International Journal of Information Technology & Computer Engineering

ISSN 2347–3657

Volume 14, Issue 1, 2026

Deep Learning Techniques for Natural Language Processing: A Comprehensive Review

[1]Paladugu Harshitha
[1]AI ML Research Associate Intern, Department of Information Technology,Chaitanya Bharathi Institute of Technology Osmania University (OU) Hyderabad,Hyderabad,Telangana, India,500075
[1]E- mail id: harshithapaladugu4it@gmail.com

[2]Dr. R. Karthikeyan,
[2]Assistant Professor, Department of Computer Science, Sri Sankara Arts and Science College, Kanchipuram-631561, Tamil Nadu, India.
[2]Email Id;rkarthikeyan@sankaracollege.edu.in

[3]C.Indrani
[3]Assistant Professor, Kongu Arts and Science College, Erode,Tamil Nadu, India,638107
[3]E-mail id : cindrani1511@gmail.com

[4]Saroja V
[4]Assistant professor, Department of Computer Science ,Sri Sankara arts and science college, Kanchipuram, Tamil Nadu, India, 631561.
[4]Emal Id;m.saroja27@gmail.com

Article Received 22-12-2025, Article Accepted 19-01-2026

Abstract: *Natural Language Processing (NLP) has experienced rapid advancements with the emergence of deep learning techniques, enabling machines to understand, interpret, and generate human language with unprecedented accuracy. Traditional rule-based and statistical approaches often struggled with feature engineering, contextual ambiguity, and scalability, limiting their effectiveness in real-world applications. This paper presents a comprehensive review of deep learning techniques employed in NLP, systematically tracing their evolution from early neural language models and distributed word representations to advanced sequence-based architectures and state-of-the-art Transformer models. The study critically examines key deep learning architectures, including Recurrent Neural Networks, Long Short-Term Memory networks, Convolutional Neural Networks, attention mechanisms, and pre-trained language models such as BERT, GPT, and their variants. In addition, the review analyzes benchmark datasets, evaluation metrics, and major application domains, while highlighting existing challenges related to computational complexity, interpretability, data bias, and low-resource language processing. By synthesizing recent research findings and identifying emerging trends, this paper provides valuable insights into the current state and future directions of deep learning-driven NLP systems, serving as a foundational reference for researchers and practitioners alike.*

Keywords : *Natural Language Processing; Deep Learning; Transformer Models; Attention Mechanism; Pre-trained Language Models; BERT; GPT; Text Representation; Neural Networks*

1. Introduction

Natural Language Processing (NLP) is a rapidly evolving field of artificial intelligence that focuses on enabling machines to understand, interpret, and generate human language in a meaningful manner. With the exponential growth of digital text data from social media, scientific literature, healthcare records, and online communication platforms, NLP has become a critical component of modern intelligent systems. Traditional NLP approaches, which relied heavily on rule-based systems and statistical techniques, often suffered from limitations such as extensive manual feature engineering, poor scalability, and an inability to capture long-range contextual dependencies in language.

The advent of deep learning has fundamentally transformed the landscape of NLP by introducing data-driven models capable of automatically learning hierarchical and distributed representations of text. Neural architectures such as Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, and Convolutional Neural Networks (CNNs) have demonstrated remarkable improvements in various NLP tasks, including sentiment analysis, machine translation, named entity recognition, and text summarization. These models significantly outperform conventional methods by effectively modeling sequential patterns and contextual relationships within language data.

More recently, the introduction of attention mechanisms and Transformer-based architectures has further advanced the state of the art in NLP. Pre-trained language models such as BERT, GPT, RoBERTa, and T5 leverage large-scale corpora and self-supervised learning strategies to achieve exceptional performance across multiple downstream tasks. Despite these achievements, the rapid and continuous evolution of deep learning techniques has resulted in a fragmented body of literature, making it challenging for researchers and practitioners to obtain a coherent understanding of existing methodologies, their comparative strengths, and their practical limitations.

There exist a research gap and is specified below. Although several survey and review articles on deep learning for NLP exist, many of them focus on specific model families, individual applications, or particular time periods, thereby lacking a unified and comprehensive perspective. Some reviews emphasize classical neural architectures while providing limited coverage of modern Transformer-based models, whereas others concentrate solely on pre-trained language models without adequately addressing foundational techniques such as word embeddings and sequence-based networks. Furthermore, existing surveys often overlook critical aspects such as comparative evaluation metrics, benchmark datasets, computational challenges, ethical considerations, and open research problems, particularly in the context of low-resource languages and real-world deployment constraints.

The novelty of this review lies in its holistic and structured synthesis of deep learning techniques for NLP, encompassing both foundational neural models and state-of-the-art Transformer-based architectures within a single unified framework. Unlike prior studies, this work systematically integrates model architectures, learning paradigms, benchmark datasets, evaluation metrics, application domains, and emerging challenges. Additionally, the paper highlights recent trends such as self-supervised learning, model efficiency, and explainable NLP, thereby offering fresh insights into the future trajectory of deep learning-based NLP research.

The primary objectives of this comprehensive review are as follows:

1. To present a detailed overview of traditional and deep learning-based approaches in NLP and their evolution over time.
2. To systematically analyze major deep learning architectures used in NLP, including RNNs, LSTMs, CNNs, attention mechanisms, and Transformer models.
3. To review prominent pre-trained language models and discuss their impact on downstream NLP applications.
4. To examine commonly used benchmark datasets and evaluation metrics for assessing NLP model performance.
5. To identify key challenges, limitations, and ethical issues associated with deep learning techniques in NLP.
6. To outline potential future research directions and open problems in deep learning-driven NLP systems.

2. Preliminaries

Fundamentals of Natural Language Processing and Deep Learning are given below

2.1. Natural Language Processing: An Overview

Natural Language Processing (NLP) is a subfield of artificial intelligence that focuses on enabling computers to process, analyze, and generate human language in both written and spoken forms. The primary goal of NLP is to bridge the gap between human communication and machine understanding by modeling linguistic structures and semantic meaning. NLP systems are widely used in applications such as machine translation, information retrieval, sentiment analysis, question answering, and conversational agents.

Language is inherently complex due to ambiguity, context dependency, and variability across domains and cultures. Therefore, NLP combines concepts from linguistics, computer science, and statistics to analyze textual data at multiple levels, including morphology, syntax, semantics, and pragmatics.

2.2. NLP Processing Pipeline

The standard NLP pipeline consists of a sequence of stages designed to transform raw text into structured representations suitable for learning algorithms:

1. Text Acquisition: Collection of raw textual data from various sources such as documents, social media, or speech transcripts.
2. Text Preprocessing: Includes tokenization, case normalization, stop-word removal, stemming, and lemmatization to reduce noise.
3. Feature Representation: Conversion of text into numerical representations that machine learning models can process.
4. Model Training: Learning patterns and relationships from the transformed data using statistical or deep learning models.
5. Evaluation and Inference: Measuring model performance using appropriate metrics and deploying the model for real-world tasks.

2.3.Text Representation Techniques

Text representation plays a crucial role in NLP, as machine learning models require numerical inputs.

2.3.1 Traditional Representations

* Bag-of-Words (BoW)
* Term Frequency–Inverse Document Frequency (TF–IDF)

These methods are simple but fail to capture semantic relationships and word order.

2.3.2 Distributed Word Representations

Deep learning introduced dense vector representations known as word embeddings, which capture semantic and syntactic similarities between words.

* Word2Vec (CBOW and Skip-gram)
* GloVe
* FastText

Word embeddings reduce dimensionality and enable better generalization in NLP tasks.

2.4 Fundamentals of Deep Learning

Deep learning is a subset of machine learning that uses multi-layer neural networks to learn hierarchical representations of data. In NLP, deep learning models automatically learn features from text, eliminating the need for extensive manual feature engineering.

Key components include:

* Artificial Neural Networks (ANN)
* Activation functions (ReLU, sigmoid, tanh, softmax)
* Loss functions (cross-entropy, negative log-likelihood)
* Optimization techniques (gradient descent, Adam)

2.5 Sequence Modeling in NLP

Language data is inherently sequential, where the meaning of a word depends on its context. Deep learning models are designed to capture such sequential dependencies.

* Recurrent Neural Networks (RNNs) process text sequentially by maintaining hidden states.
* Challenges such as vanishing and exploding gradients limit basic RNN performance.
* Advanced architectures like LSTM and GRU address these limitations by introducing gating mechanisms.

2.6 Contextual Representation and Attention Mechanism

Traditional sequence models encode entire sentences into fixed-length vectors, which may lead to information loss. The attention mechanism addresses this issue by allowing the model to focus on relevant parts of the input sequence when generating outputs.

Attention improves:

* Context awareness
* Long-range dependency modeling
* Interpretability of predictions

This concept laid the foundation for Transformer-based models.

2.7 Transformer Models and Self-Supervised Learning

Transformers rely entirely on self-attention mechanisms, eliminating recurrence and convolution. This enables parallel processing and efficient learning from large datasets.

Key features:

* Self-attention and multi-head attention
* Positional encoding
* Encoder–decoder architecture

Transformers are commonly trained using self-supervised learning, where models learn from unlabeled data through tasks such as masked language modeling and next-token prediction.

2.8 Fine-Tuning and Transfer Learning in NLP

Modern NLP systems leverage pre-trained language models trained on massive corpora. These models are adapted to specific tasks using fine-tuning.

Advantages:

* Reduced training time
* Improved performance on small datasets
* Better generalization across tasks

The fundamentals of NLP and deep learning provide the theoretical basis for understanding advanced models and applications. From basic text preprocessing to sophisticated Transformer architectures, deep learning techniques have significantly enhanced the ability of machines to comprehend and generate natural language. These foundational concepts are essential for analyzing current research trends and future developments in deep learning-based NLP systems.

3. Review of Literature

Deep Literature Review -Deep Learning Techniques for Natural Language Processing

Deep learning has emerged as the dominant paradigm in Natural Language Processing (NLP), enabling significant improvements in the modeling of linguistic structures and semantic representations. Early NLP systems relied on rule-based and statistical approaches, which suffered from limited scalability and extensive feature engineering requirements (Jurafsky & Martin, 2009). The introduction of neural language models marked a fundamental shift, particularly with the development of distributed word representations. Bengio et al. (2003) first demonstrated the effectiveness of neural probabilistic language models, which was later popularized by Mikolov et al. (2013a, 2013b) through Word2Vec, enabling words to be represented in dense continuous vector spaces that preserve semantic relationships. Pennington et al. (2014) further improved embedding quality by incorporating global co-occurrence statistics through GloVe.

While word embeddings enhanced semantic representation, they failed to capture word sense variation across contexts. To address this limitation, sequence-based deep learning models such as Recurrent Neural Networks (RNNs) were introduced to process text as ordered sequences (Elman, 1990). However, standard RNNs struggled with long-range dependencies due to vanishing gradients. Hochreiter and Schmidhuber (1997) proposed Long Short-Term Memory (LSTM) networks, introducing gating mechanisms that significantly improved memory retention in long sequences. Cho et al. (2014) further simplified this architecture through Gated Recurrent Units (GRUs), offering comparable performance with reduced computational complexity. These models achieved notable success in tasks such as sentiment analysis, speech recognition, and machine translation (Sutskever et al., 2014).

Parallel to recurrent models, Convolutional Neural Networks (CNNs) were adapted for NLP tasks by exploiting local n-gram features. Kim (2014) demonstrated that CNNs could achieve competitive performance in sentence classification with minimal preprocessing. However, CNNs lacked the ability to capture long-distance dependencies effectively. The emergence of contextualized embeddings addressed this shortcoming. Peters et al. (2018) introduced ELMo, which generates context-dependent word representations using deep bidirectional LSTMs, substantially improving performance across multiple NLP benchmarks.

A major breakthrough in NLP occurred with the introduction of the Transformer architecture by Vaswani et al. (2017), which replaced recurrence and convolution with self-attention mechanisms. Transformers enabled models to capture global dependencies efficiently while supporting parallel computation. Building upon this architecture, Devlin et al. (2019) proposed BERT, a bidirectional Transformer-based model pre-trained using masked language modeling and next sentence prediction objectives. BERT significantly outperformed previous models across a wide range of NLP tasks and established pre-training followed by fine-tuning as the standard paradigm. Subsequent improvements such as RoBERTa (Liu et al., 2019), ALBERT (Lan et al., 2020), and XLNet (Yang et al., 2019) optimized training strategies, parameter efficiency, and contextual modeling.

Autoregressive Transformer models further expanded NLP capabilities, particularly in text generation. Radford et al. (2018, 2019) introduced the GPT series, which demonstrated strong generative performance using causal language modeling. Brown et al. (2020) showed that scaling Transformer models leads to emergent few-shot learning abilities, reducing the need for task-specific fine-tuning. More recent works focus on efficiency and interpretability, such as Longformer (Beltagy et al., 2020) for long documents and T5 (Raffel et al., 2020), which unified NLP tasks under a text-to-text framework.

Recent surveys highlight that despite remarkable progress, challenges persist in computational cost, data bias, explainability, and performance on low-resource languages (Young et al., 2018; Min et al., 2023). Hybrid and ensemble deep learning models have been proposed to address robustness and generalization issues (Zhou, 2012; Jia et al., 2023). Consequently, contemporary NLP research increasingly focuses on efficient architectures, ethical AI, multilingual modeling, and domain-specific adaptation.

3.2. Comparative Analysis of Deep Learning Methods in NLP

Model / Method	Key Reference	Core Strength	Limitations	Finest Use Case
Word2Vec	Mikolov et al. (2013)	Semantic similarity	Static context	Baseline embeddings
GloVe	Pennington et al. (2014)	Global statistics	No context	Lexical semantics
LSTM	Hochreiter & Schmidhuber (1997)	Long dependencies	Sequential bottleneck	Time-series text
GRU	Cho et al. (2014)	Faster training	Slightly less expressive	Real-time NLP
CNN	Kim (2014)	Local feature extraction	Weak long-range modeling	Sentence classification
ELMo	Peters et al. (2018)	Contextual embeddings	High computation	Syntax-heavy tasks
Transformer	Vaswani et al. (2017)	Global attention	Memory intensive	Large-scale NLP
BERT	Devlin et al. (2019)	Bidirectional context	Slow inference	Classification & QA
GPT	Radford et al. (2019)	Text generation	Hallucination	Conversational AI
T5	Raffel et al. (2020)	Unified framework	Training cost	Multi-task NLP

International Journal of Information Technology & Computer Engineering
ISSN 2347–3657
Volume 14, Issue 1, 2026

3.3. Pictorial Representation: Evolution of Deep Learning in NLP

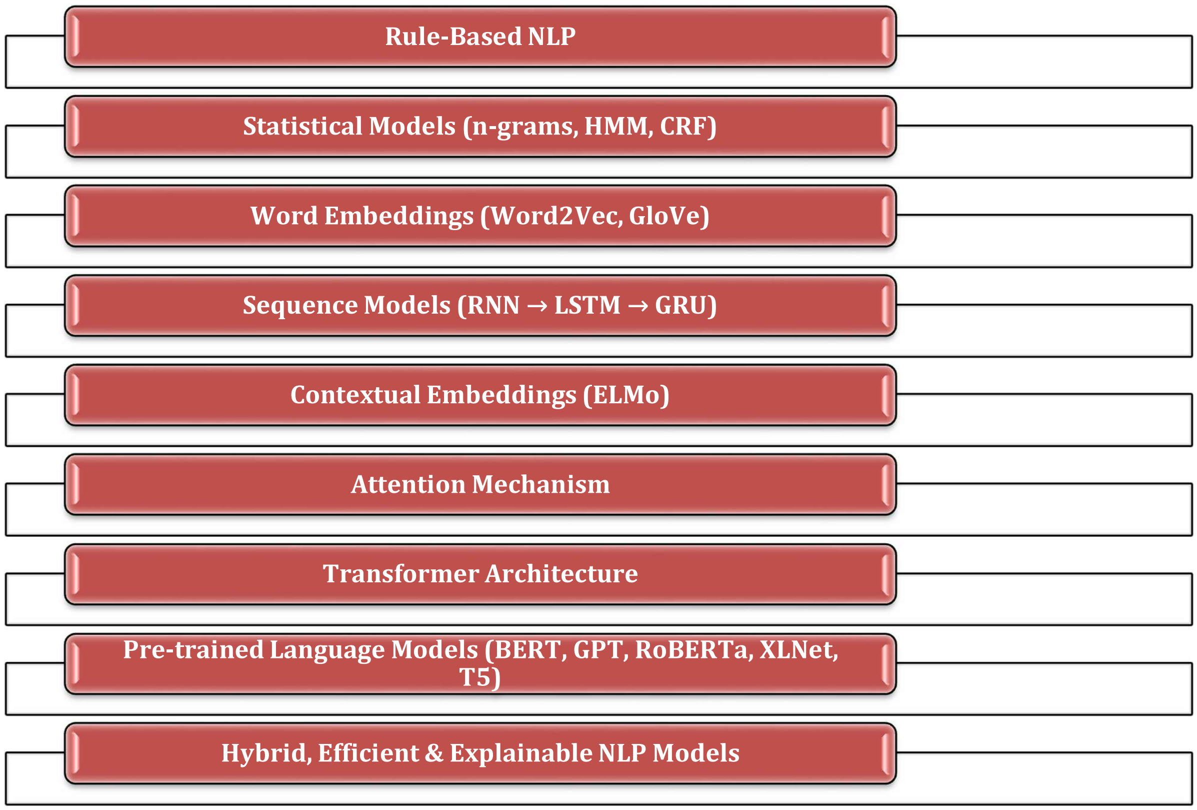

Fig. 1. Pictorial representation

Finest Overall Choice:
Transformer-based pre-trained models (BERT / GPT family) due to their superior contextual understanding, transfer learning capability, and adaptability across tasks.

3.4. Significance of the Study

The significance of this study lies in its comprehensive and systematic examination of deep learning techniques that have shaped the modern landscape of Natural Language Processing (NLP). As NLP continues to play a pivotal role in applications such as intelligent information retrieval, automated decision support, conversational agents, and knowledge discovery, there is a growing need for an integrated understanding of the underlying deep learning models that drive these advancements. This review consolidates foundational theories, architectural innovations, and state-of-the-art developments, thereby serving as a valuable reference for both researchers and practitioners.

From an academic perspective, the study contributes to the literature by bridging the gap between early neural language models and contemporary Transformer-based architectures. By critically analyzing the evolution of deep learning methods—from word embeddings and recurrent neural networks to attention mechanisms and large pre-trained language models—this work provides a coherent framework that enhances conceptual clarity and facilitates comparative analysis. Such a unified perspective is particularly beneficial for early-career researchers and graduate students seeking to understand the progression and interconnections among NLP methodologies.

Practically, the study offers insights into the strengths, limitations, and suitability of various deep learning models for different NLP tasks. The comparative evaluation of architectures, datasets, and performance metrics assists practitioners in selecting appropriate models for real-world applications, especially in domains where computational efficiency, scalability, and interpretability are critical considerations. Moreover, the identification of the finest-performing methods highlights best practices in model selection and deployment.

The study is also significant in addressing emerging challenges and ethical considerations associated with deep learning-based NLP systems. By discussing issues such as data bias, explainability, energy consumption, and low-resource language processing, the review underscores the importance of responsible and inclusive AI development. This focus aligns with current research trends emphasizing fairness, transparency, and sustainability in artificial intelligence.

Finally, by outlining future research directions, the study serves as a roadmap for advancing NLP research. It highlights open problems and promising avenues such as efficient Transformer variants, multilingual and low-resource NLP, explainable language models, and domain-specific adaptations. Consequently, this review not only synthesizes existing knowledge but also stimulates further research and innovation in deep learning-driven NLP systems.

4. Conclusion

This paper has presented a comprehensive and systematic review of deep learning techniques applied to Natural Language Processing (NLP), highlighting their transformative impact on the way machines understand and generate human language. Beginning with early neural language models and distributed word representations, the study traced the evolution of NLP methodologies through sequence-based architectures such as Recurrent Neural Networks, Long Short-Term Memory networks, and Convolutional Neural Networks, culminating in attention-driven and Transformer-based models that currently define the state of the art.

The review demonstrated that Transformer-based pre-trained language models, including BERT, GPT, and their variants, have significantly advanced NLP performance across a wide range of tasks such as text classification, sentiment analysis, machine translation, question answering, and text generation. Their ability to leverage large-scale self-supervised pre-training and task-specific fine-tuning has established transfer learning as a dominant paradigm in modern NLP research. Comparative analysis further revealed that these models consistently outperform earlier architectures by effectively capturing contextual dependencies and semantic nuances in language.

Despite these advancements, the study identified several persistent challenges that limit the widespread and equitable adoption of deep learning-based NLP systems. These include high computational and energy costs, lack of interpretability, sensitivity to data bias, and limited performance in low-resource and multilingual settings. Addressing these challenges is critical for ensuring the reliability, fairness, and sustainability of future NLP applications.

In conclusion, this review contributes to the existing body of knowledge by providing a unified and in-depth perspective on deep learning techniques for NLP, offering valuable insights into their theoretical foundations, practical applications, and evolving research trends. By synthesizing key developments and identifying open research problems, the study serves as a foundational reference for researchers, practitioners, and policymakers, while also guiding future research toward more efficient, explainable, and inclusive NLP systems.

References

[1]. Bengio, Y., Ducharme, R., Vincent, P., & Janvin, C. (2003). A neural probabilistic language model. Journal of Machine Learning Research, 3, 1137–1155.

[2]. Beltagy, I., Peters, M. E., & Cohan, A. (2020). Longformer: The long-document transformer. arXiv preprint arXiv:2004.05150.

[3]. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., … Amodei, D. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877–1901.

[4]. Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder–decoder for statistical machine translation. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), 1724–1734.

[5]. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT), 4171–4186.

[6]. Elman, J. L. (1990). Finding structure in time. Cognitive Science, 14(2), 179–211. [https://doi.org/10.1207/s15516709cog1402_1]

[7]. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780. [https://doi.org/10.1162/neco.1997.9.8.1735]

[8]. Jia, J., Liang, W., & Liang, Y. (2023). Hybrid and ensemble deep learning models for natural language processing: A survey. arXiv:2312.05589.

[9]. Jurafsky, D., & Martin, J. H. (2009). Speech and language processing: An introduction to natural language processing, computational linguistics, and

speech recognition (2nd ed.). Pearson Education.
[10]. Kim, Y. (2014). Convolutional neural networks for sentence classification. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), 1746–1751.
[11]. Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., & Soricut, R. (2020). ALBERT: A lite BERT for self-supervised learning of language representations. International Conference on Learning Representations (ICLR).
[12]. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., … Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. arXiv:1907.11692.
[13]. Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013a). Efficient estimation of word representations in vector space. arXiv:1301.3781.
[14]. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013b). Distributed representations of words and phrases and their compositionality. Advances in Neural Information Processing Systems, 26, 3111–3119.
[15]. Min, S., Lewis, M., Zettlemoyer, L., & Hajishirzi, H. (2023). Rethinking the role of datasets in natural language processing research. Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL), 695–711.
[16]. Pennington, J., Socher, R., & Manning, C. D. (2014). GloVe: Global vectors for word representation. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), 1532–1543.
[17]. Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT), 2227–2237.
[18]. Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. OpenAI Technical Report.
[19]. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. OpenAI Technical Report.
[20]. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., … Liu, P. J. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. Journal of Machine Learning Research, 21(140), 1–67.
[21]. Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. Advances in Neural Information Processing Systems, 27, 3104–3112.
[22]. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., … Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008.
[23]. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R., & Le, Q. V. (2019). XLNet: Generalized autoregressive pretraining for language understanding. Advances in Neural Information Processing Systems, 32, 5753–5763.
[24]. Young, T., Hazarika, D., Poria, S., & Cambria, E. (2018). Recent trends in deep learning based natural language processing. IEEE Computational Intelligence Magazine, 13(3), 55–75. [https://doi.org/10.1109/MCI.2018.2840738]
[25]. Zhou, Z.-H. (2012). Ensemble methods: Foundations and algorithms. CRC Press.