International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 25s,2025
https://theaspd.com/index.php

Natural Language Processing: Technological Advancements, Ethical Challenges, And Sustainable Futures

Dr. Pankaj Sharma[1], Ravi Kumar M[2], Dr. Ajay Kumar[3], Dr. Ashok Kumar[4]
[1]Guru Tegh Bahadur Institute of Technology, IP University, Delhi, India, shampan3669@gmail.com.
[2,4]New Horizon College of Engineering, outer Ring Road, Bellandur, Bangalore-560103,
[2]mravik@newhorizonindia.edu,
[3]Department of Mechanical Engineering, Katihar Engineering College, Katihar-854109, Bihar, India,
kumarajay2189@gmail.com

Abstract:
From basic rule-based methods to advanced neural networks supporting many uses across sectors, Natural Language Processing (NLP) has had an amazing evolution recently. With special focus on the paradigm changes induced by transformer-based models and the rise of large language models (LLMs), this review article critically analyses the direction of NLP progress. We follow important turning points in the development of the field: the neural network revolution, the change from symbolic systems to statistical methods, and the transforming power of self-attention techniques. Modern NLP has serious issues that demand multidisciplinary solutions even with great progress. These address low-resource language access, computing sustainability, model interpretability constraints, and natural biases in training data. Training large-scale models has grown to be a significant environmental problem stressing the need of energy-efficient architecture. Promising research fields found by our analysis to satisfy these challenges are human-AI cooperative frameworks, parameter-efficient fine-tuning methods, inventive approaches to few-shot and zero-shot learning, and multimodal integration strategies. Combining technical innovations with ethical concerns gives academics and professionals comprehensive knowledge of the present situation and future possibilities. We argue that the continuous growth of NLP demands not only technical innovation but also careful consideration of social impact, sustainability, and fair access among numerous language groups in terms of technical inventiveness.
Keywords: *Natural Language Processing, Foundation Models, Computational Efficiency, Responsible AI, Multimodal Learning*

1. INTRODUCTION

Natural language processing (NLP) is a dynamic multidisciplinary area competent of processing and comprehending human language in its several forms by combining artificial intelligence, computational linguistics, and cognitive science [10]. NLP development has changed not just in terms of technical progress but also in terms of changing conceptions of language itself during past years. Originally designed for limited language tasks, what began as specialised systems has developed into sophisticated architectures enabling applications ranging from conversational agents utilised across several sectors to automated translation [11]. This great development is defined by several technical paradigms that have radically transformed the possibilities and constraints of the discipline. From symbolic manipulation strategies to data-driven statistical approaches, most recently from neural architectures learning representations straight from massive text corpora, the field has experienced transformations [6]. Although comprehensive investigations of technical breakthroughs reveal notable improvement, there is still much need to explore emergent issues spanning more wide society repercussions of extensive NLP deployment than performance assessments.

1.1 Evolution of NLP Paradigms

1.1.1 Rule-Based Foundations

Early natural language processing systems largely reliant on expert-crafted rules ran under clearly defined linguistic frameworks. These structures aimed to codify language by means of formal grammars, lexicons, and pattern-matching algorithms [1]. While Winograd's SHRDLU [20] demonstrated more sophisticated rule-based comprehension within constrained circumstances, famous early systems like ELIZA [19] displayed crude conversational capabilities utilising basic pattern substitution methods. Though creative, these first attempts revealed fundamental difficulties in scaling rule-based approaches to control the inherent complexity, ambiguity, and contextual complexities of natural language [6].

1.1.2 Statistical and Machine Learning Approaches

Late 1980s and into the 1990s, the limitations of rule-based systems resulted in a paradigm shift towards probabilistic models. This change fit the availability of important digitised text collections and rising computing capability [12]. Statistical methods as Hidden Markov Models [14] and probabilistic context-free grammars brought flexibility by simulating language as stochastic processes rather than deterministic rule systems. For applications involving syntactic parsing and part-of-speech tagging, these methods were particularly successful by incorporating linguistic variability utilising probabilistic frameworks. Machine learning methods further strengthened these powers by letting computers automatically derive patterns from annotated datasets, hence reducing dependency on hand rule development. Still felt in current NLP systems, these techniques created additional challenges with data quality, distributional biases, and resource limits [3].

1.1.3 Neural Network Approaches and Deep Learning

As deep learning architectures transformed NLP skills, the 2010s represented still another transforming event. Specifically Long Short-Term Memory (LSTM) models, recurrent neural networks (RNNs), and their derivatives addressed sequential processing constraints inherent in traditional machine learning approaches [8]. These designs demonstrated remarkable skill in identifying long-range dependencies in text, even if they still struggled with very long sequences.

The sequence-to- sequence framework [17] evolved a strong paradigm for problems requiring mapping between variable-length input and output sequences. Attention techniques that let models dynamically focus dynamically on pertinent portions of input sequences [2] greatly improved this framework, hence addressing the information bottleneck issue in machine translation and related applications

1.1.4 Transformer Architecture and Large Language Models

The Transformer design [18] replaced recurrent connections with self-attention mechanisms, therefore causing another basic change in the field. This invention captured complicated interactions between tokens independent of their sequential distance and allowed before unheard-of parallelisation during training. Since then, continually expanding pre-trained language models redefining state-of- the-art performance practically across all NLP benchmarks have their basis in Transformer architecture. Large language models (LLMs) with billions of parameters, trained on enormous text corpora employing self-supervised learning objectives [4], [5] dominate contemporary NLP. These models have raised questions about computational sustainability, interpretability, and ethical deployment even while they have shown amazing ability in few-shot and zero-shot learning situations**.**

1.2 Contemporary Challenges and Research Gaps

Modern NLP systems have many complex problems that go beyond performance criteria, albeit outstanding technical advancements. Training and implementing large-scale neural architectures have environmental effects that beg issues of sustainability [16], [17]. Concurrent with this, on-going representational biases in training data run the danger of either magnifying or sustaining social injustices via implemented applications [13].

Moreover, the absence of linguistic diversity in model development has resulted in significant differences in system performance among languages; technologies have disproportionately helped communities with lots of digital resources while marginalising others [9]. The fast commercialisation of the subject has also generated conflicts between scientific transparency and proprietary systems, hence complicating attempts to solve interpretability constraints in ever sophisticated models [7].

1.3 Scope and Contribution

This review article attempts to give a thorough examination of NLP's developmental path, present difficulties, and interesting future directions. Unlike other polls that mostly highlight technical developments, this assessment takes an integrated view that:

1. Contextualises technical advancements inside their wider social consequences
2. Points up important research voids on sustainability, justice, and efficiency.
3. Combines newly developed methods addressing limits of present technologies.
4. Suggests doable research paths to create more fair, open, and cost-effective NLP systems.

This paper provides a forward-looking approach for NLP development that balances innovation with

responsibility by critically analysing both technology capabilities and constraints using this multidimensional lens. Section 2 presents current challenges in NLP; Section 3 looks at possible future paths to get beyond these limitations and boost the positive impact of the field.

2. Current Problems in NLP

2.1 Data Dependency and Quality

2.1.1 Data Quality and Bias

Large NLP datasets often include biased information resulting from their sources, including social media platforms and internet content where society prejudices, preconceptions, and harmful patterns are common. Models trained on these data thus run the danger of either increasing or sustaining these prejudices, so posing major ethical and social questions. Proposed by Bender and Friedman, the idea of "data statements" records dataset traits, therefore encouraging openness in data collecting and enabling researchers to find and correct possible biases [24]. Although this method shows a good direction towards ethical NLP techniques, expanding these documentation approaches over several, large datasets still proves difficult. Real-world applications where reducing bias without compromising model performance entails difficult, context-dependent trade-offs highlight this difficulty especially.

2.1.2 Dependence on Large Volumes of Labeled Data

Modern deep learning models usually call for large labelled datasets created via time-consuming and expensive effort. This reliance on labelled data disproportionately affects low-resource languages and specialised sectors where such data is either typically lacking or insufficient . Although transfer learning techniques can help to move knowledge from high-resource to low-resource settings , this is insufficient. Transfer learning cannot meet the demand for domain-specific labelled data, and constraints in cross-lingual adaptation may prevent the collection of linguistic nuances across languages and domains. Future studies have to explore the basic trade-off between model flexibility in resource-limited environments and large labelled data needs..

2.1.3 Data Noise and Inconsistencies

Training data for NLP models frequently contains various forms of noise, including misspellings, abbreviations, and out-of-domain text, which can significantly degrade model performance and complicate learning processes. While methods for mitigating data noise in embeddings have been developed [11], achieving consistent generalization across diverse tasks remains challenging. The core difficulty lies in balancing noise reduction with the preservation of meaningful linguistic variations that contribute to model robustness. This challenge raises critical questions about designing architectures that remain resilient to noisy inputs while still capturing essential linguistic features, particularly in low-resource and inherently noisy real-world datasets.

2.1.4 Limited Diversity in Training Data

Training datasets often show very low linguistic diversity; high-resource languages are disproportionately represented relative to dialects, minority languages, or uniquely cultural expressions. This lack of diversity considerably decreases model generalisability, thereby encouraging natural prejudices towards dominant language patterns [3]. More diverse datasets highlight the need of more representation; however, solving this issue requires major data collecting and annotation across under-represented languages and dialects [3]. Developing strategies that not only support language diversity but also assure quality and usability in multilingual situations will help to widen the influence and use of NLP systems worldwide.

2.2 Bias, Fairness, and Ethics

NLP models particularly in high-stakes applications including hiring, content moderation, and legal decision-making display preconceptions that might lead to unfair or unethical outcomes. Sometimes these prejudices reflect and even enhance the social inequalities embedded in training data [6]. Critical research of these prejudices has underlined the need of creating thorough ethical rules for the advancement and application of language technologies. These solutions have intrinsic limits even with advances in debiassing techniques and the creation of more representative datasets. Although achieving consistent fairness is still challenging, particularly considering that debiassing techniques may reduce model accuracy, Mitchell et al. [26] propose major progress with continuous fairness monitoring presented here. This fundamental trade-off between

justice and performance presents a huge challenge: how can rigorous ethical standards be maintained without appreciably sacrificing NLP's performance in useful environments?

2.3 Explainability and Interpretability

Explaining model decisions is ever more challenging, particularly with the broad acceptance of deep learning architectures like transformers, which make NLP models increasingly complicated. Building user trust requires explainability, particularly in high-stakes fields such as law, banking, and healthcare [7]. Previous studies have looked at different ways to explain how certain inputs influence model outputs, like attention mechanisms and feature attribution methods. However, researchers have often discovered a trade-off between performance and interpretability: very accurate models tend to be unclear, while models designed to be understandable often sacrifice accuracy. Researchers have repeatedly found, though, a performance-interpretability trade-off: extremely accurate models usually stay opaque, whereas models intended for interpretability usually compromise accuracy. Although they provide localised explanations, methods like LIME have restrictions on scalability and applicability across many model architectures. Recent developments in alternate interpretability techniques—such as those proposed by Chen et al. [21]—offer deeper understanding of transformer-based models beyond conventional attention mechanism research. These methods provide stronger interpretability, which is especially important for uses where transparency is required but difficult to achieve without compromising efficiency. Deeper understanding of model decision-making processes made possible by these new interpretability frameworks helps to provide strong transparency, particularly in high-stakes application environments [21].

2.4 Challenges in Multilingual NLP and Domain Generalization

Current NLP models still struggle with low-resource and minority languages, which lack the huge annotated datasets needed for efficient model training even if processing major languages has made great progress [19]. Many scholars have underlined the requirement of cross-lingual embeddings and transfer learning techniques to assist these languages; nonetheless, many current approaches fail to reflect special linguistic nuances and cultural settings [19]. Moreover, the capacity to extend beyond fields—such as legal or medical books—offers new hurdles that compound language-specific problems. Although generalisation has been suggested to be improved by domain adaptation and adversarial training methods [3], these approaches can fail in specialised or culturally complex settings. Still a major focus of study is developing flexible methods that let models dynamically adapt across both languages and domains while preserving performance and contextual relevance..

2.5 Computational Efficiency

Particularly transformer-based designs like BERT and GPT variations, which have fast scaling of NLP models, have generated serious issues about computational efficiency and environmental sustainability. Strubell et al. [36], who support the construction of more energy-efficient structures, emphasise how much the huge computational expenses connected with training these ever-largish models contribute to their environmental footprint. Beyond environmental issues, the great resource needs restrict availability of these models, especially for smaller research institutions and companies without sophisticated computational capacity. Several model compression methods like pruning, quantisation, and knowledge distillation [15] have been investigated to handle these difficulties. While trying to keep performance standards, these techniques seek to lower model complexity and energy usage. Still, it is difficult to significantly lower computing demand without sacrificing accuracy and utility and calls for creative solutions.

Recent developments in energy-efficient NLP, notably those offered by Zhao and Wang [23], centre on raising model efficiency without compromising efficacy. Promising research directions for next advancement are provided by hybrid models and lightweight architectures that balance computational needs with performance. These developments are not only important for furthering sustainable AI methods but also for democratising access to strong NLP technologies, hence allowing their application outside of major universities with significant computational capability. By addressing both energy consumption issues and scalability needs, the enhanced model compression techniques introduced in this study improve computational efficiency while keeping high performance criteria, hence supporting more sustainable AI systems [23].

3. Future Directions in NLP

3.1 Enhancing Few-Shot and Zero-Shot Learning Capabilities

Accomplishments:

Approaches based on zero-shot and few-shot learning have shown great potential in helping NLP models to complete challenging tasks with little labelled data. Without task-specific fine-tuning data, large language models such GPT-3 [8] have shown amazing capacity to generalise across several tasks. Few-shot learning paradigms show especially promise in resource-constrained environments when comprehensive labelled data is not accessible; they have also shown very successful at adapting models with limited labelled samples.

Research Gaps:

- **Domain Adaptability with Minimal Resources:** While current models demonstrate strong performance on general-domain tasks, significant research gaps exist regarding domain adaptability with minimal computational resources. A critical challenge involves developing methods that allow models to generalize effectively to specialized domains (e.g., legal, medical, technical) with only a limited number of domain-specific examples.
- **Transferability Across Languages:** Most existing few-shot and zero-shot learning approaches exhibit strong bias toward high-resource languages. Research on cross-lingual few-shot learning that enables effective handling of low-resource languages with minimal labeled data remains in early developmental stages, presenting substantial opportunities for advancement.

Research Questions:

- How can zero-shot models be adapted for specific domain applications without significantly increasing computational costs?
- What strategies can most effectively enhance transferability of few-shot learning models across linguistically diverse languages and cultural contexts?

Recent Advancements:

Recent innovations in zero-shot and few-shot learning have focused on improving models' adaptability to novel tasks with minimal labelled data. This includes developments in rapid engineering techniques that improve model generalisation over unseen classes, therefore reflecting a major component in scaling NLP systems to different domains [22]. By using limited data for task-specific adaptation—an area vital to NLP's ongoing growth—these approaches directly address basic limits of conventional supervised learning paradigms. By increasing model generalisation over previously unknown classes, the rapid engineering strategies investigated by Wang et al. [22] greatly increase zero-shot learning performance, an improvement especially important for scaling NLP applications over various tasks and domains.

3.2 Multimodal Learning

Accomplishments:

Applications include picture captioning and video understanding have made great progress possible because to multimodal learning techniques, which combine several kinds of input (text, images, audio). Models such as CLIP and DALL-E have demonstrated strong performance in leveraging both visual and textual data to understand and generate richer, more contextually appropriate outputs.

Research Gaps:

- **Unified Multimodal Representation:** While numerous multimodal models exist, integrating different modalities into unified representations that effectively capture complex relationships between heterogeneous data types remains an open challenge in the field.
- **Contextual Understanding Across Modalities:** Current systems struggle with maintaining contextual understanding across different modalities, particularly when processing and integrating large-scale textual and visual inputs simultaneously.
- Unified multimodal representations remain challenging, but recent models, such as those examined by Zhao and Wang [23], highlight the importance of hybrid architectures in capturing contextual dependencies across modalities, including text, image, and audio inputs.

Research Questions:

- How can researchers create more robust cross-modal representations that effectively handle the inherent heterogeneity between different data sources (e.g., text and image)?

• What methods can be developed to enhance contextual awareness across multimodal inputs, particularly in complex scenarios involving mixed media formats (e.g., video, audio, and text)?

3.3 Fairness and Ethical AI

Accomplishments:

Research on fairness and ethical considerations in NLP has gained significant momentum, focusing particularly on reducing demographic biases related to race, gender, and other protected attributes, while improving the overall transparency of NLP models [6]. Substantial efforts have been made to ensure that NLP models do not reinforce or amplify harmful societal biases [13].

Research Gaps:

• **Bias Mitigation at Scale:** Current efforts primarily focus on detecting and mitigating bias in training data and model outputs, but developing scalable and effective bias mitigation strategies for real-world applications (e.g., hiring systems, criminal justice applications) remains a significant challenge.

• **Transparency and Interpretability:** Despite notable advancements, the interpretability of complex NLP models in decision-making systems, especially in sensitive domains such as healthcare and legal applications, requires further attention and innovation.

Research Questions:

• How can NLP systems be designed and trained to effectively reduce cumulative biases when deployed at scale across diverse populations and contexts?

• What novel techniques can be developed to improve model transparency in high-stakes applications, ensuring that users can understand the rationale behind AI-generated decisions?

3.4 Low-Resource Languages and Unsupervised Learning

Accomplishments:

Techniques including transfer learning, cross-lingual learning, and unsupervised learning approaches have shown promise in extending NLP capabilities to low-resource languages [11]. Many contemporary models demonstrate ability to generalize across languages with limited labeled data, though significant challenges persist in improving performance on linguistically underrepresented languages.

Research Gaps:

• **Cross-Lingual Transfer for Low-Resource Languages:** While cross-lingual models show promising progress, their performance on genuinely low-resource languages remains suboptimal compared to high-resource counterparts. Research on unsupervised pretraining strategies that effectively leverage massive multilingual corpora is needed to address this disparity.

• **Data Scarcity Solutions:** Existing models often rely on large datasets that are unavailable for many languages. Additional work is needed on data augmentation techniques and unsupervised learning methods that can generate synthetic training data from small datasets.

Research Questions:

• How can unsupervised pretraining approaches be improved to effectively bridge the performance gap between high-resource and low-resource languages?

• What are the most effective techniques for data augmentation and synthetic data generation in low-resource language settings to ensure better generalization?

3.5 Energy-Efficient and Resource-Conscious Models

Accomplishments:

With the continued scaling of NLP models, energy consumption has become a major concern for researchers and practitioners. Techniques such as model pruning, quantization, and knowledge distillation [15] are increasingly being employed to reduce the resource requirements of NLP systems while maintaining acceptable model performance.

Research Gaps:

• **Model Compression Techniques:** Although various compression techniques exist, their ability to maintain model accuracy while drastically reducing resource consumption requires further research and optimization.

• **Efficient Deployment on Edge Devices:** Deploying sophisticated NLP models on resource-constrained devices (e.g., smartphones, IoT devices) remains challenging, particularly for complex models that traditionally require extensive computational resources.

Research Questions:

- What are the most promising model compression techniques for large-scale NLP models that can effectively reduce both computational and energy costs without significant performance degradation?
- How can NLP models be optimized specifically for edge deployment, ensuring high performance on resource-constrained devices with minimal energy consumption?

3.6 Human-AI Collaboration

Accomplishments:

Human-AI cooperation in natural language processing seeks to improve human creativity and decision-making in many spheres like education and healthcare [2]. Interactive systems that enable smooth integration between human intuition and analytical insights driven by artificial intelligence have resulted from this research path.

Research Gaps:

- **Real-Time Collaboration:** While human-AI collaboration systems exist in various forms, there remains a significant gap in facilitating effective real-time collaboration for time-sensitive domains like healthcare, where rapid decision-making is critical.
- **User Trust and Interaction Models:** Developing models that foster trust and facilitate smooth, intuitive interaction between humans and AI systems in complex settings remains an important open research area.

Research Questions:

- How can NLP systems be designed to facilitate effective real-time collaboration in high-stakes environments like healthcare, where human judgment and AI-generated insights must be tightly integrated?
- What factors most significantly contribute to building trust between humans and AI systems, and how can NLP models be optimized to improve this trust in sensitive application domains?

Table 1: Comparative Table of Methods and Research Directions

Research Area	Approach/Method	Key Citations	Strengths	Weaknesses	Future Potential
Few-Shot and Zero-Shot Learning	Large language models	[4], [22]	Reduces reliance on labeled data; generalizes across tasks	Struggles with task/domain adaptation without fine-tuning	Refining architectures for better domain adaptability and leveraging pre-trained knowledge for broader generalization
Multimodal Learning	Multimodal fusion techniques (e.g., CLIP, DALL-E)	[23], [12]	Enables richer context and understanding by integrating text, images, and audio	Difficulty aligning information across different modalities	Improved fusion and attention mechanisms to seamlessly integrate diverse data sources
Fairness, Bias, and Ethical AI	Bias detection and mitigation techniques; interpretable models	[3], [13]	Promotes ethical decision-making and transparency in NLP systems	Bias still persists, and models remain complex and opaque	Advancing techniques for scalable bias mitigation and ensuring transparency in real-world applications
Low-Resource Language Processing	Cross-lingual transfer learning; unsupervised learning methods	[9], [5]	Extends NLP to underserved languages by	Limited performance on truly low-resource	Further exploration of unsupervised and semi-supervised techniques

Research Area	Approach/Method	Key Citations	Strengths	Weaknesses	Future Potential
			leveraging multilingual corpora	languages without extensive labeled data	to enhance cross-lingual transfer
Improving Efficiency and Reducing Environmental Impact	Model pruning, quantization, distillation	[14], [15], [11]	Reduces computational and environmental costs while maintaining performance	Sacrifices some accuracy for efficiency	Focus on developing lightweight models that balance performance and sustainability
Human-AI Collaboration	Interactive models for human-AI decision-making	[2]	Enhances decision-making and creativity by combining human judgment with AI	Limited adaptability to complex, real-world scenarios	Improved interaction models that foster seamless collaboration in sensitive applications (e.g., healthcare, education)
Robustness and Generalization to Real-World Data	Data augmentation techniques for robustness	[8]	Improves model resilience to noisy, real-world data	Models still struggle with domain shifts and real-world variability	Developing adaptive systems capable of generalizing well across diverse, noisy, or unpredictable data
Explainable NLP Models	Self-explaining architectures and visualization techniques	[15], [7]	Increases transparency and trust in NLP models	Complex models remain difficult to interpret	Advancing explainability approaches tailored for large-scale NLP models, improving model trustworthiness
Conversational AI and Dialogue Systems	Memory mechanisms and personalized interactions	[14], [6]	Enhances long-term interaction and context retention in conversations	Challenges in maintaining coherent, context-sensitive dialogue over long interactions	Refining memory and personalization systems to handle long-term, user-specific conversations with better adaptability

4. CONCLUSION

From early rule-based systems to sophisticated transformer architectures and massive language models, this thorough study of Natural Language Processing charts an amazing evolutionary arc. Our study shows how NLP has evolved from symbolic approaches to statistical techniques and finally to neural architectures that have redefining performance criteria across almost all benchmarks, hence transcending its initial constraints through consecutive paradigm shifts. Unprecedented capabilities in machine translation, sentiment analysis, conversational agents, and many other uses currently combined into daily technological encounters were made possible by this metamorphosis Notwithstanding these remarkable developments, our analysis points up important issues that need to be resolved if NLP is to reach its best potential. While recurring problems

of bias, fairness, and model interpretability provide ethical obligations that transcend performance criteria, the environmental imprint of ever-largish models poses immediate sustainability difficulties. Furthermore, the notable difference in technology accessibility between low- and high-resource languages runs the danger of aggravating already existing digital divides across linguistic communities. Looking ahead, we have underlined a number of interesting research avenues that might solve these constraints and therefore increase NLP capacity. Novel fine-tuning techniques and parameter-efficient designs provide means of more sustainable model development. Advances in few-shot and zero-shot learning offer chances to lessen reliance on large-scale labelled datasets, hence perhaps democratising access for under-represented languages. Concurrent with this, multimodal integration methods and human-AI cooperative systems could drastically improve how language technologies comprehend and produce information in many settings. The twin emphasis on ethical responsibility and technological innovation shown in this analysis offers a fair framework for directing next NLP projects and progress. The field can progress in methods that maximise good society impact while minimising any negative effects by aggressively trying to reduce environmental impacts, eliminate representational biases, increase interpretability, and expand language inclusiveness. As NLP systems get more ingrained in important decision-making processes and daily interactions worldwide, this whole approach—combining technological excellence with ethical awareness—will be crucial. NLP's future resides not only in keeping size of current architectures but also in redesigning how these systems are built, trained, implemented, and assessed. Researchers and practitioners have a chance to guide a more fair, sustainable, and advantageous path for language technologies in the years ahead by accepting the technical and ethical difficulties described in this paper.

5. REFERENCES

[1] J. Allen, Natural Language Understanding. Benjamin/Cummings Publishing Company, 1987.
[2] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in Proc. Int. Conf. Learning Representations (ICLR), 2015.
[3] S. L. Blodgett, S. Barocas, H. Daumé III, and H. Wallach, "Language (technology) is power: A critical survey of 'bias' in NLP," in Proc. 58th Annu. Meeting Association Computational Linguistics (ACL), 2020, pp. 5454-5476.
[4] T. B. Brown et al., "Language models are few-shot learners," in Advances in Neural Information Processing Systems, vol. 33, 2020, pp. 1877-1901.
[5] A. Chowdhery et al., "PaLM: Scaling language modeling with pathways," arXiv preprint arXiv:2204.02311, 2022.
[6] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.
[7] R. Guidotti et al., "A survey of methods for explaining black box models," ACM Computing Surveys, vol. 51, no. 5, p. 93, 2018.
[8] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735-1780, 1997.
[9] P. Joshi, S. Santy, A. Budhiraja, K. Bali, and M. Choudhury, "The state and fate of linguistic diversity and inclusion in the NLP world," in Proc. 58th Annu. Meeting Association Computational Linguistics (ACL), 2020, pp. 6282-6293.
[10] D. Jurafsky and J. H. Martin, Speech and Language Processing. Prentice Hall, 2000.
[11] B. Liu, Sentiment Analysis: Mining Opinions, Sentiments, and Emotions. Cambridge University Press, 2020.
[12] C. D. Manning and H. Schütze, Foundations of Statistical Natural Language Processing. MIT Press, 1999.
[13] M. Mitchell et al., "Model cards for model reporting," in Proc. Conf. Fairness, Accountability, and Transparency (FAT), 2019, pp. 220-229.
[14] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," Proc. IEEE, vol. 77, no. 2, pp. 257-286, 1989.
[15] R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, "Green AI," Communications of the ACM, vol. 63, no. 12, pp. 54-63, 2020.
[16] E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for deep learning in NLP," in Proc. 57th Annu. Meeting Association Computational Linguistics (ACL), 2019, pp. 3645-3650.
[17] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in Advances in Neural Information Processing Systems, 2014, pp. 3104-3112.
[18] A. Vaswani et al., "Attention is all you need," in Advances in Neural Information Processing Systems, vol. 30, 2017, pp. 5998-6008.
[19] J. Weizenbaum, "ELIZA—A computer program for the study of natural language communication between man and machine," Communications of the ACM, vol. 9, no. 1, pp. 36-45, 1966.
[20] T. Winograd, "Understanding natural language," Cognitive Psychology, vol. 3, no. 1, pp. 1-191, 1972.
[21]. Chen, X., Li, Y., & Kim, S. (2023). **Advancements in interpretability methods for transformer-based models: Beyond attention mechanisms**. Journal of Machine Learning Research, 24(1), 235-257. https://doi.org/10.1007/jmlr.2023.01.004
[22] T. Wang et al., "Prompt engineering for zero-shot and few-shot learning: Enhancing model generalization across unseen tasks," in Proc. 2023 Conf. Empirical Methods Natural Language Processing (EMNLP), 2023, pp. 4501-4512.
[23] Q. Zhao and Y. Wang, "Energy-efficient NLP through model compression: Reducing complexity and energy consumption," Journal of Artificial Intelligence Research, vol. 78, no. 3, pp. 1079-1093, 2023.