

## Review Article

# An In-Depth Review of Natural Language Processing (NLP): Progress, Challenges, and Future Directions

Parul Kastwar<sup>1</sup>, Chaitali Vishwakarma<sup>2</sup>

<sup>1,2</sup>Student, Department of Architecture, Madhav Institute of Technology, Gwalior, India

## I N F O

**Corresponding Author:**

Parul Kastwar, Department of Architecture,  
Madhav Institute of Technology, Gwalior, India

**E-mail Id:**

kastwarparul10@gmail.com

**Orcid Id:**

<https://orcid.org/0009-0006-1324-0224>

**How to cite this article:**

Kastwar P, Vishwakarma C. An In-Depth Review of Natural Language Processing (NLP): Progress, Challenges, and Future Directions. *Int J Adv Res Artif Intell Mach Learn Rev* 2025; 1(1): 9-15.

Date of Submission: 2025-01-08

Date of Acceptance: 2025-03-08

## A B S T R A C T

Natural Language Processing (NLP) has emerged as a pivotal domain in artificial intelligence, enabling machines to understand, interpret, and generate human language. With the advent of deep learning techniques, especially transformer models like GPT and BERT, NLP has seen remarkable advancements in tasks such as language translation, sentiment analysis, and conversational agents. This review explores the key progressions in NLP, examining the impact of deep learning, transfer learning, and pre-trained models on the field. It also addresses the ongoing challenges, including language ambiguity, biases in models, and the limitations posed by resource constraints for low-resource languages. Furthermore, it highlights future directions in NLP, including multimodal integration, personalized systems, and the ethical considerations necessary to ensure fair and transparent use of these technologies. Despite the impressive achievements, the field continues to evolve, with substantial room for growth in achieving truly intuitive, accurate, and universally accessible NLP systems.

**Keywords:** Natural Language Processing (NLP), Multimodal AI, Bias Mitigation, Zero-shot Learning, Ethical AI

## Introduction

Natural Language Processing (NLP) stands at the intersection of linguistics, computer science, and artificial intelligence, aiming to equip machines with the ability to understand, interpret, and generate human language. As one of the most dynamic and rapidly advancing fields in AI, NLP has seen transformative improvements over the past decade, driven largely by the development of deep learning techniques and the availability of large-scale language datasets. These advancements have enabled machines to perform complex language-related tasks such as speech recognition, machine translation, text summarization, sentiment analysis, and conversational AI, with increasingly human-like accuracy.<sup>1</sup>

The ability of machines to understand and interact in

natural language has far-reaching implications across various industries, from healthcare and finance to customer service and entertainment. However, while NLP systems have achieved impressive milestones, they are not without challenges. Human language is inherently ambiguous, context-dependent, and culturally diverse, making it a difficult task for machines to fully comprehend. Issues like bias in language models, difficulties in processing low-resource languages, and the need for transparency in AI decision-making remain significant hurdles.

This review provides a comprehensive overview of the current state of NLP, highlighting its key advancements, ongoing challenges, and potential future directions. By examining both the technical and ethical dimensions of

NLP, we aim to provide a nuanced understanding of where the field stands today and where it is headed in the coming years.<sup>2</sup>

### What is Natural Language Processing?

Natural Language Processing (NLP) is a subfield of artificial intelligence (AI) that focuses on enabling machines to understand, interpret, and generate human language. The goal of NLP is to bridge the gap between human communication, which is often complex, ambiguous, and rich in context, and the ability of computers to process and analyze this language in a meaningful way.

Human language is full of nuances—words can have multiple meanings, sentences can be structured in various ways, and context plays a huge role in interpreting the intended message. NLP aims to allow computers to handle these complexities and perform a variety of language-related tasks, such as:

- **Text Processing:** Understanding and analyzing the structure of text (e.g., breaking down a sentence into words or phrases).
- **Speech Recognition:** Converting spoken language into text.
- **Machine Translation:** Translating text or speech from one language to another (e.g., Google Translate).
- **Sentiment Analysis:** Determining the sentiment or emotional tone of a piece of text (e.g., whether a review is positive or negative).
- **Named Entity Recognition (NER):** Identifying entities such as names, dates, locations, and more within text.
- **Part-of-Speech Tagging:** Identifying the grammatical components of a sentence, such as nouns, verbs, and adjectives.
- **Text Summarization:** Creating a concise summary of a larger text.
- **Question Answering:** Answering questions based on a given text or dataset.

NLP is built on a combination of linguistics (the study of language) and computational techniques. Initially, rule-based methods were used, relying on predefined grammatical rules and human-defined dictionaries. However, modern NLP has largely been revolutionized by machine learning techniques, particularly deep learning, which allows models to learn patterns from vast amounts of text data. This has led to substantial improvements in NLP systems, enabling more accurate and context-aware language processing.<sup>3</sup>

### Key Advancements in NLP

The field of Natural Language Processing (NLP) has made tremendous strides over the past decade, driven by advancements in machine learning, especially deep

learning techniques. These breakthroughs have significantly enhanced the performance and capabilities of NLP systems, enabling them to handle more complex tasks with greater accuracy. Below are some of the most notable advancements in NLP:

### Deep Learning and Neural Networks

One of the most transformative developments in NLP has been the adoption of deep learning models, particularly neural networks. These models are capable of learning from vast amounts of data, recognizing complex patterns, and making decisions based on that information. Before deep learning, NLP relied largely on rule-based methods, which were limited and required extensive human input. With the advent of deep learning, NLP models can now handle vast amounts of unstructured text data more effectively.

### Transformer Models

The introduction of transformer models has revolutionized NLP. Transformers, first introduced in the 2017 paper “Attention is All You Need” by Vaswani et al., replaced the traditional recurrent neural networks (RNNs) and long short-term memory (LSTM) networks. The transformer architecture uses a mechanism called “self-attention,” which allows models to process all parts of a sentence simultaneously rather than sequentially. This makes transformers much more efficient and capable of handling long-range dependencies in text. Key models based on the transformer architecture include:

- **BERT (Bidirectional Encoder Representations from Transformers):** BERT introduced bidirectional training, which means the model looks at both the left and right context of a word to better understand its meaning. This has improved performance in tasks like question answering and text classification.
- **GPT (Generative Pretrained Transformer):** GPT models, including GPT-3 and GPT-4, are generative models that excel at producing human-like text based on a given prompt. They have become renowned for their ability to write coherent essays, summarize texts, and even generate creative content.
- **T5 (Text-to-Text Transfer Transformer):** T5 treats all NLP tasks as text-to-text problems, which has simplified the framework for training models on various tasks.

### Pretrained Language Models and Transfer Learning

Pretrained models like BERT and GPT have significantly changed how NLP tasks are approached. These models are trained on massive corpora of text data and are then fine-tuned for specific tasks like sentiment analysis, translation, or summarization. This approach is known as transfer learning, and it has reduced the need for large amounts of task-specific data, allowing researchers and practitioners to develop NLP systems more efficiently. Pretraining on large

datasets enables these models to learn general language patterns and semantics, making them adaptable to a wide range of applications.

### Multilingual and Cross-lingual Models

NLP has traditionally been dominated by models trained in specific languages, especially English. However, the development of multilingual models like mBERT and XLM-R has allowed for the development of systems that understand and process multiple languages. These models are trained on data from various languages and can perform tasks in languages with fewer resources or where large-scale datasets are not available. Cross-lingual transfer learning allows models to generalize across languages, enabling more accessible NLP tools for global users.

### Zero-shot and Few-shot Learning

Zero-shot learning and few-shot learning have emerged as powerful techniques in NLP. Zero-shot learning refers to the ability of a model to perform tasks without having seen any examples of that task during training. Few-shot learning, on the other hand, enables models to make accurate predictions with only a small number of examples. This has been a breakthrough for tasks where labeled data is scarce or hard to obtain, making NLP models more flexible and versatile.

### Attention Mechanisms

Attention mechanisms have become a cornerstone in many NLP models. These mechanisms allow the model to focus on different parts of the input data when making predictions. For example, in translation tasks, the attention mechanism helps the model focus on relevant words in the source language while generating the translation in the target language. The success of transformers and their attention mechanisms has led to their widespread adoption across various NLP applications.

### Improved Language Models and Fine-Tuning

Recent advancements have also been seen in the fine-tuning of large pre-trained language models. Fine-tuning involves adjusting a pre-trained model's parameters based on task-specific data to improve performance in particular applications, such as sentiment analysis, text summarization, or question answering. Fine-tuning has led to improved performance across various NLP tasks by enabling models to adapt to specific domains or user requirements.

### Multimodal NLP

A newer area of research in NLP involves multimodal models, which combine information from various types of data, such as text, images, and even audio. These models aim to mimic human-like understanding by processing

not just textual data but also visual or auditory input. For example, CLIP (Contrastive Language–Image Pre-training) by OpenAI integrates vision and language to perform tasks like image captioning and image-text matching. This advancement brings NLP closer to understanding the world as humans do, integrating multiple senses.

### Language Generation and Creativity

Advances in language generation, particularly with models like GPT-3 and GPT-4, have demonstrated the ability to produce coherent and contextually relevant text across a wide range of topics. These models can write essays, code, poetry, and even engage in interactive conversations, making them valuable in applications such as content creation, customer support, and educational tools. The ability of models to generate human-like text has opened up new possibilities for creative AI applications.

### Ethical Considerations and Fairness

As NLP models become more powerful, there has been increased focus on the ethical implications of their use. Models trained on large-scale internet data often inherit biases that can manifest in harmful ways. Researchers are increasingly focusing on techniques to detect and mitigate bias, improve model fairness, and make sure that NLP technologies are developed responsibly. Fairness in NLP ensures that systems provide equitable outcomes for people from diverse backgrounds, minimizing harm and reinforcing positive societal values.

### Challenges in NLP

Despite the remarkable progress in Natural Language Processing (NLP), several challenges persist that prevent models from achieving true human-level understanding and processing of language. These challenges arise from the inherent complexity of human language, the limitations of current technology, and ethical considerations that must be addressed as NLP systems are deployed in real-world applications. Below are some of the key challenges that continue to shape the field of NLP:<sup>6,7</sup>

#### Ambiguity in Language

Human language is highly ambiguous. Words and phrases can have multiple meanings depending on the context, which can make accurate interpretation a complex task for machines. For example:

- **Polysemy:** A single word may have multiple meanings (e.g., “bank” can refer to a financial institution or the side of a river).
- **Syntactic Ambiguity:** A sentence can have multiple interpretations depending on its structure (e.g., “I saw the man with the telescope” could mean the man had a telescope or the speaker used a telescope to see the man).

- **Contextual Ambiguity:** The meaning of a word or sentence can change drastically depending on the surrounding context.

Although advances like contextualized word embeddings (e.g., BERT, GPT) have made progress in disambiguating language, these challenges still pose difficulties for NLP systems, particularly in tasks requiring deep contextual understanding.

### Bias and Fairness

NLP models often inherit biases present in the data they are trained on. Since large NLP models are typically trained on vast amounts of text data scraped from the internet, they can pick up on societal stereotypes, harmful language patterns, and biases based on race, gender, age, or other factors. These biases can manifest in a variety of ways, such as:

- **Gender Bias:** NLP systems might associate certain professions with specific genders, such as assuming a nurse is female or a doctor is male.
- **Racial and Ethnic Bias:** Language models may generate biased or discriminatory content based on racial or ethnic backgrounds.
- **Cultural Bias:** Models trained on data from certain cultural contexts may struggle to understand or process information from different cultural settings.

Ensuring fairness in NLP is a critical challenge, as biased models can perpetuate harmful stereotypes or produce unfair outcomes in sensitive applications such as hiring, law enforcement, and healthcare.

### Lack of Large, High-Quality Datasets for Low-Resource Languages

While NLP models have achieved impressive results for high-resource languages like English, there is a significant gap when it comes to low-resource languages. Many languages, especially those spoken by smaller populations or those with limited online presence, lack the large, annotated datasets needed to train sophisticated models. This results in the following issues:

- **Inaccessibility:** NLP tools and services remain largely inaccessible to speakers of low-resource languages.
- **Poor Performance:** Existing models perform poorly or fail entirely when applied to languages with fewer available resources.
- **Multilingual Models:** While multilingual models like mBERT have made strides, they still struggle to provide high-quality language understanding for underrepresented languages.

Addressing these gaps requires new methods for collecting and annotating data, as well as the development of models capable of performing well on languages with limited resources.

### Understanding Deep Context and World Knowledge

While current NLP models, especially transformer-based models like GPT-4, are incredibly good at understanding and generating text, they still lack true understanding. These models excel at identifying patterns in large datasets, but they do not possess real-world knowledge or the ability to reason about the world as humans do. This leads to challenges such as:

- **Factual Accuracy:** NLP systems can generate text that sounds coherent but may contain factual inaccuracies or outdated information, especially when trained on internet data that may be unreliable.
- **Commonsense Reasoning:** Current models struggle with tasks that require commonsense reasoning, such as understanding human emotions, making logical inferences, or comprehending cultural and societal norms.

To address this challenge, researchers are working to integrate structured knowledge bases, external reasoning systems, and dynamic learning methods into NLP models, but true deep comprehension remains a difficult hurdle to overcome.

### Real-Time and Interactive NLP

Many NLP applications require real-time processing, such as in conversational AI systems (e.g., chatbots or virtual assistants). In these settings, it is essential for models to:

- **Respond Quickly:** NLP models must process input and generate responses in real time, often in environments with limited computational resources.
- **Maintain Coherent Dialogue:** In longer conversations, models must track context and ensure that their responses are contextually relevant and coherent over multiple exchanges.

Achieving real-time, contextually aware, and coherent interactions in NLP remains a difficult challenge, particularly when the model is expected to handle complex or open-ended conversations.

### Multimodal Understanding

While NLP traditionally focuses on textual data, a more comprehensive understanding of human communication involves integrating multiple modalities, such as text, images, audio, and video. However, bridging the gap between these different forms of information remains a major challenge:

- **Multimodal Integration:** Combining text and images (e.g., in visual question answering) or text and audio (e.g., in spoken dialogue systems) requires sophisticated models that can understand the relationships between these modalities.



- **Contextual Understanding Across Modalities:** Machines need to integrate knowledge from different sources to accurately interpret multimodal input, which is challenging due to the vast variability in human communication.

Progress is being made with models like OpenAI's CLIP and DALL-E, but truly robust multimodal systems that can perform a wide variety of tasks across different modalities are still in development.

### Ethical Considerations and Accountability

As NLP systems become more pervasive in everyday life, concerns around the ethical implications of their use are gaining prominence. Some key ethical concerns include:

- **Privacy:** NLP systems, especially those integrated with voice assistants and messaging platforms, may inadvertently expose private or sensitive information.
- **Manipulation and Misinformation:** Models capable of generating human-like text can be misused for malicious purposes, such as creating misleading content, deepfakes, or fake news.
- **Transparency and Accountability:** Many NLP models, particularly deep learning models, are viewed as "black boxes," making it difficult to understand how decisions are made. This lack of transparency raises concerns in applications that require accountability, such as in healthcare, law, or finance.

Addressing these ethical concerns requires developing guidelines, regulatory frameworks, and methods for making NLP models more transparent, accountable, and aligned with societal values.<sup>8,9</sup>

### Energy Consumption and Environmental Impact

Large NLP models, especially those based on deep learning, require vast amounts of computational resources for training. This can have significant environmental impacts:

- **High Energy Usage:** Training models like GPT-3 can consume enormous amounts of electricity and computational power, contributing to a significant carbon footprint.
- **Sustainability:** As the scale of NLP models continues to increase, researchers are looking for ways to make them more energy-efficient without sacrificing performance.

Balancing model complexity with sustainability will be crucial for the future development of NLP technologies.

### The Future of NLP

The future of Natural Language Processing (NLP) holds great promise, with advancements that are likely to transform both the technology itself and its applications across various industries. While there are still many challenges to address,

the next decade is poised to bring exciting innovations, particularly with the evolution of AI models, ethical considerations, and the expansion of NLP's capabilities. Below are some key trends and directions shaping the future of NLP:

### Multimodal NLP and Cross-Modal Understanding

As human communication is inherently multimodal (involving both text and other forms of communication such as images, video, and sound), the future of NLP will likely involve multimodal models that integrate various types of data. For instance, combining text, images, and audio can enhance the understanding of meaning by considering context across multiple senses. This could lead to more sophisticated systems capable of tasks like:

- **Visual Question Answering (VQA):** Combining image recognition with language processing to answer questions about visual content.
- **Speech-to-Text with Contextual Understanding:** Systems that integrate speech recognition with real-time contextual understanding across video or audio content, enabling more natural conversations.
- **Content Generation with Rich Media:** Models that can generate text, images, and videos based on the same prompt, producing a more immersive experience.

This shift toward multimodal understanding is expected to bring AI closer to human-like reasoning, making interactions more intuitive and human-centered.

### Zero-shot and Few-shot Learning at Scale

Zero-shot learning (the ability to perform tasks with no prior examples) and few-shot learning (where only a small amount of data is needed) have already shown promising results in NLP. As these techniques improve, we can expect:

- **General-purpose models:** More versatile models that can seamlessly switch between a wide range of tasks, such as answering questions, generating text, translating languages, and summarizing content, with minimal task-specific fine-tuning.
- **More efficient learning:** NLP systems will be able to learn and adapt to new tasks with minimal data, reducing the dependency on large annotated datasets and making AI more accessible to a wider range of applications.
- **Broader applicability:** Zero-shot and few-shot learning will allow NLP models to be applied in a variety of domains, including specialized fields like healthcare, law, and finance, where large labeled datasets may not be readily available.

### Improved Ethical Standards and Bias Mitigation

As NLP systems are increasingly deployed in sensitive areas like hiring, criminal justice, and healthcare, the need for

ethical and transparent AI will become even more pressing. Future NLP research will focus on:

- **Bias mitigation:** Developing models that actively reduce biases in data and outputs. Researchers are exploring techniques like adversarial training, debiasing algorithms, and fairness metrics to address issues related to gender, race, and other forms of discrimination.
- **Explainability and transparency:** Future NLP models will likely become more interpretable, allowing users to understand the reasoning behind AI-generated outputs. This is particularly important in high-stakes applications where accountability is critical.
- **Ethical frameworks:** The development of regulatory frameworks and guidelines will ensure that NLP technologies are used responsibly. This may include setting standards for data collection, model deployment, and ensuring privacy rights are maintained.

### Real-time and Interactive NLP Systems

The next generation of NLP models will likely push the boundaries of real-time processing and interaction. This includes:

- **Conversational AI:** Chatbots and virtual assistants will become increasingly adept at managing multi-turn conversations, understanding complex contexts, and offering personalized interactions.
- **Real-time translation and transcription:** Seamless, real-time language translation will improve communication across linguistic and cultural barriers, revolutionizing sectors like travel, international business, and education.
- **Interactive content creation:** NLP systems will help generate and refine content in real time, such as in creative writing, video editing, or marketing copy creation, allowing users to engage in dynamic content creation.

These advancements will enhance the user experience and make NLP tools more ubiquitous and valuable across various domains.

### Smarter Personal Assistants and Customized NLP Systems

Future NLP systems will likely become more personalized, understanding individual preferences, behavior, and nuances. These systems could:

- **Understand user intent more deeply:** Virtual assistants and chatbots will be able to anticipate and respond to user needs with a much higher degree of accuracy, recognizing tone, context, and past interactions.
- **Customization:** NLP systems will adapt to specific user needs, tailoring responses based on individual

learning styles, habits, or preferences. This could lead to personalized learning tools, health assistants, and customer service experiences.

- **Cross-domain capabilities:** NLP systems will be able to seamlessly switch between multiple domains, providing a more cohesive experience in environments like healthcare (e.g., from scheduling appointments to providing medical advice) or finance (e.g., from managing personal finances to answering complex financial questions).<sup>10</sup>

### Domain-Specific and Expert Systems

While general-purpose models like GPT-4 can generate text on a wide variety of topics, the future of NLP will also see a rise in domain-specific models tailored to particular industries. For example:

- **Healthcare:** NLP models trained on medical literature and patient data could assist doctors in diagnosing diseases, offering treatment suggestions, and helping with medical research.
- **Legal:** NLP models can assist lawyers by automatically reviewing contracts, summarizing legal documents, and even providing insights into case law.
- **Finance:** NLP models can analyze financial statements, detect fraud, and offer personalized financial advice.

These specialized systems will improve productivity and decision-making in highly technical fields by providing more accurate, context-specific insights.<sup>11</sup>

### Resource-Efficient and Sustainable NLP

The energy consumption of large NLP models, especially those requiring extensive computational resources, has raised concerns about their environmental impact. In response, the future of NLP will likely focus on making models more resource-efficient:

- **Smaller, faster models:** There will be an increased focus on developing lightweight models that maintain high performance but require less computational power. Techniques like model pruning, quantization, and knowledge distillation can reduce the size of models while retaining their effectiveness.
- **Green AI:** Efforts to minimize the carbon footprint of training and deploying NLP models will become a priority. More energy-efficient hardware, algorithms, and training processes will reduce the environmental impact of AI research and application.

A balance between performance and sustainability will become essential in ensuring that NLP technology remains environmentally viable in the long term.

### Human-AI Collaboration

The future of NLP will increasingly involve collaborative systems that work alongside humans rather than replace

them. NLP systems will act as tools to augment human abilities, facilitating more efficient workflows and decision-making:

- **Assisting creative professionals:** Writers, marketers, and content creators will use NLP-powered tools to brainstorm, generate content, and refine ideas, while still maintaining creative control.
- **Enhancing productivity:** In sectors like customer service, NLP systems will assist agents by automating repetitive tasks and providing real-time insights, allowing human workers to focus on more complex issues.
- **Decision support:** NLP models will help humans make better decisions by analyzing vast amounts of unstructured data (e.g., medical records, financial reports) and offering relevant insights, which will be especially valuable in fields like healthcare and finance.<sup>12,13</sup>

## Conclusion

The future of Natural Language Processing (NLP) is incredibly promising, with technologies evolving to be more accurate, ethical, sustainable, and adaptable across various domains. While challenges remain—especially in terms of bias, resource consumption, and deep contextual understanding—the direction of NLP research and application points toward a more integrated, intelligent, and accessible future. As NLP continues to advance, its role in improving communication, decision-making, and human-computer interaction will only grow, transforming industries and enhancing our ability to process and understand natural language at scale.

Moreover, the integration of multimodal understanding and cross-domain capabilities will push the boundaries of what's possible, enabling machines to not only understand text but also interact with and interpret the broader world through multiple sensory inputs. These advancements will unlock new opportunities for personalized AI assistants, healthcare innovations, and creative collaborations, leading to smarter and more intuitive systems that better understand human intent.

The future will also bring more ethical considerations to the forefront. As NLP systems become integral to more sensitive and critical applications, ensuring fairness, transparency, and privacy will become increasingly important. Researchers and policymakers will need to collaborate to build safeguards that ensure NLP technologies benefit society while minimizing potential harms. The increasing focus on green AI and energy-efficient systems will also shape the next generation of NLP models, balancing cutting-edge advancements with environmental sustainability.

## References

1. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I. Attention is all you need. *Advances in neural information processing systems*. 2017;30.
2. Devlin J, Chang MW, Lee K, Toutanova K. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)* 2019 Jun (pp. 4171-4186).
3. Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A, Agarwal S. Language models are few-shot learners. *Advances in neural information processing systems*. 2020;33:1877-901.
4. Raffel C, Shazeer N, Roberts A, Lee K, Narang S, Matena M, Zhou Y, Li W, Liu PJ. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*. 2020;21(140):1-67.
5. Schuster S, Manning CD. Enhanced english universal dependencies: An improved representation for natural language understanding tasks. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)* 2016 May (pp. 2371-2378).
6. Conneau A, Khandelwal K, Goyal N, Chaudhary V, Wenzek G, Guzmán F, Grave E, Ott M, Zettlemoyer L, Stoyanov V. Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*. 2019 Nov 5.
7. Radford A, Wu J, Amodei D, Amodei D, Clark J, Brundage M, Sutskever I. Better language models and their implications. *OpenAI blog*. 2019 Feb 14;1(2).
8. Howard J, Ruder S. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*. 2018 Jan 18.
9. Wang J, Huang JX, Tu X, Wang J, Huang AJ, Laskar MT, Bhuiyan A. Utilizing bert for information retrieval: Survey, applications, resources, and challenges. *ACM Computing Surveys*. 2024 Apr 9;56(7):1-33.
10. McCann B, Bradbury J, Xiong C, Socher R. Learned in translation: Contextualized word vectors. *Advances in neural information processing systems*. 2017;30.
11. Wan CX, Li B. Financial causal sentence recognition based on BERT-CNN text classification. *The Journal of Supercomputing*. 2022 Apr 1:1-25.
12. Hendrycks D, Zou A, Mazeika M, Tang L, Li B, Song D, Steinhardt J. Pixmix: Dreamlike pictures comprehensively improve safety measures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 2022 (pp. 16783-16792).
13. Thedinga K, Herwig R. A gradient tree boosting and network propagation derived pan-cancer survival network of the tumor microenvironment. *Iscience*. 2022 Jan 21;25(1).