasks, including text classification. BERT applies the bi-directional encoder designed to pre-train the bi-directional representation of depth by jointly adjusting the context in all layers. It can utilize contextual information when predicting which words are masked. It is fine-tuned by adding just an additional output layer to construct models for multiple NLP tasks, such as SA, QA, and machine translation. Comparing with these three models, ELMo is a feature-based method using LSTM, and BERT and OpenAI GPT are fine-tuning approaches using Transformer. Furthermore, ELMo and BERT are bidirectional training models and OpenAI GPT is training from left to right. Therefore, BERT gets a better result, which combines the advantages of ELMo and OpenAI GPT.

Transformer-based models can parallelize computation without considering the sequential information suitable for large scale datasets, making it popular for NLP tasks. Thus, some other works are used for text classification tasks and get excellent performance. RoBERTa [140], is an improved version of BERT, and adopts the dynamic masking method that generates the masking pattern every time with a sequence to be fed into the model. It uses more data for longer pre-training and estimates the influence of various essential hyperparameters and the size of training data. To be specific: (1) The training time is longer (a total of nearly 200,000 training set, nearly 1.6 billion training data have been seen), the batch size (8K) is larger, and the training data is more (30G Chinese training, including 300 million sentences and 10 billion words); (2) It removes the **next sentence prediction (NSP)** task; (3) It employs more extended training sequence; 4) It dynamically adjusts the masking mechanism and use the full word mask.

XLNet [138] is a generalized autoregressive pre-training approach. Unlike BERT, the denoising autoencoder with the mask is not used in the first stage, but the autoregressive LM is used. It maximizes the expected likelihood across the whole factorization order permutations to learn the bidirectional context. Furthermore, it can overcome the weaknesses of BERT by an autoregressive formulation and integrate ideas from Transformer-XL [200] into pre-training.

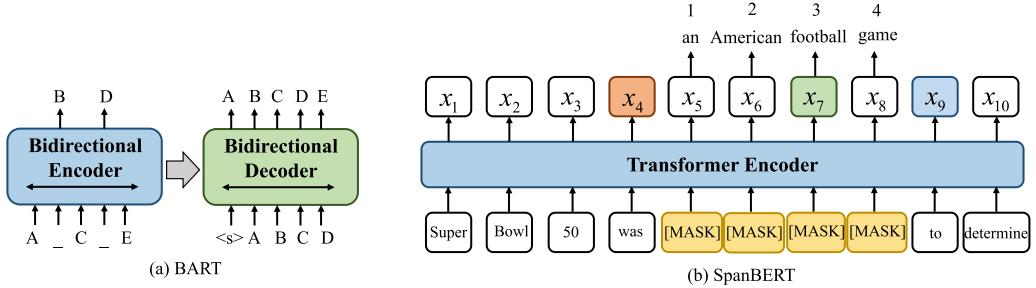


Fig. 11. The architecture of BART [202] and SpanBERT [203].

The BERT model has many parameters. In order to reduce the parameters, ALBERT [146] uses two-parameter simplification schemes. It reduces the fragmentation vector's length and shares parameters with all encoders. It also replaces the next sentence matching task with the next sentence order task and continuously blocks fragmentation. When the ALBERT model is pre-trained on a massive Chinese corpus, the parameters are fewer than BERT and also have better performance. In general, these methods adopt unsupervised objective functions for pre-training, including the next sentence prediction, masking technology, and permutation. These target functions based on the word prediction demonstrate a strong ability to learn the word dependence and semantic structure [201].

BART [202] is a denoising autoencoder based on the Seq2Seq model, as shown in Figure 11(a). The pre-training of BART consists of two steps. Firstly, it uses a noise function to destroy the text. Secondly, the Seq2Seq model is used to reconstruct the original text. In various noise methods, the model obtains optimal performance by randomly shuffling the order of the original sentence and then using the first new text filling method. The new text filling method is replacing the text fragment with a single mask token. It uses only a specific masked token to indicate that a token is masked.

SpanBERT [203] is specially designed to better represent and predict spans of text, as shown in Figure 11(b). It optimizes BERT from three aspects and achieves good results in multiple tasks such as QA. The specific optimization is embodied in three aspects. Firstly, the span mask scheme is proposed to mask a continuous paragraph of text randomly. Secondly, **Span Boundary Objective (SBO)** is added to predict span by the token next to the span boundary to get the better performance to finetune stage. Thirdly, the NSP pre-training task is removed.

ERNIE [204] is based on the method of knowledge enhancement. It learns the semantic relations in the real world by modeling the prior semantic knowledge such as entity concepts in massive datasets. Specifically, ERNIE enables the model to learn the semantic representation of complete concepts by masking semantic units such as words and entities. It mainly consists of a Transformer encoder and task embedding. In the Transformer encoder, the context information of each token is captured by the self-attention mechanism, and the context representation is generated for embedding. Task embedding is used for tasks with different characteristics.

2.2.7 GNN-based Methods. The DNN models like CNN get great performance on regular structure, but not for arbitrarily structured graphs. Some researchers study how to expand on arbitrarily structured graphs [205, 206]. With the increasing attention of **Graph Neural Networks (GNNs)**, GNN-based models [207, 208] obtain excellent performance by encoding syntactic structure of sentences on semantic role labeling task [209], relation classification task [210], and machine translation task [211]. It turns text classification into a graph node classification task. We show a GCN model for text classification with four input texts, as shown in Figure 12. Firstly, the four input

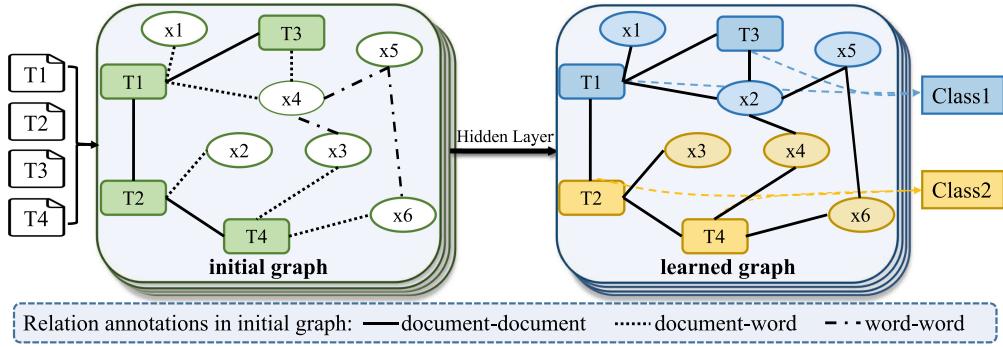


Fig. 12. The GNN-based model. The different initial graph depends on how the graph is designed. We give an example to establish edges between documents and documents, documents and sentences, and words to words.

texts $T = [T_1, T_2, T_3, T_4]$ and the words $X = [x_1, x_2, x_3, x_4, x_5, x_6]$ in the text, defined as nodes, are constructed into the graph structures. The graph nodes are connected by bold black edges, which indicates document-word edges and word-word edges. The weight of each word-word edge usually means their co-occurrence frequency in the corpus. Then, the words and texts are represented through the hidden layer. Finally, the label of all input texts can be predicted by the graph.

The GNN-based models can learn the syntactic structure of sentences, making some researchers study using GNN for text classification. DGCNN [153] is a graph-CNN converting text to graph-of-words, having the advantage of learning different levels of semantics with CNN models. Yao et al. [155] propose the **Text Graph Convolutional Network (TextGCN)**, which builds a heterogeneous word text graph for a whole dataset and captures global word co-occurrence information. To enable GNN-based models to underpin online testing, Huang et al. [159] build graphs for each text with global parameter sharing, not a corpus-level graph structure, to help preserve global information and reduce the burden. TextING [162] builds individual graphs for each document and learns text-level word interactions by GNN to effectively produce embeddings for obscure words in the new text.

Graph ATtention network (GAT) [212] employs masked self-attention layers by attending over its neighbors. Thus, some GAT-based models are proposed to compute the hidden representations of each node. The **Heterogeneous Graph ATtention networks (HGAT)** [213] with a dual-level attention mechanism learns the importance of different neighboring nodes and node types in the current node. The model propagates information on the graph and captures the relations to address the semantic sparsity for semi-supervised short text classification. MAGNET [166] is proposed to capture the correlation among the labels based on GATs, which learns the crucial dependencies between the labels and generates classifiers by a feature matrix and a correlation matrix.

Event Prediction (EP) can be divided into generated event prediction and selective event prediction (also known as script event prediction). EP, referring to scripted event prediction in this review, infers the subsequent event according to the existing event context. Unlike other text classification tasks, texts in EP are composed of a series of sequential subevents. Extracting features of the relationship among such subevents is of critical importance. SGNN [214] is proposed to model event interactions and learn better event representations by constructing an event graph to utilize the event network information better. The model makes full use of dense event connections for the EP task.

2.2.8 Others. In addition to all the above models, there are some other individual models. Here we introduce some exciting models.

Siamese Neural Network. The siamese neural network [215] is also called a **twin neural network (Twin NN)**. It utilizes equal weights while working in tandem using two distinct input vectors to calculate comparable output vectors. Mueller et al. [216] present a siamese adaptation of the LSTM network comprised of couples of variable-length sequences. The model is employed to estimate the semantic similarity among texts, exceeding carefully handcrafted features and proposed neural network models of higher complexity. The model further represents text employing neural networks whose inputs are word vectors learned separately from a vast dataset. To settle unbalanced data classification in the medical domain, Jayadeva et al. [217] use a Twin NN model to learn from enormous unbalanced corpora. The objective functions achieve the Twin SVM approach with non-parallel decision boundaries for the corresponding classes, and decrease the Twin NN complexity, optimizing the feature map to better discriminate among classes.

Virtual Adversarial Training (VAT). Deep learning methods require many extra hyperparameters, which increase the computational complexity. VAT [218], regularization based on local distributional smoothness can be used in semi-supervised tasks, requires only some hyperparameters, and can be interpreted directly as robust optimization. Miyato et al. [85] use VAT to effectively improve the robustness and generalization ability of the model and word embedding performance.

Reinforcement Learning (RL). RL learns the best action in a given environment through maximizing cumulative rewards. Zhang et al. [219] offer an RL approach to establish structured sentence representations via learning the structures related to tasks. The model has **Information Distilled LSTM (ID-LSTM)** and **Hierarchical Structured LSTM (HS-LSTM)** representation models. The ID-LSTM learns the sentence representation by choosing essential words relevant to tasks, and the HS-LSTM is a two-level LSTM for modeling sentence representation.

Memory Networks. Memory networks [220] learn to combine the inference components and the long-term memory component. Li et al. [221] use two LSTMs with extended memories and neural memory operations for jointly handling the extraction tasks of aspects and opinions via memory interactions. **Topic Memory Networks (TMN)** [169] is an end-to-end model that encodes latent topic representations indicative of class labels.

QA Style for Sentiment Classification Task. It is an interesting attempt to treat the sentiment classification task as a QA task. Shen et al. [222] create a high-quality annotated corpus. A three-stage hierarchical matching network was proposed to consider the matching information between questions and answers.

External Commonsense Knowledge. Due to the insufficient information of the event itself to distinguish the event for the EP task, Ding et al. [223] consider that the event extracted from the original text lacked common knowledge, such as the intention and emotion of the event participants. The model improves the effect of stock prediction, EP, and so on.

Quantum Language Model. In the quantum language model, the words and dependencies among words are represented through fundamental quantum events. Zhang et al. [224] design a quantum-inspired sentiment representation method to learn both the semantic and the sentiment information of subjective text. By inputting density matrices to the embedding layer, the performance of the model improves.

Summary. RNN computes sequentially and cannot be calculated in parallel. The shortcoming of RNN makes it more challenging to become mainstream in the current trend that models tend to have deeper and more parameters. CNN extracts features from text vectors through the convolution kernel. The number of features captured by the convolution kernel is related to its size. CNN is deep enough that, in theory, it can capture features at long distances. Due to insufficient optimization methods for parameters of the deep network and the loss of location information due to the pooling layer, the deeper layer does not bring significant improvement. Compared with RNN, CNN has parallel computing capability and can effectively retain location information for the improved version of CNN. Still, it has weak feature capture capability for long-distance. GNN builds a graph for text. When a valid graph structure is designed, the learned representation can better capture the structural information. Transformer treats the input text as a fully connected graph, with attention score weights on the edges. It is capable of parallel computing and is highly efficient in extracting features between different words by self-attention, solving short-term memory problems. However, the attention mechanism in Transformer is computation-heavy, especially when dealing with long sequences. Some improved models [146, 225] for computing complexity in Transformer have recently been proposed. Overall, Transformer is a better choice for text classification. Deep Learning consists of multiple hidden layers in a neural network with a higher level of complexity and can be trained on unstructured data. Deep learning can learn language features and master higher level and more abstract language features based on words and vectors. Deep learning architecture can learn feature representations directly from the input without too many manual interventions and prior knowledge. However, deep learning technology is a data-driven method that requires enormous data to achieve high performance. Although self-attention based models can bring some interpretability among words for DNNs, it is not enough comparing with traditional models to explain why and how it works well.

3 DATASETS AND EVALUATION METRICS

3.1 Datasets

The availability of labeled datasets for text classification has become the main driving force behind the fast advancement of this research field. In this section, we summarize the characteristics of these datasets in terms of domains and give an overview in Table 2, including the number of categories, average sentence length, the size of each dataset, related papers, data sources to access, and applications.

3.1.1 Sentiment Analysis (SA). SA is the process of analyzing and reasoning the subjective text within emotional color. It is crucial to get information on whether it supports a particular point of view from the text that is distinct from the traditional text classification that analyzes the objective content of the text. SA can be binary or multi-class. Binary SA is to divide the text into two categories, including positive and negative. Multi-class SA classifies text to multi-level or fine-grained labels. The SA datasets include **Movie Review (MR)** [226, 257], **Stanford Sentiment Treebank (SST)** [227], **Multi-Perspective Question Answering (MPQA)** [229, 258], **IMDB** [230], **Yelp** [231], **Amazon Reviews (AM)** [93], NLP&CC 2013 [111], **Subj** [250], **CR** [251], **SS-Twitter** [259], **SS-Youtube** [259], **SE1604** [260], and so on. Here we detail several datasets.

MR. The MR is a movie review dataset, each of which corresponds to a sentence. The corpus has 5,331 positive data and 5,331 negative data. 10-fold cross-validation by random splitting is commonly used to test MR.

SST. The SST is an extension of MR. It has two categories. SST-1 with fine-grained labels with five classes. It has 8,544 training texts and 2,210 test texts, respectively. Furthermore, SST-2 has

Table 2. Summary Statistics for the Datasets

Datasets	#C	#L	#N	Language	Related Papers	Sources	Applications
MR	2	20	10,662	English	[5, 17, 101, 155]	[226]	SA
SST-1	5	18	11,855	English	[17, 64] [1, 2] [112]	[227]	SA
SST-2	2	19	9,613	English	[6, 17, 64] [19, 79]	[228]	SA
MPQA	2	3	10,606	English	[17, 60, 120]	[229]	SA
IMDB	2	294	50,000	English	[69] [108] [6] [79] [85]	[230]	SA
Yelp.P	2	153	598,000	English	[93, 98]	[231]	SA
Yelp.F	5	155	700,000	English	[93, 98, 108]	[231]	SA
Amz.P	2	91	4,000,000	English	[93, 122]	[232]	SA
Amz.F	5	93	3,650,000	English	[93, 108, 122]	[232]	SA
Twitter	3	19	11,209	English	[5] [100]	[233]	SA
NLP&CC 2013	2	-	115,606	Multi-language	[110]	[111]	SA
20NG	20	221	18,846	English	[72, 75, 106, 155, 157]	[34]	NC
AG News	4	45/7	127,600	English	[98, 100] [101, 138]	[234]	NC
R8	8	66	7,674	English	[155, 157] [159]	[235]	NC
R52	52	70	9,100	English	[155, 157] [159]	[235]	NC
Sogou	6	578	510,000	Chinese	[93]	[236]	NC
Newsgroup	20	-	18,846	English	[237]	[237]	NC
DBpedia	14	55	630,000	English	[85, 93, 98, 136]	[238]	TL
Ohsuemed	23	136	7,400	English	[155, 157, 159]	[239]	TL
YahooA	10	112	1,460,000	English	[93, 108]	[93]	TL
EUR-Lex	3,956	1,239	19,314	English	[96] [134, 161] [134]	[240]	TL
Amazon670K	670	244	643,474	English	[103, 122]	[241]	TL
Google news	152	6	11,109	English	[3, 242, 243]	[242]	TL
TweetSet 2011-2012	89	-	2,472	English	[242, 243]	[242]	TL
TweetSet 2011-2015	269	8	30,322	English	[3, 4]	[4]	TL
Bing	4	20	34,871	English	[100]	[244]	TL
Fudan	20	2981	18,655	Chinese	[72]	[245]	TL
SQuAD	-	5,000	5,570	English	[118, 118, 140, 146]	[246]	QA
TREC-QA	-	1,162	68	English	[197]	[247]	QA
TREC	6	10	5,952	English	[5, 6, 17] [100]	[248]	QA
WikiQA	-	873	243	English	[197, 249]	[249]	QA
Subj	2	23	10,000	English	[17, 79, 101]	[250]	QA
CR	2	19	3,775	English	[17, 101]	[251]	QA
Reuters	90	168	10,788	English	[101, 166]	[252]	ML
Reuters10	10	168	9,979	English	[253]	[254]	ML
RCV1	103	240	807,595	English	[75, 103, 134, 153]	[255]	ML
RCV1-V2	103	124	804,414	English	[116, 166]	[256]	ML
AAPD	54	163	55,840	English	[116, 166]	[117]	ML

C: Number of target classes. L: Average sentence length. N: Dataset size.

9,613 texts with binary labels being partitioned into 6,920 training texts, 872 development texts, and 1,821 testing texts.

MPQA. The MPQA is an opinion dataset. It has two class labels and also an MPQA dataset of opinion polarity detection sub-tasks. MPQA includes 10,606 sentences extracted from news articles from various news sources. It should be noted that it contains 3,311 positive texts and 7,293 negative texts without labels of each text.

IMDB reviews. The IMDB review is developed for binary sentiment classification of film reviews with the same amount in each class. It can be separated into training and test groups on average, by 25,000 comments per group.

Yelp reviews. The Yelp review is summarized from the Yelp Dataset Challenges in 2013, 2014, and 2015. This dataset has two categories. Yelp-2 of these were used for negative and positive

emotion classification tasks, including 560,000 training texts and 38,000 test texts. Yelp-5 is used to detect fine-grained affective labels with 650,000 training and 50,000 test texts in all classes.

AM. The AM is a popular corpus formed by collecting Amazon website product reviews [232]. This dataset has two categories. The Amazon-2 with two classes includes 3,600,000 training sets and 400,000 testing sets. Amazon-5, with five classes, includes 3,000,000 and 650,000 comments for training and testing.

3.1.2 News Classification (NC). News content is one of the most crucial information sources which has a critical influence on people. The NC system facilitates users to get vital knowledge in real-time. News classification applications mainly encompass recognizing news topics and recommending related news according to user interest. The news classification datasets include **20 Newsgroups (20NG)** [34], **AG News (AG)** [93, 234], **R8** [235], **R52** [235], Sogou News (Sogou) [136], and so on. Here we detail several datasets.

20NG. The 20NG is a newsgroup text dataset. It has 20 categories with the same number of each category and includes 18,846 texts.

AG. The AG News is a search engine for news from academia, choosing the four largest classes. It uses the title and description fields of each news. AG contains 120,000 texts for training and 7,600 texts for testing.

R8 and R52. R8 and R52 are two subsets which are the subset of Reuters [252]. R8 has 8 categories, divided into 2,189 test files and 5,485 training courses. R52 has 52 categories, split into 6,532 training files and 2,568 test files.

Sogou. The Sogou combines two datasets, including SogouCA and SogouCS news sets. The label of each text is the domain name in the URL.

3.1.3 Topic Labeling (TL). The topic analysis attempts to get the meaning of the text by defining the sophisticated text theme. The topic labeling is one of the essential components of the topic analysis technique, intending to assign one or more subjects for each document to simplify the topic analysis. The topic labeling datasets include DBpedia [238], Ohsumed [239], Yahoo answers (YahooA) [93], EUR-Lex [240], Amazon670K [241], Bing [244], Fudan [245], and PubMed [261]. Here we detail several datasets.

DBpedia. The DBpedia is a large-scale multi-lingual knowledge base generated using Wikipedia's most ordinarily used infoboxes. It publishes DBpedia each month, adding or deleting classes and properties in every version. DBpedia's most prevalent version has 14 classes and is divided into 560,000 training data and 70,000 test data.

Ohsumed. The Ohsumed belongs to the MEDLINE database. It includes 7,400 texts and has 23 cardiovascular disease categories. All texts are medical abstracts and are labeled into one or more classes.

YahooA. The YahooA is a topic labeling task with 10 classes. It includes 140,000 training data and 5,000 test data. All texts contains three elements, being question titles, question contexts, and best answers, respectively.

3.1.4 Question Answering (QA). The QA task can be divided into two types: the extractive QA and the generative QA. The extractive QA gives multiple candidate answers for each question to choose which one is the right answer. Thus, the text classification models can be used for the extractive QA task. The QA discussed in this paper is all extractive QA. The QA system can apply the text classification model to recognize the correct answer and set others as candidates. The question answering datasets include **Stanford Question Answering Dataset (SQuAD)** [246], **TREC-QA** [248], **WikiQA** [249], **Subj** [250], **CR** [251], **MS MARCO** [262], and **Quora** [263]. Here we detail several datasets.

SQuAD. The SQuAD is a set of question and answer pairs obtained from Wikipedia articles. The SQuAD has two categories. SQuAD1.1 contains 536 pairs of 107,785 Q&A items. SQuAD2.0 combines 100,000 questions in SQuAD1.1 with more than 50,000 unanswerable questions that crowd workers face in a form similar to answerable questions [264].

TREC-QA. The TREC-QA includes 5,452 training texts and 500 testing texts. It has two versions. TREC-6 contains 6 categories, and TREC-50 has 50 categories.

WikiQA. The WikiQA dataset includes questions with no correct answer, which needs to evaluate the answer.

MS MARCO. The MS MARCO contains questions and answers. The questions and part of the answers are sampled from actual web texts by the Bing search engine. Others are generative. It is used for developing generative QA systems released by Microsoft.

3.1.5 Natural Language Inference (NLI). NLI is used to predict whether the meaning of one text can be deduced from another. Paraphrasing is a generalized form of NLI. It uses the task of measuring the semantic similarity of sentence pairs to decide whether one sentence is the interpretation of another. The NLI datasets include **Stanford Natural Language Inference (SNLI)** [181], **Multi-Genre Natural Language Inference (MNLI)** [265], **Sentences Involving Compositional Knowledge (SICK)** [266], **Microsoft Research Paraphrase (MSRP)** [267], **Semantic Textual Similarity (STS)** [268], **Recognising Textual Entailment (RTE)** [269], SciTail [270], etc. Here we detail several of the primary datasets.

SNLI. The SNLI is generally applied to NLI tasks. It contains 570,152 human-annotated sentence pairs, including training, development, and test sets, which are annotated with three categories: neutral, entailment, and contradiction.

MNLI. The MNLI is an expansion of SNLI, embracing a broader scope of written and spoken text genres. It includes 433,000 sentence pairs annotated by textual entailment labels.

SICK. The SICK contains almost 10,000 English sentence pairs. It consists of neutral, entailment, and contradictory labels.

MSRP. The MSRP consists of sentence pairs, usually for the text-similarity task. Each pair is annotated by a binary label to discriminate whether they are paraphrases. It respectively includes 1,725 training and 4,076 test sets.

3.1.6 Multi-Label (ML) Datasets. In multi-label classification, an instance has multiple labels, and each label can only take one of the multiple classes. There are many datasets based on multi-label text classification. It includes **Reuters** [252], **Reuters Corpus Volume I (RCV1)** [255], **RCV1-2K** [255], **Arxiv Academic Paper Dataset (AAPD)** [117], **Patent**, **Web of Science (WOS-11967)** [271], **AmazonCat-13K** [272], **BlurbGenreCollection (BGC)** [273], etc. Here we detail several datasets.

Reuters. The Reuters is a popularly used dataset for text classification from Reuters financial news services. It has 90 training classes, 7,769 training texts, and 3,019 testing texts, containing multiple labels and single labels. There are also some Reuters sub-sets of data, such as R8, BR52, RCV1, and RCV1-v2.

RCV1 and RCV1-2K. The RCV1 is collected from Reuters News articles from 1996-1997, which is human-labeled with 103 categories. It consists of 23,149 training and 784,446 testing texts, respectively. The RCV1-2K dataset has the same features as the RCV1. However, the label set of RCV1-2K has been expanded with some new labels. It contains 2,456 labels.

AAPD. The AAPD is a large dataset in the computer science field for the multi-label text classification from website.¹ It has 55,840 papers, including the abstract and the corresponding subjects

¹<https://arxiv.org/>.

Table 3. The Notations Used in Evaluation Metrics

Notations	Descriptions
TP	true positive
FP	false positive
TN	true negative
FN	false negative
TP_t	true positive of the t th label on a text
FP_t	false positive of the t th label on a text
TN_t	true negative of the t th label on a text
FN_t	false negative of the t th label on a text
S	label set of all samples
Q	the number of predicted labels on each text

with 54 labels in total. The aim is to predict the corresponding subjects of each paper according to the abstract.

Patent Dataset. The Patent Dataset is obtained from USPTO,² which is a patent system granting U.S. patents containing textual details such title and abstract. It contains 100,000 US patents awarded in the real-world with multiple hierarchical categories.

WOS-11967. The WOS-11967 is crawled from the Web of Science, consisting of abstracts of published papers with two labels for each example. It is more shallow, but significantly broader, with fewer classes in total.

3.1.7 Others. There are some datasets for other applications, such as SemEval-2010 Task 8 [274], ACE 2003-2004 [275], TACRED [276], and NYT-10 [277], FewRel [278], **Dialog State Tracking Challenge 4 (DSTC 4)** [279], ICSI **Meeting Recorder Dialog Act (MRDA)** [280], and **Switchboard Dialog Act (SwDA)** [281], and so on.

3.2 Evaluation Metrics

In terms of evaluating text classification models, accuracy and F1 score are the most used to assess the text classification methods. Later, with the increasing difficulty of classification tasks or the existence of some particular tasks, the evaluation metrics are improved. For example, evaluation metrics such as $P@K$ and $Micro-F1$ are used to evaluate multi-label text classification performance, and MRR is usually used to estimate the performance of QA tasks. In Table 3, we give the notations used in evaluation metrics.

3.2.1 Single-label Metrics. Single-label text classification divides the text into one of the most likely categories applied in NLP tasks such as QA, SA, and dialogue systems [7]. For single-label text classification, one text belongs to just one catalog, making it possible not to consider the relations among labels. Here, we introduce some evaluation metrics used for single-label text classification tasks.

Accuracy and ErrorRate. The Accuracy and ErrorRate are the fundamental metrics for a text classification model. The *Accuracy* and *ErrorRate* are respectively defined as

$$Accuracy = \frac{(TP + TN)}{N}, \quad (1)$$

$$ErrorRate = 1 - Accuracy = \frac{(FP + FN)}{N}. \quad (2)$$

²<https://www.uspto.gov/>.

Precision, Recall and F1. These are vital metrics utilized for unbalanced test sets, regardless of the standard type and error rate. For example, most of the test samples have a class label. $F1$ is the harmonic average of *Precision* and *Recall*. *Precision*, *Recall*, and $F1$ as defined

$$Precision = \frac{TP}{TP + FP}, \quad Recall = \frac{TP}{TP + FN}, \quad (3)$$

$$F1 = \frac{2Precision \times Recall}{Precision + Recall}. \quad (4)$$

The desired results will be obtained when the accuracy, $F1$ and *Recall* value reach 1. On the contrary, when the values become 0, the worst result is obtained. For the multi-class classification problem, the precision and recall value of each class can be calculated separately, and then the performance of the individual and whole can be analyzed.

Exact Match (EM). The EM [29] is a metric for QA tasks, measuring the prediction that matches all the ground-truth answers precisely. It is the primary metric utilized on the SQuAD dataset.

Mean Reciprocal Rank (MRR). The MRR [282] is usually applied for assessing the performance of ranking algorithms on QA and **Information Retrieval (IR)** tasks. *MRR* is defined as

$$MRR = \frac{1}{Q} \sum_{i=1}^Q \frac{1}{rank(i)}, \quad (5)$$

where $rank(i)$ is the ranking of the ground-truth answer at answer i -th.

Hamming-Loss (HL). The HL [57] assesses the score of misclassified instance-label pairs where a related label is omitted or an unrelated label is predicted.

Among these single-label evaluation metrics, the Accuracy is the earliest metric that calculates the proportion of the sample size that is predicted correctly and is not considered whether the predicted sample is a positive or a negative sample. *Precision* calculates how many of the positive samples are actually positive, and the *Recall* calculates how many of the positive examples in the sample are predicted correctly. Furthermore, $F1$ is the harmonic average of them, which is the most commonly used evaluation metrics.

3.2.2 Multi-label Metrics. Compared with single-label text classification, multi-label text classification divides the text into multiple category labels, and the number of category labels is variable. These metrics are designed for single label text classification, which are not suitable for multi-label tasks. Thus, there are some metrics designed for multi-label text classification.

Micro-F1. The *Micro-F1* [283] is a measure that considers the overall accuracy and recall of all labels. The *Micro-F1* is defined as:

$$Micro-F1 = \frac{2P_t \times R_t}{P + R}, \quad (6)$$

where:

$$P = \frac{\sum_{t \in S} TP_t}{\sum_{t \in S} TP_t + FP_t}, \quad R = \frac{\sum_{t \in S} TP_t}{\sum_{t \in S} TP_t + FN_t}. \quad (7)$$

Macro-F1. The *Macro-F1* [283] calculates the average $F1$ of all labels. Unlike *Micro-F1*, which sets even weight to every example, *Macro-F1* sets the same weight to all labels in the average process. Formally, *Macro-F1* is defined as:

$$Macro-F1 = \frac{1}{S} \sum_{t \in S} \frac{2P_t \times R_t}{P_t + R_t}, \quad (8)$$

where:

$$P_t = \frac{TP_t}{TP_t + FP_t}, \quad R_t = \frac{TP_t}{TP_t + FN_t}. \quad (9)$$

In addition to the above evaluation metrics, there are some rank-based evaluation metrics for extreme multi-label classification tasks, including $P@K$ and $NDCG@K$.

Precision at Top K (P@K). The $P@K$ [96] is the precision at the top k . For $P@K$, each text has a set of \mathcal{L} ground truth labels $L_t = \{l_0, l_1, l_2 \dots, l_{\mathcal{L}-1}\}$, in order of decreasing probability $P_t = [p_0, p_1, p_2 \dots, p_{Q-1}]$. The precision at k is

$$P@K = \frac{1}{k} \sum_{j=0}^{\min(\mathcal{L}, k)-1} \text{rel}_{L_i}(P_t(j)), \quad (10)$$

$$\text{rel}_L(p) = \begin{cases} 1 & \text{if } p \in L \\ 0 & \text{otherwise} \end{cases}, \quad (11)$$

where \mathcal{L} is the number of ground truth labels or possible answers on each text and k is the number of selected labels on extreme multi-label text classification.

Normalized Discounted Cumulated Gains (NDCG@K). The $NDCG@K$ [96] is

$$NDCG@K = \frac{1}{IDCG(L_i, k)} \sum_{j=0}^{n-1} \frac{\text{rel}_{L_i}(P_t(j))}{\ln(j+1)}, \quad (12)$$

where $IDCG$ is ideal discounted cumulative gain and the particular rank position n is

$$n = \min(\max(|P_i|, |L_i|), k). \quad (13)$$

Among these multi-label evaluation metrics, *Micro-F1* considers the number of categories, which makes it suitable for the unbalanced data distribution. *Macro-F1* does not take into account the amount of data that treats each class equally. Thus, it is easily affected by the classes with high Recall and Precision. When the number of categories is large or extremely large, either $P@K$ or $NDCG@K$ is used.

4 QUANTITATIVE RESULTS

There are many differences between sentiment analysis, news classification, topic labeling, and natural language inference tasks, which can not be simplified modeled as a text classification task. In this section, we tabulate the performance of the main models given in their articles on classic datasets evaluated by classification accuracy, as shown in Table 4, including MR, SST-2, IMDB, Yelp.P, Yelp.F, Amazon.F, 20NG, AG, DBpedia, and SNLI.

We give the performance of NB and SVM algorithms from RNTN [64] due to that the less traditional text classification model has been an experiment on datasets in Table 4. The accuracy of NB and SVM are 81.8% and 79.4% on SST-2, respectively. We can see that, in the SST-2 data set with only two categories, the accuracy of NB is better than that of SVM. It may be because NB has relatively stable classification efficiency on new datasets. The performance is also stable on small datasets. Compared with the deep learning model, the performance of NB is lower. NB has the advantage of lower computational complexity than deep models. However, it requires manual classification features, making it difficult to migrate the model directly to other data sets.

For deep learning models, pre-trained models get better results on most datasets. It means that if you need to implement a text classification task, you can preferentially try pre-trained models,

Table 4. Accuracy of Text Classification Models on Primary Datasets Evaluated by Classification Accuracy (in Terms of Publication Year)

Model	Sentiment						News		Topic	NLI
	MR	SST-2	IMDB	Yelp.P	Yelp.F	Amz.F	20NG	AG	DBpedia	SNLI
NB [8]	-	81.80	-	-	-	-	-	-	-	-
SVM [42]	-	79.40	-	-	-	-	-	-	-	-
Tree-CRF [284]	77.30	-	-	-	-	-	-	-	-	-
RAE [60]	77.70	82.40	-	-	-	-	-	-	-	-
MV-RNN [62]	79.00	82.90	-	-	-	-	-	-	-	-
RNTN [64]	75.90	85.40	-	-	-	-	-	-	-	-
DCNN [5]		86.80	89.40	-	-	-	-	-	-	-
Paragraph-Vec [67]		87.80	92.58	-	-	-	-	-	-	-
TextCNN [17]	81.50	88.10	-	-	-	-	-	-	-	-
TextRCNN [72]	-	-	-	-	-	-	96.49	-	-	-
DAN [69]	-	86.30	89.40	-	-	-	-	-	-	-
Tree-LSTM [1]		88.00	-	-	-	-	-	-	-	-
CharCNN [93]	-	-	-	95.12	62.05	-	-	90.49	98.45	-
HAN [108]	-	-	49.40	-	-	63.60	-	-	-	-
SeqTextRCNN [7]	-	-	-	-	-	-	-	-	-	-
oh-2LSTMp [75]	-	-	94.10	97.10	67.61	-	86.68	93.43	99.16	-
LSTMN [112]	-	87.30	-	-	-	-	-	-	-	-
Multi-Task [79]	-	87.90	91.30	-	-	-	-	-	-	-
BLSTM-2DCNN [77]	82.30	89.50	-	-	-	-	96.50	-	-	-
TopicRNN [83]	-	-	93.72	-	-	-	-	-	-	-
DPCNN [98]	-	-	-	97.36	69.42	65.19	-	93.13	99.12	-
KPCNN [100]	83.25	-	-	-	-	-	-	88.36	-	-
RNN-Capsule [87]	83.80	-	-	-	-	-	-	-	-	-
ULMFiT [285]	-	-	95.40	97.84	71.02	-	-	94.99	99.20	-
LEAM [286]	76.95	-	-	95.31	64.09	-	81.91	92.45	99.02	-
TextCapsule [101]	82.30	86.80	-	-	-	-	-	92.60	-	-
TextGCN [155]	76.74	-	-	-	-	-	86.34	67.61	-	-
BERT-base [19]	-	93.50	95.63	98.08	70.58	61.60	-	-	-	91.00
BERT-large [19]	-	94.90	95.79	98.19	71.38	62.20	-	-	-	91.70
MT-DNN [287]	-	95.60	83.20	-	-	-	-	-	-	91.50
XLNet-Large [138]	-	96.80	96.21	98.45	72.20	67.74	-	-	-	-
XLNet [138]	-	97.00	-	-	-	-	-	95.51	99.38	-
RoBERTa [140]	-	96.40	-	-	-	-	-	-	-	92.60

Bold is the most accurate.

such as BERT, RoBERTa, and XLNET, etc., except MR and 20NG, which have not been experimented on BERT-based models. Pre-trained models are essential to NLP. It uses a deep model to learn a better feature of the text. It also demonstrates that the accuracy of NLP tasks can be significantly improved by a profound model that can be pre-trained from unlabeled datasets. For the MR dataset, the accuracy of RNN-Capsule [87] is 83.8%, obtaining the best result. It suggests that RNN-Capsule builds a capsule in each category for sentiment analysis. It can output words including sentiment trends indicating attributes of capsules with no applying linguistic knowledge. For 20NG dataset, BLSTM-2DCNN [77] gets 96.5% score with the best accuracy score. It may demonstrate the effectiveness of applying the 2D max-pooling operation to obtain a fixed-length representation of the text and utilize 2D convolution to sample more meaningful matrix information.

5 FUTURE RESEARCH CHALLENGES

Text classification – as efficient information retrieval and mining technology – plays a vital role in managing text data. It uses NLP, data mining, machine learning, and other techniques to automatically classify and discover different text types. Text classification takes multiple types of text as input, and the text is represented as a vector by the pre-training model. Then the vector is fed into the DNN for training until the termination condition is reached, and finally, the performance of the training model is verified by the downstream task. Existing models have already shown their usefulness in text classification, but there are still many possible improvements to explore.

Although some new text classification models repeatedly brush up the accuracy index of most classification tasks, they cannot indicate whether the model “understands” the text from the semantic level like human beings can. Moreover, with the emergence of the noise sample, the small sample noise may cause the decision confidence to change substantially or even lead to decision reversal. Therefore, the semantic representation ability and robustness of the model need to be proven in practice. Besides, the pre-trained semantic representation model represented by word vectors can often improve the performance of downstream NLP tasks. The existing research on the transfer strategy of context-free word vectors is still relatively preliminary. Thus, we conclude from a data, models, and performance perspective, text classification mainly faces the following challenges.

5.1 Challenges from Data Perspective

For a text classification task, data is essential to model performance, whether it is a traditional or deep learning method. The text data mainly studied includes multi-chapter, short text, cross-language, multi-label, less sample text. For the characteristics of these data, the existing technical challenges are as follows:

Zero-shot/Few-shot learning. Zero-shot or few-shot learning for text classification aims to classify text having no or few same labeled class data. However, the current models are too dependent on numerous labeled data. The performance of these models is significantly affected by zero-shot or few-shot learning. Thus, some works focus on tackling these problems. The main idea is to infer the features through learning kinds of semantic knowledge, such as learning relationships among classes [288] and incorporating class descriptions [170]. Furthermore, latent features generation [289], meta-Learning [106, 290, 291], and dynamic memory mechanism [292] are also efficient methods. Nevertheless, with the limitation of little unseen class data and different data distribution between seen class and unseen class, there is still a long way to go to reach the learning ability comparable to that of humans.

The external knowledge. As we all know, the more beneficial information is input into a DNN, the better the performance. For example, a question answering system incorporating common-sense knowledge can answer questions about the real world and help solve problems with incomplete information. Therefore, adding external knowledge (knowledge base or knowledge graph) [293, 294] is an efficient way to promote the model’s performance. The existing knowledge includes conceptual information [100, 127, 204], commonsense knowledge [223], knowledge base information [295, 296], general knowledge graph [170], and so on, which enhances the semantic representation of texts. Nevertheless, with the limitation of input scale, how and what to add for different tasks is still a challenge.

Special domain with many terminologies. Most of the existing models are supervised models, which over-rely on numerous labeled data. When the sample size is too small, or zero samples occur, the performance of the model will be significantly affected. New data set annotation takes a lot of time. Therefore, unsupervised learning and semi-supervised learning have great potential for

text classification. Furthermore, texts in a particular field [297, 298], such as financial and medical texts, contain many specific words or domain experts intelligible slang, abbreviations, etc., which make the existing pre-trained word vectors challenging to work on.

The multi-label text classification task. Multi-label text classification requires full consideration of the semantic relationship among labels, and the embedding and encoding of the model is a process of lossy compression [299, 300]. Therefore, how to reduce the loss of hierarchical semantics and retain rich and complex document semantic information during training is still a problem to be solved.

5.2 Challenges from Model Perspective

Most existing structures of traditional and deep learning models are tried for text classification, including integration methods. BERT learns a language representation that can be used to fine-tune for many NLP tasks. The primary method is to increase data, improve computation power, and design training procedures for getting better results [301–303]. Understanding the tradeoff between data and compute resources and prediction performance is worth studying.

Text representation. The text representation method based on the vector space model is simple and effective in the text preprocessing stage. However, it will lose the semantic information of the text, so the application performance based on this method is limited. The proposed semantically based text representation method is too time-consuming. Therefore, the efficient semantically based text representation method still needs further research. In the text representation of text classification based on deep learning, word embedding is the main concept, while the representation unit is described differently in different languages. Then, a word is represented in the form of a vector by learning mapping rules through the model. Therefore, how to design adaptive data representation methods is more conducive to the combination of deep learning and specific classification tasks.

Model integration. Most structures of traditional and deep learning models are tried for text classification, including integration methods. RNN requires recursive step by step to get global information. CNN can obtain local information, and the sensing field can be increased through the multi-layer stack to capture more comprehensive contextual information. Attention mechanisms learn global dependency among words in a sentence. The transformer model is dependent on attention mechanisms to establish the depth of the global dependency relationship between the input and output. Therefore, designing an integrated model is worth trying to take advantage of these models.

Model efficiency. Text classification models based on deep learning are highly effective, such as CNNs, RNNs, and GNNs. However, there are many technical limitations, such as the depth of the network layer, regularization problem, network learning rate, etc. Therefore, there is still more broad space for development to optimize the algorithm and improve the speed of model training.

5.3 Challenges from Performance Perspective

The traditional model and the deep model can achieve good performance in most text classification tasks, but the anti-interference ability of their results needs to be improved [176, 304, 305]. How to realize the interpretation of the deep model is also a technical challenge.

The semantic robustness of the model. In recent years, researchers have designed many models to enhance the accuracy of text classification models. However, when there are some adversarial samples in the datasets, the model's performance decreases significantly. Adversarial training is a crucial method to improve the robustness of the pre-training model. For example, a popular

approach is converting attack into defense and using the adversarial sample training model. Consequently, how to improve the robustness of models is a current research hotspot and challenge.

The interpretability of the model. DNNs have unique advantages in feature extraction and semantic mining and have achieved excellent text classification tasks. Only a better understanding of the theories behind these models can accurately design better models for various applications. However, deep learning is a black-box model, the training process is challenging to reproduce, and the implicit semantics and output interpretability are poor. It makes the improvement and optimization of the model lose clear guidelines. Why does one model outperform another on one data set but underperform on others? What does the deep learning model learn? Furthermore, we cannot accurately explain why the model improves performance.

6 CONCLUSION

This paper principally introduces the existing models for text classification tasks from traditional models to deep learning. Firstly, we introduce some primary traditional models and deep learning models with a summary table. The traditional model improves text classification performance mainly by improving the feature extraction scheme and classifier design. In contrast, the deep learning model enhances performance by improving the presentation learning method, model structure, and additional data and knowledge. Then, we introduce the datasets with a summary table and evaluation metrics for single-label and multi-label tasks. Furthermore, we give the quantitative results of the leading models in a summary table under different applications for classic text classification datasets. Finally, we summarize the possible future research challenges of text classification.

ACKNOWLEDGMENTS

Thanks for computing infrastructure provided by Huawei MindSpore platform.

REFERENCES

- [1] Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved semantic representations from tree-structured long short-term memory networks. In *Proc. ACL, 2015*. 1556–1566. <http://dx.doi.org/10.3115/v1/p15-1150>
- [2] Xiao-Dan Zhu, Parinaz Sobhani, and Hongyu Guo. 2015. Long short-term memory over recursive structures. In *Proc. ICML, 2015*. 1604–1612. <http://proceedings.mlr.press/v37/zhub15.html>.
- [3] Junyang Chen, Zhiguo Gong, and Weiwen Liu. 2020. A dirichlet process biterm-based mixture model for short text stream clustering. *Appl. Intell.* 50, 5 (2020), 1609–1619. <http://dx.doi.org/10.1007/s10489-019-01606-1>
- [4] Junyang Chen, Zhiguo Gong, and Weiwen Liu. 2019. A nonparametric model for online topic discovery with word embeddings. *Inf. Sci.* 504 (2019), 32–47. <http://dx.doi.org/10.1016/j.ins.2019.07.048>
- [5] Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proc. ACL, 2014*. 655–665. <http://dx.doi.org/10.3115/v1/p14-1062>
- [6] Pengfei Liu, Xipeng Qiu, Xinchi Chen, Shiyu Wu, and Xuanjing Huang. 2015. Multi-timescale long short-term memory neural network for modelling sentences and documents. In *Proc. EMNLP, 2015*. 2326–2335. <http://dx.doi.org/10.18653/v1/d15-1280>
- [7] Ji Young Lee and Franck Dernoncourt. 2016. Sequential short-text classification with recurrent and convolutional neural networks. In *Proc. NAACL, 2016*. 515–520. <http://dx.doi.org/10.18653/v1/n16-1062>
- [8] M. E. Maron. 1961. Automatic indexing: An experimental inquiry. *J. ACM* 8, 3 (1961), 404–417. <http://dx.doi.org/10.1145/321075.321084>
- [9] Thomas M. Cover and Peter E. Hart. 1967. Nearest neighbor pattern classification. *IEEE Trans. Inf. Theory* 13, 1 (1967), 21–27. <http://dx.doi.org/10.1109/TIT.1967.1053964>
- [10] Thorsten Joachims. 1998. Text categorization with support vector machines: Learning with many relevant features. In *Proc. ECML, 1998*. 137–142. <http://dx.doi.org/10.1007/BFb0026683>
- [11] Rami Aly, Steffen Remus, and Chris Biemann. 2019. Hierarchical multi-label classification of text with capsule networks. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28 - August 2, 2019, Volume 2: Student Research Workshop*. 323–330. <http://dx.doi.org/10.18653/v1/p19-2045>

- [12] Kamran Kowsari, Kiana Jafari Meimandi, Mojtaba Heidarysafa, Sanjana Mendu, Laura E. Barnes, and Donald E. Brown. 2019. Text classification algorithms: A survey. *Information* 10, 4 (2019), 150. <http://dx.doi.org/10.3390/info10040150>
- [13] Shervin Minaee, Nal Kalchbrenner, Erik Cambria, Narjes Nikzad, Meysam Chenaghlu, and Jianfeng Gao. 2020. Deep learning based text classification: A comprehensive review. *CoRR* abs/2004.03705 (2020). arXiv:2004.03705 <https://arxiv.org/abs/2004.03705>
- [14] Leo Breiman. 2001. Random forests. *Mach. Learn.* 45, 1 (2001), 5–32. <http://dx.doi.org/10.1023/A:1010933404324>
- [15] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A scalable tree boosting system. In *Proc. ACM SIGKDD, 2016*. 785–794. <http://dx.doi.org/10.1145/2939672.2939785>
- [16] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. LightGBM: A highly efficient gradient boosting decision tree. In *Proc. NeurIPS, 2017*. 3146–3154. <http://papers.nips.cc/paper/6907-lightgbm-a-highly-efficient-gradient-boosting-decision-tree>
- [17] Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proc. EMNLP, 2014*. 1746–1751. <http://dx.doi.org/10.3115/v1/d14-1181>
- [18] Saad Albawi, Tareq Abed Mohammed, and Saad Al-Zawi. 2017. Understanding of a convolutional neural network. In *2017 International Conference on Engineering and Technology (ICET)*. IEEE, 1–6.
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proc. NAACL, 2019*. 4171–4186. <https://www.aclweb.org/anthology/N19-1423>.
- [20] Yin Zhang, Rong Jin, and Zhi-Hua Zhou. 2010. Understanding bag-of-words model: A statistical framework. *International Journal of Machine Learning and Cybernetics* 1, 1–4 (2010), 43–52.
- [21] William B. Cavnar, John M. Trenkle, et al. 1994. N-gram-based text categorization. In *Proceedings of SDAIR-94, 3rd Annual Symposium on Document Analysis and Information Retrieval*, Vol. 161175. Citeseer.
- [22] Ricardo Baeza-Yates, and Berthier Ribeiro-Neto. 1999. Modern information retrieval. ACM press, Vol. 463.
- [23] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *Proc. ICLR, 2013*. <http://arxiv.org/abs/1301.3781>.
- [24] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global vectors for word representation. In *Proc. EMNLP, 2014*. 1532–1543. <http://dx.doi.org/10.3115/v1/d14-1162>
- [25] Min-Ling Zhang and Kun Zhang. 2010. Multi-label learning by exploiting label dependency. In *Proc. ACM SIGKDD, 2010*. 999–1008. <http://dx.doi.org/10.1145/1835804.1835930>
- [26] Karl-Michael Schneider. 2004. A new feature selection score for multinomial Naïve Bayes text classification based on KL-divergence. In *Proc. ACL, 2004*. <https://www.aclweb.org/anthology/P04-3024/>.
- [27] Thomas M. Cover and Joy A. Thomas. 2006. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, USA.
- [28] Wenyuan Dai, Gui-Rong Xue, Qiang Yang, and Yong Yu. 2007. Transferring Naïve Bayes classifiers for text classification. In *Proc. AAAI, 2007*. 540–545. <http://www.aaai.org/Library/AAAI/2007/aaai07-085.php>.
- [29] A. P. Dempster, N. M. Laird, and D. B. Rubin. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society* (1977).
- [30] M. Granik and V. Mesyura. 2017. Fake news detection using Naïve Bayes classifier. In *2017 IEEE First Ukraine Conference on Electrical and Computer Engineering (UKRCON)*. 900–903. <http://dx.doi.org/10.1109/UKRCON.2017.8100379>
- [31] Mohamad Syahrul Mubarok, Kang Adiwijaya, and Muhammad Aldhi. 2017. Aspect-based sentiment analysis to review products using Naïve Bayes. *AIP Conference Proceedings* 1867 (08 2017), 020060. <http://dx.doi.org/10.1063/1.4994463>
- [32] Shuo Xu. 2018. Bayesian Naïve Bayes classifiers to text classification. *J. Inf. Sci.* 44, 1 (2018), 48–59. <http://dx.doi.org/10.1177/0165551516677946>
- [33] G. Singh, B. Kumar, L. Gaur, and A. Tyagi. 2019. Comparison between multinomial and Bernoulli Naïve Bayes for text classification. In *2019 International Conference on Automation, Computational and Technology Management (ICACTM)*. 593–596. <http://dx.doi.org/10.1109/ICACTM.2019.8776800>
- [34] 2007. 20NG Corpus. <http://ana.cachopo.org/datasets-for-single-label-text-categorization>. (2007).
- [35] Mark Craven, Dan DiPasquo, Dayne Freitag, Andrew McCallum, Tom M. Mitchell, Kamal Nigam, and Séan Slattery. 1998. Learning to extract symbolic knowledge from the World Wide Web. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence and Tenth Innovative Applications of Artificial Intelligence Conference, AAAI 98, IAAI 98, July 26-30, 1998, Madison, Wisconsin, USA*. 509–516. <http://www.aaai.org/Library/AAAI/1998/aaai98-072.php>.
- [36] T. Jo. 2017. Using K nearest neighbors for text segmentation with feature similarity. In *2017 International Conference on Communication, Control, Computing and Electronics Engineering (ICCCCEE)*. 1–5. <http://dx.doi.org/10.1109/ICCCCEE.2017.7866706>

- [37] Li Baoli, Lu Qin, and Yu Shiwen. 2004. An adaptive $<\mathbf{i}>\text{-}\mathbf{k}\text{-}<\mathbf{i}>$ -nearest neighbor text categorization strategy. *ACM Transactions on Asian Language Information Processing* 3, 4 (Dec. 2004), 215–226. <http://dx.doi.org/10.1145/1039621.1039623>
- [38] Shufeng Chen. 2018. K-nearest neighbor algorithm optimization in text categorization. *IOP Conference Series: Earth and Environmental Science* 108 (Jan 2018), 052074. <http://dx.doi.org/10.1088/1755-1315/108/5/052074>
- [39] Shengyi Jiang, Guansong Pang, Meiling Wu, and Limin Kuang. 2012. An improved K-nearest-neighbor algorithm for text categorization. *Expert Syst. Appl.* 39, 1 (2012), 1503–1509. <http://dx.doi.org/10.1016/j.eswa.2011.08.040>
- [40] Pascal Soucy and Guy W. Mineau. 2001. A simple KNN algorithm for text categorization. In *Proc. ICDM, 2001*. 647–648. <http://dx.doi.org/10.1109/ICDM.2001.989592>
- [41] Songbo Tan. 2005. Neighbor-weighted K-nearest neighbor for unbalanced text corpus. *Expert Syst. Appl.* 28, 4 (2005), 667–671. <http://dx.doi.org/10.1016/j.eswa.2004.12.023>
- [42] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Mach. Learn.* 20, 3 (1995), 273–297. <http://dx.doi.org/10.1007/BF00994018>
- [43] Christina Leslie, Eleazar Eskin, and William Stafford Noble. 2001. The spectrum kernel: A string kernel for SVM protein classification. In *Biocomputing 2002*. World Scientific, 564–575.
- [44] Hirotoshi Taira and Masahiko Haruno. 1999. Feature selection in SVM text categorization. In *AAAI/IAAI*. 480–486.
- [45] Xin Li and Yuhong Guo. 2013. Active learning with multi-label SVM classification. In *IJCAI*. Citeseer, 1479–1485.
- [46] Tao Peng, Wanli Zuo, and Fengling He. 2008. SVM based adaptive learning method for text classification from positive and unlabeled documents. *Knowledge and Information Systems* 16, 3 (2008), 281–301.
- [47] Thorsten Joachims. 2001. A statistical learning model of text classification for support vector machines. In *Proc. SIGIR, 2001*. 128–136. <http://dx.doi.org/10.1145/383952.383974>
- [48] T. Joachims. 1999. Transductive inference for text classification using support vector machines. In *International Conference on Machine Learning*.
- [49] Tom M. Mitchell. 1997. *Machine Learning*. McGraw-Hill. <http://www.worldcat.org/oclc/61321007>.
- [50] Rajeev Rastogi and Kyuseok Shim. 2000. PUBLIC: A decision tree classifier that integrates building and pruning. *Data Min. Knowl. Discov.* 4, 4 (2000), 315–344. <http://dx.doi.org/10.1023/A:1009887311454>
- [51] Ross J. Quinlan. 1986. Induction of decision trees. *Machine Learning* 1, 1 (1986), 81–106.
- [52] J. Ross Quinlan. 1993. *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [53] Micheline Kamber, Lara Winstone, Wan Gong, Shan Cheng, and Jiawei Han. 1997. Generalization and decision tree induction: Efficient classification in data mining. In *Proceedings Seventh International Workshop on Research Issues in Data Engineering: High Performance Database Management for Large-Scale Applications*. IEEE, 111–120.
- [54] David E. Johnson, Frank J. Oles, Tong Zhang, and Thilo Götz. 2002. A decision-tree-based symbolic rule induction system for text categorization. *IBM Syst. J.* 41, 3 (2002), 428–437. <http://dx.doi.org/10.1147/sj.413.0428>
- [55] Peerapon Vateekul and Miroslav Kubat. 2009. Fast induction of multiple decision trees in text categorization from large scale, imbalanced, and multi-label data. In *Proc. ICDM Workshops, 2009*. 320–325. <http://dx.doi.org/10.1109/ICDMW.2009.94>
- [56] Yoav Freund and Robert E. Schapire. 1995. A decision-theoretic generalization of on-line learning and an application to boosting. In *Proc. EuroCOLT, 1995*. 23–37. http://dx.doi.org/10.1007/3-540-59119-2_166
- [57] Robert E. Schapire and Yoram Singer. 1999. Improved boosting algorithms using confidence-rated predictions. *Mach. Learn.* 37, 3 (1999), 297–336. <http://dx.doi.org/10.1023/A:1007614523901>
- [58] Ameni Bouaziz, Christel Dartigues-Pallez, Célia da Costa Pereira, Frédéric Precioso, and Patrick Lloret. 2014. Short text classification using semantic random forest. In *Proc. DAWAK, 2014*. 288–299. http://dx.doi.org/10.1007/978-3-319-10160-6_26
- [59] Md. Zahidul Islam, Jixue Liu, Jiuyong Li, Lin Liu, and Wei Kang. 2019. A semantics aware random forest for text classification. In *Proc. CIKM, 2019*. 1061–1070. <http://dx.doi.org/10.1145/3357384.3357891>
- [60] Richard Socher, Jeffrey Pennington, Eric H. Huang, Andrew Y. Ng, and Christopher D. Manning. 2011. Semi-supervised recursive autoencoders for predicting sentiment distributions. In *Proc. EMNLP, 2011*. 151–161. <https://www.aclweb.org/anthology/D11-1014/>
- [61] 2011. A MATLAB Implementation of RAE. <https://github.com/vin00/Semi-Supervised-Recursive-Autoencoders-for-Predicting-Sentiment-Distributions>. (2011).
- [62] Richard Socher, Brody Huval, Christopher D. Manning, and Andrew Y. Ng. 2012. Semantic compositionality through recursive matrix-vector spaces. In *Proc. EMNLP, 2012*. 1201–1211. <https://www.aclweb.org/anthology/D12-1110/>
- [63] 2012. A Tensorflow Implementation of MV_RNN. https://github.com/github-pengge/MV_RNN. (2012).
- [64] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proc. EMNLP, 2013*. 1631–1642. <https://www.aclweb.org/anthology/D13-1170/>.

- [65] 2013. A MATLAB Implementation of RNTN. <https://github.com/pondruska/DeepSentiment>. (2013).
- [66] Ozan Irsoy and Claire Cardie. 2014. Deep recursive neural networks for compositionality in language. In *Proc. NIPS, 2014*. 2096–2104. <http://papers.nips.cc/paper/5551-deep-recursive-neural-networks-for-compositionality-in-language>.
- [67] Quoc V. Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proc. ICML, 2014*. 1188–1196. <http://proceedings.mlr.press/v32/le14.html>.
- [68] 2014. A PyTorch Implementation of Paragraph Vectors (doc2vec). <https://github.com/inejc/paragraph-vectors>. (2014).
- [69] Mohit Iyyer, Varun Manjunatha, Jordan L. Boyd-Graber, and Hal Daumé III. 2015. Deep unordered composition rivals syntactic methods for text classification. In *Proc. ACL, 2015*. 1681–1691. <http://dx.doi.org/10.3115/v1/p15-1162>
- [70] 2015. An Implementation of DAN. <https://github.com/miyyer/dan>. (2015).
- [71] 2015. A PyTorch Implementation of Tree-LSTM. <https://github.com/stanfordnlp/treelstm>. (2015).
- [72] Siwei Lai, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Recurrent convolutional neural networks for text classification (AAAI'15). AAAI Press, 2267–2273.
- [73] 2015. A Tensorflow Implementation of TextRCNN. <https://github.com/roomylee/rcnn-text-classification>. (2015).
- [74] 2015. An Implementation of MT-LSTM. <https://github.com/AlexAntn/MTLSTM>. (2015).
- [75] Rie Johnson and Tong Zhang. 2016. Supervised and semi-supervised text categorization using LSTM for region embeddings. In *Proc. ICML, 2016*. 526–534. <http://proceedings.mlr.press/v48/johnson16.html>.
- [76] 2015. An Implementation of oh-2LSTMp. http://riejohnson.com/cnn_2download.html. (2015).
- [77] Peng Zhou, Zhenyu Qi, Suncong Zheng, Jiaming Xu, Hongyun Bao, and Bo Xu. 2016. Text classification improved by integrating bidirectional LSTM with two-dimensional max pooling. In *Proc. COLING, 2016*. 3485–3495. <https://www.aclweb.org/anthology/C16-1329>.
- [78] 2016. An Implementation of BLSTM-2DCNN. <https://github.com/ManuelVs/NNForTextClassification>. (2016).
- [79] Pengfei Liu, Xipeng Qiu, and Xuanjing Huang. 2016. Recurrent neural network for text classification with multi-task learning. In *Proc. IJCAI, 2016*. 2873–2879. <http://www.ijcai.org/Abstract/16/408>.
- [80] 2016. A PyTorch Implementation of Multi-Task. https://github.com/baixl/text_classification. (2016).
- [81] Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm. In *Proc. EMNLP, 2017*. 1615–1625. <https://www.aclweb.org/anthology/D17-1169>.
- [82] 2018. A Keras Implementation of DeepMoji. <https://github.com/bfelbo/DeepMoji>. (2018).
- [83] Adji B. Dieng, Chong Wang, Jianfeng Gao, and John W. Paisley. 2017. TopicRNN: A recurrent neural network with long-range semantic dependency. In *Proc. ICLR, 2017*. <https://openreview.net/forum?id=rJbbOLce>.
- [84] 2017. A PyTorch Implementation of TopicRNN. <https://github.com/dangitstam/topic-rnn>. (2017).
- [85] Takeru Miyato, Andrew M. Dai, and Ian J. Goodfellow. 2017. Adversarial training methods for semi-supervised text classification. In *Proc. ICLR, 2017*. https://openreview.net/forum?id=r1X3g2_xl.
- [86] 2017. A Tensorflow Implementation of Virtual Adversarial Training. https://github.com/tensorflow/models/tree/master/adversarial_text. (2017).
- [87] Yeqian Wang, Aixin Sun, Jialong Han, Ying Liu, and Xiaoyan Zhu. 2018. Sentiment analysis by capsules. In *Proc. WWW, 2018*. 1165–1174. <http://dx.doi.org/10.1145/3178876.3186015>
- [88] 2018. A PyTorch Implementation of RNN-Capsule. <https://github.com/wangjiosw/Sentiment-Analysis-by-Capsules>. (2018).
- [89] Yi Zhao, Yanyan Shen, and Junjie Yao. 2019. Recurrent neural network for text classification with hierarchical multiscale dense connections. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10–16, 2019*. 5450–5456. <http://dx.doi.org/10.24963/ijcai.2019/757>
- [90] 2019. An Implementation of HM-DenseRNNs. <https://github.com/zhaoyizhaoyi/hm-densernns>. (2019).
- [91] 2014. A Keras Implementation of TextCNN. <https://github.com/alexander-rakhlin/CNN-for-Sentence-Classification-in-Keras>. (2014).
- [92] 2014. A Tensorflow Implementation of DCNN. https://github.com/kinimod23/ATS_Project. (2014).
- [93] Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Proc. NeurIPS, 2015*. 649–657. <http://papers.nips.cc/paper/5782-character-level-convolutional-networks-for-text-classification>.
- [94] 2015. A Tensorflow Implementation of CharCNN. <https://github.com/mhjabreel/CharCNN>. (2015).
- [95] 2016. A Keras Implementation of SeqTextRCNN. <https://github.com/limugur/short-text-classification>. (2016).
- [96] Jingzhou Liu, Wei-Cheng Chang, Yuexin Wu, and Yiming Yang. 2017. Deep learning for extreme multi-label text classification. In *Proc. ACM SIGIR, 2017*. 115–124. <http://dx.doi.org/10.1145/3077136.3080834>
- [97] 2017. A Pytorch Implementation of XML-CNN. <https://github.com/siddsax/XML-CNN>. (2017).
- [98] Rie Johnson and Tong Zhang. 2017. Deep pyramid convolutional neural networks for text categorization. In *Proc. ACL, 2017*. 562–570. <http://dx.doi.org/10.18653/v1/P17-1052>

- [99] 2017. A PyTorch Implementation of DPCNN. <https://github.com/Cheneng/DPCNN>. (2017).
- [100] Jin Wang, Zhongyuan Wang, Dawei Zhang, and Jun Yan. 2017. Combining knowledge with deep convolutional neural networks for short text classification. In *Proc. IJCAI, 2017*. 2915–2921. <http://dx.doi.org/10.24963/ijcai.2017/406>
- [101] Min Yang, Wei Zhao, Jianbo Ye, Zeyang Lei, Zhou Zhao, and Soufei Zhang. 2018. Investigating capsule networks with dynamic routing for text classification. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*. 3110–3119. <http://dx.doi.org/10.18653/v1/d18-1350>
- [102] 2018. A Tensorflow Implementation of TextCapsule. https://github.com/andyweizhao/capsule_text_classification. (2018).
- [103] Kazuya Shimura, Jiyi Li, and Fumiyo Fukumoto. 2018. HFT-CNN: Learning hierarchical category structure for multi-label short text categorization. In *Proc. EMNLP, 2018*. 811–816. <http://dx.doi.org/10.18653/v1/d18-1093>
- [104] 2018. An Implementation of HFT-CNN. <https://github.com/ShimShim46/HFT-CNN>. (2018).
- [105] Jingyun Xu and Yi Cai. 2019. Incorporating context-relevant knowledge into convolutional neural networks for short text classification. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*. 10067–10068. <http://dx.doi.org/10.1609/aaai.v33i01.330110067>
- [106] Yujia Bao, Menghua Wu, Shiyu Chang, and Regina Barzilay. 2020. Few-shot text classification with distributional signatures. In *Proc. ICLR, 2020*. <https://openreview.net/forum?id=H1emfT4twB>.
- [107] 2020. A PyTorch Implementation of Few-shot Text Classification with Distributional Signatures. <https://github.com/YujiaBao/Distributional-Signatures>. (2020).
- [108] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J. Smola, and Eduard H. Hovy. 2016. Hierarchical attention networks for document classification. In *Proc. NAACL, 2016*. 1480–1489. <https://www.aclweb.org/anthology/N16-1174/>.
- [109] 2014. A Keras Implementation of TextCNN. <https://github.com/richliao/textClassifier>. (2014).
- [110] Xinjie Zhou, Xiaojun Wan, and Jianguo Xiao. 2016. Attention-based LSTM network for cross-lingual sentiment classification. In *Proc. EMNLP, 2016*. 247–256. <http://dx.doi.org/10.18653/v1/d16-1024>
- [111] 2013. NLP&CC Corpus. <http://tccf.cc.org.cn/conference/2013/index.html>. (2013).
- [112] Jianpeng Cheng, Li Dong, and Mirella Lapata. 2016. Long short-term memory-networks for machine reading. In *Proc. EMNLP, 2016*. 551–561. <http://dx.doi.org/10.18653/v1/d16-1053>
- [113] 2016. A Tensorflow Implementation of LSTMNN. <https://github.com/JRC1995/Abstractive-Summarization>. (2016).
- [114] Zhouchen Lin, Minwei Feng, Cáceres Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. A structured self-attentive sentence embedding. In *Proc. ICLR, 2017*. https://openreview.net/forum?id=BJC_jUqxe.
- [115] 2017. A PyTorch Implementation of Structured-Self-Attention. <https://github.com/kaushalshetty/Structured-Self-Attention>. (2017).
- [116] Pengcheng Yang, Xu Sun, Wei Li, Shuming Ma, Wei Wu, and Houfeng Wang. 2018. SGM: Sequence generation model for multi-label classification. In *Proc. COLING, 2018*. 3915–3926. <https://www.aclweb.org/anthology/C18-1330/>.
- [117] 2018. A PyTorch Implementation of SGM. <https://github.com/lancopku/SGM>. (2018).
- [118] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proc. NAACL, 2018*. 2227–2237. <https://www.aclweb.org/anthology/N18-1202/>.
- [119] 2018. A PyTorch Implementation of ELMo. <https://github.com/flairNLP/flair>. (2018).
- [120] Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, and Chengqi Zhang. 2018. Bi-directional block self-attention for fast and memory-efficient sequence modeling. In *Proc. ICLR, 2018*. <https://openreview.net/forum?id=H1cWzoxA->.
- [121] 2018. A PyTorch Implementation of BiBloSA. <https://github.com/galsang/BiBloSA-pytorch>. (2018).
- [122] Ronghui You, Zihan Zhang, Ziye Wang, Suyang Dai, Hiroshi Mamitsuka, and Shanfeng Zhu. 2019. AttentionXML: Label tree-based attention-aware deep model for high-performance extreme multi-label text classification. In *Proc. NeurIPS, 2019*. 5812–5822. <http://papers.nips.cc/paper/8817-attentionxml-label-tree-based-attention-aware-deep-model-for-high-performance-extreme-multi-label-text-classification>.
- [123] 2019. A PyTorch Implementation of AttentionXML. <https://github.com/yourh/AttentionXML>. (2019).
- [124] Shengli Sun, Qingfeng Sun, Kevin Zhou, and Tengchao Lv. 2019. Hierarchical attention prototypical networks for few-shot text classification. In *Proc. EMNLP, 2019*. 476–485. <http://dx.doi.org/10.18653/v1/D19-1045>
- [125] Tianyu Gao, Xu Han, Zhiyuan Liu, and Maosong Sun. 2019. Hybrid attention-based prototypical networks for noisy few-shot relation classification. In *Proc. AAAI, 2019*. 6407–6414. <http://dx.doi.org/10.1609/aaai.v33i01.33016407>
- [126] 2019. A PyTorch Implementation of HATT-Proto. <https://github.com/thunlp/HATT-Proto>. (2019).
- [127] Jindong Chen, Yizhou Hu, Jingping Liu, Yanghua Xiao, and Haiyun Jiang. 2019. Deep short text classification with knowledge powered attention. In *Proc. AAAI, 2019*. 6252–6259. <http://dx.doi.org/10.1609/aaai.v33i01.33016252>

- [128] 2019. A PyTorch Implementation of STCKA. <https://github.com/AlRobotZhang/STCKA>. (2019).
- [129] Kaize Ding, Jianling Wang, Jundong Li, Dingcheng Li, and Huan Liu. 2020. Be more with less: Hypergraph attention networks for inductive text classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*. 4927–4936. <http://dx.doi.org/10.18653/v1/2020.emnlp-main.399>
- [130] 2020. A Pytorch Implementation of HyperGAT. <https://github.com/kaize0409/HyperGAT>. (2020).
- [131] Qipeng Guo, Xipeng Qiu, Pengfei Liu, Xiangyang Xue, and Zheng Zhang. 2020. Multi-scale self-attention for text classification. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. 7847–7854. <https://aaai.org/ojs/index.php/AAAI/article/view/6290>.
- [132] Seungtaek Choi, Haeju Park, Jinyoung Yeo, and Seung-won Hwang. 2020. Less is more: Attention supervision with counterfactuals for text classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*. 6695–6704. <http://dx.doi.org/10.18653/v1/2020.emnlp-main.543>
- [133] 2019. A Tensorflow Implementation of BERT. <https://github.com/google-research/bert>. (2019).
- [134] Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, and Ion Androutsopoulos. 2019. Large-scale multi-label text classification on EU legislation. In *Proc. ACL, 2019*. 6314–6322. <http://dx.doi.org/10.18653/v1/p19-1636>
- [135] 2019. A Tensorflow Implementation of BERT-BASE. <https://github.com/iliashalkidis/lmtc-eurlex57k>. (2019).
- [136] Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to fine-tune BERT for text classification? In *Proc. CCL, 2019*. 194–206. http://dx.doi.org/10.1007/978-3-030-32381-3_16
- [137] 2019. A Tensorflow Implementation of BERT4doc-Classification. <https://github.com/xuyige/BERT4doc-Classification>. (2019).
- [138] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized autoregressive pretraining for language understanding. In *Proc. NeurIPS, 2019*. 5754–5764. <http://papers.nips.cc/paper/8812-xlnet-generalized-autoregressive-pretraining-for-language-understanding>.
- [139] 2019. A Tensorflow Implementation of XLNet. <https://github.com/zihangdai/xlnet>. (2019).
- [140] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *CoRR* abs/1907.11692 (2019). arXiv:1907.11692 <http://arxiv.org/abs/1907.11692>
- [141] 2019. A PyTorch Implementation of RoBERTa. <https://github.com/pytorch/fairseq>. (2019).
- [142] Danilo Croce, Giuseppe Castellucci, and Roberto Basili. 2020. GAN-BERT: Generative adversarial learning for robust text classification with a bunch of labeled examples. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*. 2114–2119. <http://dx.doi.org/10.18653/v1/2020.acl-main.191>
- [143] 2020. A Pytorch Implementation of GAN-BERT. <https://github.com/crux82/ganbert>. (2020).
- [144] Siddhant Garg and Goutham Ramakrishnan. 2020. BAE: BERT-based adversarial examples for text classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*. 6174–6181. <http://dx.doi.org/10.18653/v1/2020.emnlp-main.498>
- [145] 2020. An Implementation of BAE. https://github.com/QData/TextAttack/blob/master/textattack/attack_recipes/bae_garg_2019.py. (2020).
- [146] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A lite BERT for self-supervised learning of language representations. In *Proc. ICLR, 2020*. <https://openreview.net/forum?id=H1eA7AEtvS>.
- [147] 2020. A Tensorflow Implementation of ALBERT. <https://github.com/google-research/ALBERT>. (2020).
- [148] Haopeng Zhang and Jiawei Zhang. 2020. Text graph transformer for document classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*. 8322–8327. <http://dx.doi.org/10.18653/v1/2020.emnlp-main.668>
- [149] Wei-Cheng Chang, Hsiang-Fu Yu, Kai Zhong, Yiming Yang, and Inderjit S. Dhillon. 2020. Taming pretrained transformers for extreme multi-label text classification. In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*. 3163–3171. <http://dx.doi.org/10.1145/3394486.3403368>
- [150] 2020. An Implementation of X-Transformer. <https://github.com/OctoberChang/X-Transformer>. (2020).
- [151] Ting Jiang, Deqing Wang, Leilei Sun, Huayi Yang, Zhengyang Zhao, and Fuzhen Zhuang. 2021. LightXML: Transformer with dynamic negative sampling for high-performance extreme multi-label text classification. *CoRR* abs/2101.03305 (2021). arXiv:2101.03305 <https://arxiv.org/abs/2101.03305>.
- [152] 2021. An Implementation of LightXML. <https://github.com/kongds/LightXML>. (2021).

[153] Hao Peng, Jianxin Li, Yu He, Yaopeng Liu, Mengjiao Bao, Lihong Wang, Yangqiu Song, and Qiang Yang. 2018. Large-scale hierarchical text classification with recursively regularized deep graph-CNN. In *Proc. WWW, 2018*. 1063–1072. <http://dx.doi.org/10.1145/3178876.3186005>

[154] 2018. A Tensorflow Implementation of DeepGraphCNNforTexts. <https://github.com/HKUST-KnowComp/DeepGraphCNNforTexts>. (2018).

[155] Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Graph convolutional networks for text classification. In *Proc. AAAI, 2019*. 7370–7377. <http://dx.doi.org/10.1609/aaai.v33i01.33017370>

[156] 2019. A Tensorflow Implementation of TextGCN. https://github.com/yao8839836/text_gcn. (2019).

[157] Felix Wu, Amauri H. Souza Jr., Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Q. Weinberger. 2019. Simplifying graph convolutional networks. In *Proc. ICML, 2019*. 6861–6871. <http://proceedings.mlr.press/v97/wu19e.html>.

[158] 2019. An Implementation of SGC. <https://github.com/Tiiger/SGC>. (2019).

[159] Lianzhe Huang, Dehong Ma, Sujian Li, Xiaodong Zhang, and Houfeng Wang. 2019. Text level graph neural network for text classification. In *Proc. EMNLP, 2019*. 3442–3448. <http://dx.doi.org/10.18653/v1/D19-1345>

[160] 2019. An Implementation of TextLevelGNN. <https://github.com/LindgeW/TextLevelGNN>. (2019).

[161] Hao Peng, Jianxin Li, Senzhang Wang, Lihong Wang, Qiran Gong, Renyu Yang, Bo Li, Philip Yu, and Lifang He. 2019. Hierarchical taxonomy-aware and attentional graph capsule RCNNs for large-scale multi-label text classification. *IEEE Transactions on Knowledge and Data Engineering* (2019).

[162] Yufeng Zhang, Xueli Yu, Zeyu Cui, Shu Wu, Zhongzhen Wen, and Liang Wang. 2020. Every document owns its structure: Inductive text classification via graph neural networks. In *Proc. ACL, 2020*. 334–339. <https://www.aclweb.org/anthology/2020.acl-main.31>.

[163] 2019. A Tensorflow Implementation of TextING. <https://github.com/CRIPAC-DIG/TextING>. (2019).

[164] Xien Liu, Xinxin You, Xiao Zhang, Ji Wu, and Ping Lv. 2020. Tensor graph convolutional networks for text classification. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. 8409–8416. <https://aaai.org/ojs/index.php/AAAI/article/view/6359>.

[165] 2019. A Tensorflow Implementation of TensorGCN. <https://github.com/THUMLP/TensorGCN>. (2019).

[166] Ankit Pal, Muru Selvakumar, and Malaikannan Sankarasubbu. 2020. MAGNET: Multi-label text classification using attention-based graph neural network. In *Proc. ICAART, 2020*. 494–505. <http://dx.doi.org/10.5220/0008940304940505>

[167] 2020. A Repository of MAGNET. <https://github.com/monk1337/MAGnet>. (2020).

[168] 2017. A Tensorflow Implementation of Miyato et al. <https://github.com/TobiasLee/Text-Classification>. (2017).

[169] Jichuan Zeng, Jing Li, Yan Song, Cuiyun Gao, Michael R. Lyu, and Irwin King. 2018. Topic memory networks for short text classification. In *Proc. EMNLP, 2018*. 3120–3131. <http://dx.doi.org/10.18653/v1/D18-1351>

[170] Jingqing Zhang, Piyawat Lertvittayakumjorn, and Yike Guo. 2019. Integrating semantic knowledge to tackle zero-shot text classification. In *Proc. NAACL, 2019*. 1031–1040. <http://dx.doi.org/10.18653/v1/n19-1108>

[171] 2019. A Tensorflow Implementation of KG4ZeroShotText. <https://github.com/JingqingZ/KG4ZeroShotText>. (2019).

[172] M. K. Alsmadi, K. B. Omar, S. A. Noah, and I. Almarashdah. 2009. Performance comparison of multi-layer perceptron (back propagation, delta rule and perceptron) algorithms in neural networks. In *2009 IEEE International Advance Computing Conference*. 296–299.

[173] Samira Pouyanfar, Saad Sadiq, Yilin Yan, Haiman Tian, Yudong Tao, Maria E. Presa Reyes, Mei-Ling Shyu, Shu-Ching Chen, and S. S. Iyengar. 2019. A survey on deep learning: Algorithms, techniques, and applications. *ACM Comput. Surv.* 51, 5 (2019), 92:1–92:36. <http://dx.doi.org/10.1145/3234150>

[174] Libo Qin, Wanxiang Che, Yangming Li, Minheng Ni, and Ting Liu. 2020. DCR-Net: A deep co-interactive relation network for joint dialog act recognition and sentiment classification. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. 8665–8672. <https://aaai.org/ojs/index.php/AAAI/article/view/6391>.

[175] Zhongfen Deng, Hao Peng, Dongxiao He, Jianxin Li, and Philip S. Yu. 2021. HTCIInfoMax: A global model for hierarchical text classification via information maximization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tür, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tammooy Chakraborty, and Yichao Zhou (Eds.). Association for Computational Linguistics, 3259–3265. <http://dx.doi.org/10.18653/v1/2021.nacl-main.260>

[176] Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is BERT really robust? A strong baseline for natural language attack on text classification and entailment. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. 8018–8025. <https://aaai.org/ojs/index.php/AAAI/article/view/6311>.

[177] Chen Li, Xutan Peng, Hao Peng, Jianxin Li, and Lihong Wang. 2021. TextGTL: Graph-based transductive learning for semi-supervised TextClassification via structure-sensitive interpolation. In *IJCAI 2021*. ijcai.org.

[178] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. 2017. Virtual adversarial training: A regularization method for supervised and semi-supervised learning. *CoRR* abs/1704.03976 (2017). arXiv:1704.03976 <http://arxiv.org/abs/1704.03976>.

[179] Geoffrey E. Hinton, Alex Krizhevsky, and Sida D. Wang. 2011. Transforming auto-encoders. In *Proc. ICANN, 2011*, Timo Honkela, Włodzisław Duch, Mark Girolami, and Samuel Kaski (Eds.). Springer Berlin, Berlin, 44–51.

[180] S. Hochreiter and J. Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9, 8 (1997), 1735–1780.

[181] Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proc. EMNLP, 2015*. 632–642. <http://dx.doi.org/10.18653/v1/d15-1075>

[182] Zhiguo Wang, Wael Hamza, and Radu Florian. 2017. Bilateral multi-perspective matching for natural language sentences. In *Proc. IJCAI, 2017*. 4144–4150. <http://dx.doi.org/10.24963/ijcai.2017/579>

[183] Rie Johnson and Tong Zhang. 2015. Semi-supervised convolutional neural networks for text categorization via region embedding. In *Proc. NeurIPS, 2015*. 919–927. <http://papers.nips.cc/paper/5849-semi-supervised-convolutional-neural-networks-for-text-categorization-via-region-embedding>.

[184] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Identity mappings in deep residual networks. In *Proc. ECCV, 2016*. 630–645. http://dx.doi.org/10.1007/978-3-319-46493-0_38

[185] Issam Bazzi. 2002. *Modelling out-of-vocabulary words for robust speech recognition*. Ph.D. Dissertation. Massachusetts Institute of Technology.

[186] Huy Nguyen and Minh-Le Nguyen. 2017. A deep neural architecture for sentence-level sentiment classification in Twitter social networking. In *Proc. PACLING, 2017*. 15–27. http://dx.doi.org/10.1007/978-981-10-8438-6_2

[187] Benjamin Adams and Grant McKenzie. 2018. Crowdsourcing the character of a place: Character-level convolutional networks for multilingual geographic text classification. *Trans. GIS* 22, 2 (2018), 394–408. <http://dx.doi.org/10.1111/tgis.12317>

[188] Zhuang Chen and Tieyun Qian. 2019. Transfer capsule network for aspect level sentiment classification. In *Proc. ACL, 2019*. 547–556. <http://dx.doi.org/10.18653/v1/p19-1052>

[189] Wei Xue, Wubai Zhou, Tao Li, and Qing Wang. 2017. MTNA: A neural multi-task model for aspect category classification and aspect term extraction on restaurant reviews. In *Proc. IJCNLP, 2017*. 151–156. <https://www.aclweb.org/anthology/I17-2026>.

[190] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proc. ICLR, 2015*. <http://arxiv.org/abs/1409.0473>.

[191] Zikun Hu, Xiang Li, Cunchao Tu, Zhiyuan Liu, and Maosong Sun. 2018. Few-shot charge prediction with discriminative legal attributes. In *Proc. COLING, 2018*. 487–498. <https://www.aclweb.org/anthology/C18-1041>.

[192] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proc. NeurIPS, 2017*. 5998–6008. <http://papers.nips.cc/paper/7181-attention-is-all-you-need>.

[193] Yukun Ma, Haiyun Peng, and Erik Cambria. 2018. Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive LSTM. In *Proc. AAAI, 2018*. 5876–5883. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16541>.

[194] Yequan Wang, Minlie Huang, Xiaoyan Zhu, and Li Zhao. 2016. Attention-based LSTM for aspect-level sentiment classification. In *Proc. EMNLP, 2016*. 606–615. <http://dx.doi.org/10.18653/v1/d16-1058>

[195] Feifan Fan, Yansong Feng, and Dongyan Zhao. 2018. Multi-grained attention network for aspect-level sentiment classification. In *Proc. EMNLP, 2018*. 3433–3442. <http://dx.doi.org/10.18653/v1/d18-1380>

[196] Ming Tan, Cicero dos Santos, Bing Xiang, and Bowen Zhou. 2016. Improved representation learning for question answer matching. In *Proc. ACL, 2016*. Association for Computational Linguistics, Berlin, Germany, 464–473. <http://dx.doi.org/10.18653/v1/P16-1044>

[197] Cicero Nogueira dos Santos, Ming Tan, Bing Xiang, and Bowen Zhou. 2016. Attentive pooling networks. *CoRR* abs/1602.03609 (2016). arXiv:1602.03609 <http://arxiv.org/abs/1602.03609>.

[198] Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. Pre-trained models for natural language processing: A survey. *CoRR* abs/2003.08271 (2020). arXiv:2003.08271 <https://arxiv.org/abs/2003.08271>.

[199] Alec Radford. 2018. Improving language understanding by generative pre-training.

[200] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Viet Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive language models beyond a fixed-length context. In *Proc. ACL, 2019*. 2978–2988. <http://dx.doi.org/10.18653/v1/p19-1285>

[201] Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. What does BERT learn about the structure of language? In *Proc. ACL, 2019*. 3651–3657. <http://dx.doi.org/10.18653/v1/p19-1356>

[202] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language

generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*. 7871–7880. <http://dx.doi.org/10.18653/v1/2020.acl-main.703>

[203] Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. 2020. SpanBERT: Improving pre-training by representing and predicting spans. *Trans. Assoc. Comput. Linguistics* 8 (2020), 64–77. <https://transacl.org/ojs/index.php/tacl/article/view/1853>.

[204] Yu Sun, Shuohuan Wang, Yu-Kun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019. ERNIE: Enhanced representation through knowledge integration. *CoRR* abs/1904.09223 (2019). arXiv:1904.09223 <http://arxiv.org/abs/1904.09223>.

[205] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *Proc. NeurIPS, 2016*. 3837–3845. <http://papers.nips.cc/paper/6081-convolutional-neural-networks-on-graphs-with-fast-localized-spectral-filtering>.

[206] Hao Peng, Ruitong Zhang, Yingtong Dou, Renyu Yang, Jingyi Zhang, and Philip S. Yu. 2021. Reinforced neighborhood selection guided multi-relational graph neural networks. *arXiv preprint arXiv:2104.07886* (2021).

[207] Hao Peng, Renyu Yang, Zheng Wang, Jianxin Li, Lifang He, Philip Yu, Albert Zomaya, and Raj Ranjan. 2021. Lime: Low-cost incremental learning for dynamic heterogeneous information networks. *IEEE Trans. Comput.* (2021).

[208] Jianxin Li, Hao Peng, Yuwei Cao, Yingtong Dou, Hekai Zhang, Philip Yu, and Lifang He. 2021. Higher-order attribute-enhancing heterogeneous graph neural networks. *IEEE Transactions on Knowledge and Data Engineering* (2021).

[209] Diego Marcheggiani and Ivan Titov. 2017. Encoding sentences with graph convolutional networks for semantic role labeling. In *Proc. EMNLP, 2017*. 1506–1515. <http://dx.doi.org/10.18653/v1/d17-1159>

[210] Yifu Li, Ran Jin, and Yuan Luo. 2019. Classifying relations in clinical narratives using segment graph convolutional and recurrent neural networks (Seg-GCRNs). *JAMIA* 26, 3 (2019), 262–268. <http://dx.doi.org/10.1093/jamia/ocy157>

[211] Joost Bastings, Ivan Titov, Wilker Aziz, Diego Marcheggiani, and Khalil Sima'an. 2017. Graph convolutional encoders for syntax-aware neural machine translation. In *Proc. EMNLP, 2017*. 1957–1967. <http://dx.doi.org/10.18653/v1/d17-1209>

[212] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *Proc. ICLR, 2018*. <https://openreview.net/forum?id=rJXMpikCZ>.

[213] Linmei Hu, Tianchi Yang, Chuan Shi, Houye Ji, and Xiaoli Li. 2019. Heterogeneous graph attention networks for semi-supervised short text classification. In *Proc. EMNLP, 2019*. 4820–4829. <http://dx.doi.org/10.18653/v1/D19-1488>

[214] Zhongyang Li, Xiao Ding, and Ting Liu. 2018. Constructing narrative event evolutionary graph for script event prediction. In *Proc. IJCAI, 2018*. 4201–4207. <http://dx.doi.org/10.24963/ijcai.2018/584>

[215] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. 1993. Signature verification using a siamese time delay neural network. In *Proc. NeurIPS, 1993*. 737–744. <http://papers.nips.cc/paper/769-signature-verification-using-a-siamese-time-delay-neural-network>.

[216] Jonas Mueller and Aditya Thyagarajan. 2016. Siamese recurrent architectures for learning sentence similarity. In *Proc. AAAI, 2016*. 2786–2792. <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12195>.

[217] Jayadeva, Himanshu Pant, Mayank Sharma, and Sumit Soman. 2019. Twin neural networks for the classification of large unbalanced datasets. *Neurocomputing* 343 (2019), 34–49. <http://dx.doi.org/10.1016/j.neucom.2018.07.089> Learning in the Presence of Class Imbalance and Concept Drift.

[218] Takeru Miyato, Shinichi Maeda, Masanori Koyama, Ken Nakae, and Shin Ishii. 2015. Distributional Smoothing with Virtual Adversarial Training. (2015). arXiv:stat.ML/1507.00677.

[219] Tianyang Zhang, Minlie Huang, and Li Zhao. 2018. Learning structured representation for text classification via reinforcement learning. In *Proc. AAAI, 2018*. 6053–6060. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16537>.

[220] Jason Weston, Sumit Chopra, and Antoine Bordes. 2015. Memory Networks. (2015). arXiv:cs.AI/1410.3916.

[221] Xin Li and Wai Lam. 2017. Deep multi-task learning for aspect term extraction with memory interaction. In *Proc. EMNLP, 2017*. 2886–2892. <http://dx.doi.org/10.18653/v1/d17-1310>

[222] Chenlin Shen, Changlong Sun, Jingjing Wang, Yangyang Kang, Shoushan Li, Xiaozhong Liu, Luo Si, Min Zhang, and Guodong Zhou. 2018. Sentiment classification towards question-answering with hierarchical matching network. In *Proc. EMNLP, 2018*. 3654–3663. <http://dx.doi.org/10.18653/v1/d18-1401>

[223] Xiao Ding, Kuo Liao, Ting Liu, Zhongyang Li, and Junwen Duan. 2019. Event representation learning enhanced with external commonsense knowledge. In *Proc. EMNLP, 2019*. 4893–4902. <http://dx.doi.org/10.18653/v1/D19-1495>

[224] Yazhou Zhang, Dawei Song, Peng Zhang, Xiang Li, and Panpan Wang. 2019. A quantum-inspired sentiment representation model for Twitter sentiment analysis. *Applied Intelligence* (2019).

[225] Ofir Zafrir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. 2019. Q8BERT: Quantized 8Bit BERT. In *Fifth Workshop on Energy Efficient Machine Learning and Cognitive Computing - NeurIPS Edition, EMC2@NeurIPS 2019, Vancouver, Canada, December 13, 2019*. 36–39. <http://dx.doi.org/10.1109/EMC2-NIPS53020.2019.00016>

[226] 2002. MR Corpus. <http://www.cs.cornell.edu/people/pabo/movie-review-data/>. (2002).

- [227] 2013. SST Corpus. <http://nlp.stanford.edu/sentiment>. (2013).
- [228] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proc. EMNLP, 2013*. Association for Computational Linguistics, Seattle, Washington, USA, 1631–1642. <https://www.aclweb.org/anthology/D13-1170>.
- [229] 2005. MPQA Corpus. <http://www.cs.pitt.edu/mpqa/>. (2005).
- [230] Qiming Diao, Minghui Qiu, Chao-Yuan Wu, Alexander J. Smola, Jing Jiang, and Chong Wang. 2014. Jointly modeling aspects, ratings and sentiments for movie recommendation (JMARS). In *Proc. ACM SIGKDD, 2014*. 193–202. <http://dx.doi.org/10.1145/2623330.2623758>
- [231] Duyu Tang, Bing Qin, and Ting Liu. 2015. Document modeling with gated recurrent neural network for sentiment classification. In *Proc. EMNLP, 2015*. 1422–1432. <http://dx.doi.org/10.18653/v1/d15-1167>
- [232] 2015. Amazon Review Corpus. <https://www.kaggle.com/datasets/akondguru/consumer-reviews-of-amazon-products>. (2015).
- [233] 2013. Twitter Corpus. <https://www.cs.york.ac.uk/semeval-2013/task2/>. (2013).
- [234] 2004. AG Corpus. http://www.di.unipi.it/~gulli/AG_corpus_of_news_articles.html. (2004).
- [235] 2007. Reuters Corpus. [https://www.cs.umb.edu/~smimarog/textmining/datasets/](http://www.cs.umb.edu/~smimarog/textmining/datasets/). (2007).
- [236] Canhui Wang, Min Zhang, Shaoping Ma, and Liyun Ru. 2008. Automatic online news issue construction in web environment. In *Proc. WWW, 2008*. 457–466. <http://dx.doi.org/10.1145/1367497.1367560>
- [237] Ximing Li, Changchun Li, Jinjin Chi, Jihong Ouyang, and Chenliang Li. 2018. Dataless text classification: A topic modeling approach with document manifold. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM 2018, Torino, Italy, October 22–26, 2018*. 973–982. <http://dx.doi.org/10.1145/3269206.3271671>
- [238] Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N. Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick van Kleef, Sören Auer, and Christian Bizer. 2015. DBpedia - A large-scale, multilingual knowledge base extracted from Wikipedia. *Semantic Web* 6, 2 (2015), 167–195. <http://dx.doi.org/10.3233/SW-140134>
- [239] 2015. Ohsumed Corpus. <http://davis.wpi.edu/xmdv/datasets/ohsumed.html>. (2015).
- [240] 2019. EUR-Lex Corpus. <http://www.kenst.de/resources/eurlex/eurlex.html>. (2019).
- [241] 2016. Amazon670K Corpus. <http://manikvarma.org/downloads/XC/XMLRepository.html>. (2016).
- [242] Jianhua Yin and Jianyong Wang. 2014. A dirichlet multinomial mixture model-based approach for short text clustering. In *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD'14, New York, NY, USA - August 24 - 27, 2014*. 233–242. <http://dx.doi.org/10.1145/2623330.2623715>
- [243] J. Chen, Z. Gong, W. Wang, W. Liu, M. Yang, and C. Wang. 2020. TAM: Targeted analysis model with reinforcement learning on short texts. *IEEE Transactions on Neural Networks and Learning Systems* (2020), 1–10. <http://dx.doi.org/10.1109/TNNLS.2020.3009247>
- [244] Fang Wang, Zhongyuan Wang, Zhoujun Li, and Ji-Rong Wen. 2014. Concept-based short text classification and ranking. In *Proc. CIKM, 2014*. 1069–1078. <http://dx.doi.org/10.1145/2661829.2662067>
- [245] 2015. Fudan Corpus. www.datatang.com/data/44139and43543. (2015).
- [246] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100, 000+ questions for machine comprehension of text. In *Proc. EMNLP, 2016*. 2383–2392. <http://dx.doi.org/10.18653/v1/d16-1264>
- [247] Xuchen Yao, Benjamin Van Durme, Chris Callison-Burch, and Peter Clark. 2013. Answer extraction as sequence tagging with tree edit distance. In *Proc. NAACL, 2013*. 858–867. <https://www.aclweb.org/anthology/N13-1106>.
- [248] 2002. TREC Corpus. <https://cogcomp.seas.upenn.edu/Data/QA/QC/>. (2002).
- [249] Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. WikiQA: A challenge dataset for open-domain question answering. In *Proc. EMNLP, 2015*. 2013–2018. <http://dx.doi.org/10.18653/v1/d15-1237>
- [250] Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proc. ACL, 2004*. Barcelona, Spain, 271–278. <http://dx.doi.org/10.3115/1218955.1218990>
- [251] Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proc. ACM SIGKDD, 2004*. 168–177. <http://dx.doi.org/10.1145/1014052.1014073>
- [252] 2017. Reuters Corpus. <https://martin-thoma.com/nlp-reuters>. (2017).
- [253] Jaeyoung Kim, Sion Jang, Eunjeong L. Park, and Sungchul Choi. 2020. Text classification using capsules. *Neurocomputing* 376 (2020), 214–221. <http://dx.doi.org/10.1016/j.neucom.2019.10.033>
- [254] 2020. Reuters10 Corpus. <http://www.nltk.org/book/ch02.html>. (2020).
- [255] David D. Lewis, Yiming Yang, Tony G. Rose, and Fan Li. 2004. RCV1: A new benchmark collection for text categorization research. *J. Mach. Learn. Res.* 5 (2004), 361–397. <http://jmlr.org/papers/volume5/lewis04a/lewis04a.pdf>.
- [256] 2004. RCV1-V2 Corpus. http://www.ai.mit.edu/projects/jmlr/papers/volume5/lewis04a/lyrl2004_rcv1v2_README.htm. (2004).
- [257] Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. In *Proc. ACL, 2005*. 115–124. <https://www.aclweb.org/anthology/P05-1015>.

[258] Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation* 39, 2–3 (2005), 165–210. <http://dx.doi.org/10.1007/s10579-005-7880-9>

[259] Mike Thelwall, Kevan Buckley, and Georgios Paltoglou. 2012. Sentiment strength detection for the social web. *J. Assoc. Inf. Sci. Technol.* 63, 1 (2012), 163–173. <http://dx.doi.org/10.1002/asi.21662>

[260] Preslav Nakov, Alan Ritter, Sara Rosenthal, Fabrizio Sebastiani, and Veselin Stoyanov. 2016. SemEval-2016 task 4: Sentiment analysis in Twitter. In *Proc. SemEval, 2016*.

[261] Zhiyong Lu. 2011. PubMed and beyond: A survey of web tools for searching biomedical literature. *Database* 2011 (2011). <http://dx.doi.org/10.1093/database/baq036>

[262] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A human generated machine reading comprehension dataset. In *Proc. NeurIPS, 2016*. http://ceur-ws.org/Vol-1773/CoCoNIPS_2016_paper9.pdf.

[263] <https://data.quora.com/First-Quora-Dataset-Release-QuestionPairs>. ([n. d.]).

[264] Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don't know: Unanswerable questions for SQuAD. In *Proc. ACL, 2018*, 784–789. <http://dx.doi.org/10.18653/v1/P18-2124>

[265] Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proc. NAACL, 2018*, 1112–1122. <http://dx.doi.org/10.18653/v1/n18-1101>

[266] Marco Marelli, Luisa Bentivogli, Marco Baroni, Raffaella Bernardi, Stefano Menini, and Roberto Zamparelli. 2014. SemEval-2014 task 1: Evaluation of compositional distributional semantic models on full sentences through semantic relatedness and textual entailment. In *Proc. SemEval, 2014*, 1–8. <http://dx.doi.org/10.3115/v1/s14-2001>

[267] Bill Dolan, Chris Quirk, and Chris Brockett. 2004. Unsupervised construction of large paraphrase corpora: Exploiting massively parallel news sources. In *Proc. COLING, 2004*. <https://www.aclweb.org/anthology/C04-1051>.

[268] Daniel M. Cer, Mona T. Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. SemEval-2017 task 1: Semantic textual similarity - multilingual and cross-lingual focused evaluation. *CoRR* abs/1708.00055 (2017). arXiv:1708.00055 <http://arxiv.org/abs/1708.00055>.

[269] Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The PASCAL recognising textual entailment challenge. In *Machine Learning Challenges, Evaluating Predictive Uncertainty, Visual Object Classification and Recognizing Textual Entailment, First PASCAL Machine Learning Challenges Workshop, MLCW 2005, Southampton, UK, April 11-13, 2005, Revised Selected Papers*, 177–190. http://dx.doi.org/10.1007/11736790_9

[270] Tushar Khot, Ashish Sabharwal, and Peter Clark. 2018. SciTaiL: A textual entailment dataset from science question answering. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, 5189–5197. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17368>.

[271] Kamran Kowsari, Donald E. Brown, Mojtaba Heidarysafa, Kiana Jafari Meimandi, Matthew S. Gerber, and Laura E. Barnes. 2017. HDLTex: Hierarchical deep learning for text classification. In *Proc. ICMLA, 2017*, 364–371. <http://dx.doi.org/10.1109/ICMLA.2017.0-134>

[272] 2018. AmazonCat-13K Corpus. https://drive.google.com/open?id=1VwHAbri6y6oh8lkpZ6sSY_b1FRNnCLFL. (2018).

[273] 2017. BlurbGenreCollection-EN Corpus. <https://www.inf.uni-hamburg.de/en/inst/ab/lt/resources/data/blurb-genre-collection.html>. (2017).

[274] Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Penacchietti, Lorenza Romano, and Stan Szpakowicz. 2009. SemEval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. In *Proc. NAACL, 2009*, 94–99. <https://www.aclweb.org/anthology/W09-2415>.

[275] Stephanie M. Strassel, Mark A. Przybocki, Kay Peterson, Zhiyi Song, and Kazuaki Maeda. 2008. Linguistic resources and evaluation techniques for evaluation of cross-document automatic content extraction. In *Proc. LREC, 2008*. <http://www.lrec-conf.org/proceedings/lrec2008/summaries/677.html>.

[276] Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D. Manning. 2017. Position-aware attention and supervised data improve slot filling. In *Proc. EMNLP, 2017*, 35–45. <http://dx.doi.org/10.18653/v1/d17-1004>

[277] Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. Modeling relations and their mentions without labeled text. In *Proc. ECML PKDD, 2010*, 148–163. http://dx.doi.org/10.1007/978-3-642-15939-8_10

[278] 2019. FewRel Corpus. <https://github.com/thunlp/FewRel>. (2019).

[279] Seokhwan Kim, Luis Fernando D’Haro, Rafael E. Banchs, Jason D. Williams, and Matthew Henderson. 2016. The fourth dialog state tracking challenge. In *Proc. IWSDS, 2016*, 435–449. http://dx.doi.org/10.1007/978-981-10-2585-3_36

[280] Jeremy Ang, Yang Liu, and Elizabeth Shriberg. 2005. Automatic dialog act segmentation and classification in multi-party meetings. In *Proc. ICASSP, 2005*, 1061–1064. <http://dx.doi.org/10.1109/ICASSP.2005.1415300>

[281] Dan Jurafsky and Elizabeth Shriberg. 1997. Switchboard SWBD-DAMSL shallow-discourse-function annotation coders manual. (01 1997).

[282] Aliaksei Severyn and Alessandro Moschitti. 2015. Learning to rank short text pairs with convolutional deep neural networks. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, Santiago, Chile, August 9-13, 2015*, Ricardo Baeza-Yates, Mounia Lalmas, Alistair Moffat, and Berthier A. Ribeiro-Neto (Eds.). ACM, 373–382. <http://dx.doi.org/10.1145/2766462.2767738>

[283] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press. <http://dx.doi.org/10.1017/CBO9780511809071>

[284] Tetsuji Nakagawa, Kentaro Inui, and Sadao Kurohashi. 2010. Dependency tree-based sentiment classification using CRFs with hidden variables. In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 2-4, 2010, Los Angeles, California, USA*. 786–794. <https://aclanthology.org/N10-1120>.

[285] Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *Proc. ACL, 2018*. 328–339. <http://dx.doi.org/10.18653/v1/P18-1031>

[286] Guoyin Wang, Chunyuan Li, Wenlin Wang, Yizhe Zhang, Dinghan Shen, Xinyuan Zhang, Ricardo Henao, and Lawrence Carin. 2018. Joint embedding of words and labels for text classification. In *Proc. ACL, 2018*. 2321–2331. <http://dx.doi.org/10.18653/v1/P18-1216>

[287] Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019. Multi-task deep neural networks for natural language understanding. In *Proc. ACL, 2019*. 4487–4496. <http://dx.doi.org/10.18653/v1/p19-1441>

[288] Pushpankar Kumar Pushp and Muktabh Mayank Srivastava. 2017. Train once, test anywhere: Zero-shot learning for text classification. *CoRR* abs/1712.05972 (2017). arXiv:1712.05972 <http://arxiv.org/abs/1712.05972>.

[289] Congzheng Song, Shanghang Zhang, Najmeh Sadoughi, Pengtao Xie, and Eric P. Xing. 2020. Generalized zero-shot text classification for ICD coding. In *Proc. IJCAI, 2020*. 4018–4024. <http://dx.doi.org/10.24963/ijcai.2020/556>

[290] Ruiying Geng, Binhua Li, Yongbin Li, Xiaodan Zhu, Ping Jian, and Jian Sun. 2019. Induction networks for few-shot text classification. In *Proc. EMNLP, 2019*. 3902–3911. <http://dx.doi.org/10.18653/v1/D19-1403>

[291] Shumin Deng, Ningyu Zhang, Zhanlin Sun, Jiaoyan Chen, and Huajun Chen. 2020. When low resource NLP meets unsupervised language model: Meta-pretraining then meta-learning for few-shot text classification (student abstract). In *Proc. AAAI, 2020*. 13773–13774. <https://aaai.org/ojs/index.php/AAAI/article/view/7158>.

[292] Ruiying Geng, Binhua Li, Yongbin Li, Jian Sun, and Xiaodan Zhu. 2020. Dynamic memory induction networks for few-shot text classification. In *Proc. ACL, 2020*. 1087–1094. <https://www.aclweb.org/anthology/2020.acl-main.102/>.

[293] Kervy Rivas Rojas, Gina Bustamante, Arturo Oncevay, and Marco Antonio Sobrevilla Cabezudo. 2020. Efficient strategies for hierarchical text classification: External knowledge and auxiliary tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5–10, 2020*. 2252–2257. <http://dx.doi.org/10.18653/v1/2020.acl-main.205>

[294] Niloofer Shanavas, Hui Wang, Zhiwei Lin, and Glenn I. Hawe. 2021. Knowledge-driven graph similarity for text classification. *Int. J. Mach. Learn. Cybern.* 12, 4 (2021), 1067–1081. <http://dx.doi.org/10.1007/s13042-020-01221-4>

[295] Yanchao Hao, Yuanzhe Zhang, Kang Liu, Shizhu He, Zhanyi Liu, Hua Wu, and Jun Zhao. 2017. An end-to-end model for question answering over knowledge base with cross-attention combining global knowledge. In *Proc. ACL, 2017*. 221–231. <http://dx.doi.org/10.18653/v1/P17-1021>

[296] Rima Türker, Lei Zhang, Maria Kourtraki, and Harald Sack. 2018. TECNE: Knowledge based text classification using network embeddings. In *Proc. EKAW, 2018*. 53–56. <http://ceur-ws.org/Vol-2262/ekaw-demo-18.pdf>.

[297] Xin Liang, Dawei Cheng, Fangzhou Yang, Yifeng Luo, Weining Qian, and Aoying Zhou. 2020. F-HMTC: Detecting financial events for investment decisions based on neural hierarchical multi-label text classification. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*. 4490–4496. <http://dx.doi.org/10.24963/ijcai.2020/619>

[298] P. B. Shanthi, Shivani Modi, K. S. Hareesha, and Sampath Kumar. 2021. Classification and comparison of malignancy detection of cervical cells based on nucleus and textural features in microscopic images of uterine cervix. *Int. J. Medical Eng. Informatics* 13, 1 (2021), 1–13. <http://dx.doi.org/10.1504/IJMEI.2021.111861>

[299] Tianshi Wang, Li Liu, Naiwen Liu, Huaxiang Zhang, Long Zhang, and Shanshan Feng. 2020. A multi-label text classification method via dynamic semantic representation model and deep neural network. *Appl. Intell.* 50, 8 (2020), 2339–2351. <http://dx.doi.org/10.1007/s10489-020-01680-w>

[300] Boyan Wang, Xuegang Hu, Pei-Pei Li, and Philip S. Yu. 2021. Cognitive structure learning model for hierarchical multi-label text classification. *Knowl. Based Syst.* 218 (2021), 106876. <http://dx.doi.org/10.1016/j.knosys.2021.106876>

[301] Jinhua Du, Yan Huang, and Karo Moilanen. 2020. Pointing to select: A fast pointer-LSTM for long text classification. In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8–13, 2020*. 6184–6193. <http://dx.doi.org/10.18653/v1/2020.coling-main.544>

[302] Jie Du, Chi-Man Vong, and C. L. Philip Chen. 2021. Novel efficient RNN and LSTM-like architectures: Recurrent and gated broad learning systems and their applications for text classification. *IEEE Trans. Cybern.* 51, 3 (2021), 1586–1597. <http://dx.doi.org/10.1109/TCYB.2020.2969705>

- [303] Teja Kanchinadam, Qian You, Keith Westpfahl, James Kim, Siva Gunda, Sebastian Seith, and Glenn Fung. 2021. A simple yet brisk and efficient active learning platform for text classification. *CoRR* abs/2102.00426 (2021). arXiv:2102.00426 <https://arxiv.org/abs/2102.00426>.
- [304] Yichao Zhou, Jyun-Yu Jiang, Kai-Wei Chang, and Wei Wang. 2019. Learning to discriminate perturbations for blocking adversarial attacks in 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, Hong Kong, China, November 3-7, 2019*. 4903–4912. <http://dx.doi.org/10.18653/v1/D19-1496>
- [305] Ahmadreza Azizi, Ibrahim Asadullah Tahmid, Asim Waheed, Neal Mangaokar, Jiameng Pu, Mobin Javed, Chandan K. Reddy, and Bimal Viswanath. 2021. T-Miner: A generative approach to defend against trojan attacks on DNN-based text classification. *CoRR* abs/2103.04264 (2021). arXiv:2103.04264 <https://arxiv.org/abs/2103.04264>.

Received April 2021; revised August 2021; accepted October 2021