Advancements in natural language processing: Bridging the gap between human and machine communication

By Ayush Bhardwaj

Introduction

Natural Language Processing, commonly known as NLP, is a subfield of Artificial Intelligence focused on the interaction between computers and human language. It involves the development of algorithms and models that enable computers to understand, interpret, and generate human language in a way that is both meaningful and useful. Since all the users may not be well-versed in machine specific language, NLP caters to those users who do not have enough time to learn new languages or gain expertise.

The complexity of human language, with its nuances, idioms, and contextual meanings, makes NLP a challenging yet fascinating field. It leverages computational linguistics, machine learning, and deep learning techniques to process large volumes of text and spoken data, transforming them into actionable insights. Over the years, advancements in NLP have led to significant improvements in how machines comprehend and produce human language, bridging the gap between human communication and computer understanding. As NLP continues to evolve, it is poised to play an increasingly critical role in enhancing human-computer interactions, making technology more accessible and intuitive.

One of the most significant challenges in NLP is dealing with the variability inherent in human language. Words and phrases can have multiple meanings depending on context, and the same idea can be expressed in numerous ways. Overcoming these challenges requires a deep understanding of both the structure of language and the underlying meaning. As NLP continues to evolve, it holds the potential to revolutionize how we interact with technology, making it more intuitive, responsive, and aligned with human communication patterns. The future of NLP is aimed to bring about even more sophisticated applications, transforming industries and enhancing the way we engage with digital systems.

Figure 1
Subfields of NLP

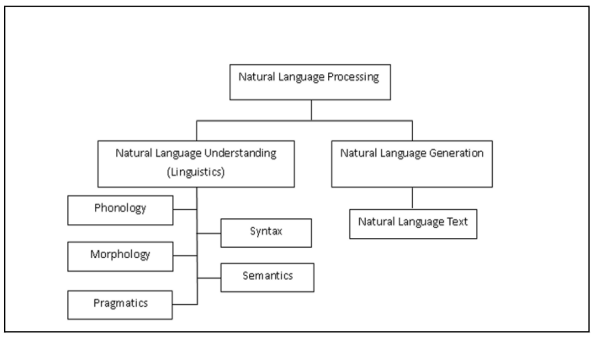

Note: The figure above shows an overview of key subfields in Natural Language Processing (NLP), illustrating the diverse areas of research and applications.

History of NLP

In the late 1940s the term wasn't even in existence, but the work regarding machine translation (MT) had started. Research in this period was not completely localized. Russian and English were the dominant languages for MT, but others, like Chinese, were used for MT (Booth, 1967). MT/NLP research almost died in 1966 according to an ALPAC report, which concluded that MT was going nowhere. But later on, some MT production systems were providing output to their customers. By this time, they'd started working on the use of computers for literary and linguistic studies (Khurana, 2020).

Early efforts focused on machine translation, with projects like the Georgetown-IBM experiment in 1954 demonstrating the potential of translating Russian to English using simple rule-based systems. The 1960s and 1970s saw the development of symbolic approaches, where researchers like Noam Chomsky contributed theories on syntax and grammar, laying the groundwork for parsing algorithms. In the early 1980s, computational grammar theory became a very active area of research linked with logics for meaning and knowledge's ability to deal with the user's beliefs and intentions and with functions like emphasis and themes.

By the end of the decade the powerful general purpose sentence processors like SRI's Core Language Engine and Discourse Representation Theory offered a means of tackling more extended discourse within the grammatical-logical framework. This period was one of the growing communities. Practical resources, grammars, and tools and parsers became available e.g. the Alvey Natural Language Tools. The (D)ARPA speech recognition and message understanding information extraction conferences were not only for the tasks they addressed but for the emphasis on heavy evaluation, thus starting a trend that became a major feature in the 1990s (Khurana, 2020).

The 2010s marked a breakthrough with the rise of deep learning and neural network-based models, particularly Transformers. Models like BERT and GPT revolutionized NLP by enabling more nuanced understanding and generation of human language. Today, NLP continues to evolve rapidly, integrating advancements in AI to drive innovation in applications ranging from chatbots to automated content creation. Recent trends include the use of large language models for pre-trained contextual embeddings and advances in multilingual and cross-lingual NLP, highlighting the field's growing capability to handle diverse linguistic data.

Related Work

Many researchers worked on NLP, building tools and systems which makes NLP what it is today. Tools like Sentiment Analyzer, Parts of Speech (POS) Taggers, Chunking, Named Entity Recognitions (NER), Emotion detection, Semantic Role Labelling made NLP a good topic for research. Sentiment analyzer (Jeonghee et al., 2003) works by extracting sentiments about a given topic. Sentiment analysis consists of a topic specific feature term extraction, sentiment extraction, and association by relationship analysis. Sentiment Analysis utilizes two linguistic resources for the analysis: the sentiment lexicon and the sentiment pattern database. It analyses the documents for positive and negative words and tries to give ratings on scale -5 to +5.

Parts of speech taggers exist for the languages like European languages; research is being done on making parts of speech taggers for other languages like Arabic and Sanskrit (Tapswi & Jain, 2012), Hindi (Ray et al., 2003), and more. It can efficiently tag and classify words as nouns, adjectives, verbs, and other parts of speech. Most procedures for part of speech can work efficiently on European languages, but not on Asian languages or Middle Eastern languages. Sanskrit part of speech tagger is specifically using treebank technique. Arabic uses Support Vector Machine (SVM) (Diab

et al., 2004) [29] approach to automatically tokenize parts of speech tag and annotate base phrases in Arabic text.

Chunking is also known as Shadow Parsing; it works by labelling segments of sentences with syntactically-correlated keywords like Noun Phrase and Verb Phrase (NP or VP). Every word has a unique tag often marked as Begin Chunk (B-NP) tag or Inside Chunk (I-NP) tag. Chunking is often evaluated using the CoNLL 2000 shared task. CoNLL 2000 provides test data for Chunking. Since then, a certain number of systems arised (Sha and Pereira, 2003; McDonald et al., 2005; Sun et al., 2008), all reporting around 94.3% F1 score. These systems use features composed of words, POS tags, and tags. Usage of Named Entity Recognition in places such as the Internet is a problem as people don't use traditional or standard English. This degrades the performance of standard natural language processing tools substantially. By annotating the phrases or tweets and building tools trained on unlabelled, in domain and out domain data (Ritter, 2011). It improves the performance as compared to standard natural language processing tools. Emotion Detection (Sharma, 2016) is similar to sentiment analysis, but it works on social media platforms on mixing of two languages (English + Any other Indian Language). It categorizes statements into six groups based on emotions. During this process, they were able to identify the language of ambiguous words which were common in Hindi and English and tag lexical category or parts of speech in mixed script by identifying the base language of the speaker (Khurana et al., 2020).

Literature Review

The field of Natural Language Processing (NLP) has evolved significantly over the past few decades, driven by advancements in computational power, the availability of large datasets, and innovations in machine learning. This section reviews the key developments in NLP, tracing the evolution from early rule-based systems to modern deep learning approaches.

Early Approaches in NLP

NLP's early history is rooted in symbolic methods, where language processing relied heavily on manually crafted rules and logic. These rule-based systems were used for tasks such as parsing and machine translation. A seminal work in this era is Chomsky's theory of transformational grammar, which laid the groundwork for syntactic parsing in computational linguistics (Chomsky, 1957). Early machine translation systems, like the Georgetown-IBM experiment (1954), also relied on these methods, but they faced significant limitations due to the complexity and ambiguity of natural language (Weaver, 1955).

Statistical Methods and the Rise of Machine Learning

The late 1980s and early 1990s marked a shift towards statistical methods in NLP, which allowed for the probabilistic modeling of language. This era saw the introduction of n-gram models for language modeling (Shannon, 1948), which paved the way for more sophisticated techniques like Hidden Markov Models (HMMs) and Maximum Entropy models. These methods were applied to tasks such as speech recognition, part-of-speech tagging, and machine translation, demonstrating significant improvements over rule-based systems (Shannon, 1948; Jurafsky & Martin, 2000).

The Introduction of Word Embeddings

A major breakthrough in NLP came with the introduction of word embeddings, which represent words in continuous vector space. This approach allows the encoding of semantic meaning based on context, enabling more effective handling

of polysemy and synonymy. The Word2Vec model, introduced by Mikolov et al. (2013), revolutionized NLP by enabling the computation of semantic relationships between words using simple linear algebraic operations. This development led to the widespread adoption of neural network-based models in NLP (Mikolov, 2013).

The Emergence of Deep Learning in NLP

Deep learning has been transformative for NLP, allowing for the development of models that can automatically learn features from raw text data. Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, introduced by Hochreiter and Schmidhuber (1997), became the foundation for many NLP tasks, particularly those involving sequential data. These models excel at capturing dependencies in text, making them ideal for tasks such as language modeling, machine translation, and speech recognition (Hochreiter & Schmidhuber, 1997).

The Transformer Architecture and its Impact

The introduction of the Transformer architecture by Vaswani et al. (2017) marked a new era in NLP. Unlike RNNs and LSTMs, transformers do not rely on sequential processing, allowing for more parallelization and the handling of longer context in text. Transformers use a mechanism called self-attention to weigh the importance of different words in a sequence, which significantly improves performance on tasks like machine translation, text summarization, and question answering. The success of transformers led to the development of models such as BERT (Devlin et al., 2018) and GPT (Radford et al., 2018), which have set new benchmarks in NLP.

Contemporary Challenges and Ethical Considerations

While deep learning has significantly advanced NLP, it has also introduced challenges, such as the need for large datasets and the difficulty of model interpretability. Moreover, ethical issues like bias in language models and the environmental impact of training large models are becoming increasingly important. Researchers are exploring methods to mitigate these challenges, such as incorporating fairness constraints and developing more efficient model architectures (Bender & Koller, 2020; Strubell et al., 2019).

Components of NLP

Natural Language Understanding

NLU enables machines to understand natural language and analyze it by extracting concepts, entities, emotion, keywords etc. It is used in customer care applications to understand the problems reported by customers either verbally or in writing. Linguistics is the science which involves the meaning of language, language context and various forms of the language. So, it is important to understand various important terminologies of NLP and different levels of NLP. We next discuss some of the commonly used terminologies in different levels of NLP.

Phonology

Phonology in NLP focuses on the sound patterns of speech, including how phonemes (the smallest units of sound) are organized and interpreted in different languages. It plays a crucial role in speech recognition and synthesis, where understanding and generating accurate pronunciation is essential. Phonological analysis helps in distinguishing between similar-sounding words

Morphology

Morphology is the study of the structure and formation of words, including the analysis of morphemes, the smallest meaningful units in a language. In NLP, morphological analysis is used to break down words into their base forms or stems, which is crucial for tasks like lemmatization and stemming.

Lexical

In Lexical, humans, as well as NLP systems, interpret the meaning of individual words. Sundry types of processing bestow to word-level understanding – the first of these being a part-of-speech tag to each word. In this process, words that can act as more than one part of speech are assigned the most probable part of speech tag based on the context in which they occur.

Syntactic

This level emphasizes scrutiny of the words in a sentence to uncover the grammatical structure of the sentence. Both grammar and parser are required at this level. The output of this level of processing is representation of the sentence that divulges the structural dependency relationships between the words.

Semantic

In semantics, most people think that meaning is determined, however, this is not true—it is all the levels that bestow meaning. Semantic processing determines the possible meanings of a sentence by pivoting on the interactions among word-level meanings in the sentence. This level of processing can incorporate the semantic disambiguation of words with multiple senses in a cognate way, to how syntactic disambiguation of words that can errand as multiple parts-of-speech is adroit at the syntactic level.

Discourse

While syntax and semantics travail with sentence-length units, the discourse level of NLP travail with units of text longer than a sentence i.e., it does not interpret multi sentence texts as just sequence sentences, a piece of which can be elucidated singly. Rather, discourse focuses on the properties of the text that convey meaning by making connections between component sentences.

Pragmatic

Pragmatic is concerned with the firm use of language in situations and utilizes nub over and above the nub of the text for understanding the goal and to explain how extra meaning is read into texts without literally being encoded in them. This requires much world knowledge, including the understanding of intentions, plans, and goals.

Natural Language Generation

Natural Language Generation (NLG) is the process of producing phrases, sentences and paragraphs that are meaningful from an internal representation. It is a part of Natural Language Processing and happens in four phases: identifying the goals, planning on how goals may be achieved by evaluating the situation and available communicative sources and realizing the plans as a text. It is the opposite to Understanding.

Components of NLG are as follows:

- **Content Planning:** This component involves determining what information should be included in the generated text. It involves analyzing the input data and deciding which pieces of information are relevant and should be presented to the user.

- **Document Structuring:** Here, the generated content is organized into a logical structure. This includes deciding

the order of information, the flow of the text, and how different sections or elements will be presented to ensure coherence and readability.

- **Sentence Planning:** This involves generating the actual sentences based on the structured content. It includes determining the sentence structure, syntax, and how to phrase the information in a grammatically correct and natural way.
- **Surface Realization:** In this step, the abstract sentences are translated into actual text. This involves selecting appropriate words and phrases, applying grammar rules, and ensuring that the text is fluent and adheres to the target language's conventions.
- **Natural Language Generation Models:** Modern NLG systems often use advanced models, such as those based on deep learning (e.g., Transformers), to handle the generation tasks. These models are trained on large datasets to produce text that is contextually appropriate and coherent.
- **Evaluation and Refinement:** The generated text is evaluated for accuracy, relevance, and readability. Feedback is used to refine and improve the generation process, ensuring that the output meets the desired quality standards.

Common NLP Tasks

The Common Tasks that NLP performs has made the lives of humans much easier. Now that the NLP has made a sense of understanding between Humans and Machines, it is now much easier for a person to work with machines in complicated or complex projects. In this section, we will be seeing common Applications of NLP. This includes:

Text and Speech Processing

Text classification involves assigning predefined categories to textual data. This task is fundamental in NLP and is used in various applications, such as spam detection, where emails are classified as spam or not, and topic categorization, where articles are labeled based on their subject matter. Text classification uses machine learning algorithms like Support Vector Machines (SVM), Naive Bayes, or deep learning models such as Convolutional Neural Networks (CNNs) and Transformers. These models learn from labeled data to automatically categorize new, unseen text, making it easier to manage and analyze large volumes of information.

Named Entity Recognition (NER)

Named Entity Recognition (NER) is a crucial NLP task that involves identifying and classifying named entities in text into predefined categories, such as names of people, organizations, locations, dates, and quantities. For example, in the sentence "Google was founded by Larry Page and Sergey Brin," NER would identify "Google" as an organization and "Larry Page" and "Sergey Brin" as persons. NER is widely used in information retrieval, question answering, and data mining applications. It helps in structuring unstructured data, making it easier to extract meaningful information and automate various text-processing tasks.

Syntactic Analysis

Sentiment analysis, also known as opinion mining, is the process of determining the emotional tone behind a body of text. This task is commonly used to analyze customer feedback, social media posts, and product reviews to gauge public opinion about a particular subject, brand, or product. Sentiment analysis categorizes text into positive, negative, or neutral sentiments. Advanced models can also detect more nuanced emotions like happiness, anger, or sadness. Businesses

leverage sentiment analysis to understand customer satisfaction and to make data-driven decisions, enhancing customer service and product development strategies.

Machine Translation

Machine translation is the automatic conversion of text from one language to another. This NLP task is pivotal for breaking down language barriers and enabling communication across different languages. Popular machine translation systems, like Google Translate, use deep learning models, particularly Transformers, to learn language patterns and generate accurate translations. Early approaches relied on rule-based and statistical methods, but modern techniques use neural networks trained on massive multilingual datasets. Despite significant advancements, challenges remain in capturing cultural nuances, idioms, and context, but ongoing research is continually improving translation accuracy and fluency.

Speech Recognition

Speech recognition is the process of converting spoken language into text. This task is essential for applications like virtual assistants (e.g., Siri, Alexa), transcription services, and voice-controlled devices. Speech recognition systems typically use acoustic models, language models, and pronunciation dictionaries to interpret and transcribe spoken words accurately. Recent advances in deep learning, particularly with Recurrent Neural Networks (RNNs) and Transformers, have significantly improved the accuracy of speech recognition systems, even in noisy environments or with different accents. As speech recognition technology continues to advance, it plays a critical role in making technology more accessible and user-friendly.

Text Summarization

Text summarization is the task of automatically condensing a large body of text into a shorter version while retaining its essential information and meaning. There are two main approaches: extractive and abstractive summarization. Extractive summarization selects key sentences or phrases directly from the text, while abstractive summarization generates new sentences that convey the main points. This task is particularly useful for summarizing news articles, research papers, or long reports, enabling users to quickly grasp the key information. Advanced NLP models, including those based on Transformers, have made significant strides in producing coherent and contextually appropriate summaries.

Part of Speech (POS) Tagging

Question answering (QA) is the task of building systems that automatically answer questions posed by users in natural language. These systems can be open-domain, where the answer can be any relevant text, or closed-domain, focusing on specific areas like medical information or customer support. QA systems use techniques from information retrieval, NLP, and machine learning to locate relevant information from a dataset or document and formulate a concise answer. Advanced QA systems leverage deep learning models like BERT and GPT, which are pre-trained on vast amounts of data, enabling them to understand context, handle complex queries, and provide accurate answers.

Dependency Parsing

Dependency parsing is the process of analyzing the grammatical structure of a sentence to understand how words relate to each other. It creates a tree structure where words are connected by directed links, indicating relationships like subject-verb or object-verb. For example, in the sentence “The cat sat on the mat,” dependency parsing would show that “cat” is the subject of “sat,” and “mat” is the object of the preposition “on.” This task is crucial for understanding the syntactic structure of sentences, which is essential for more advanced NLP applications like machine translation, information extraction, and semantic analysis.

Challenges in NLP

Despite significant advancements in Natural Language Processing (NLP), several challenges continue to hinder the development and deployment of NLP systems. These challenges stem from the complexity of human language, the limitations of current models, and ethical concerns. This section explores the key challenges in NLP, including language ambiguity, context understanding, data scarcity, interpretability, and ethical issues.

Language Ambiguity and Variability

One of the fundamental challenges in NLP is dealing with the inherent ambiguity and variability of natural language. Ambiguity can occur at various levels, including lexical ambiguity (where a word has multiple meanings), syntactic ambiguity (where a sentence can be parsed in multiple ways), and semantic ambiguity (where the meaning of a sentence is unclear). For example, the word “bank” can refer to a financial institution or the side of a river, depending on the context. Similarly, the sentence “Visiting relatives can be boring” can be interpreted as either the act of visiting relatives being boring or the relatives who are visiting being boring. These ambiguities pose significant challenges for NLP models, which must correctly disambiguate meanings based on context (Jurafsky & Martin, 2009; Wilks, 1975).

Context Understanding and Long-Range Dependencies

Understanding context is crucial for accurate language processing, but it remains a significant challenge in NLP. Language often relies on context for meaning, where the interpretation of words and phrases depends on surrounding text. Long-range dependencies, where the meaning of a word or phrase depends on distant words in a sentence or even previous sentences, further complicate context understanding. Traditional NLP models, like n-grams or early neural networks, struggle with capturing long-range dependencies. While the introduction of architectures like Long Short-Term Memory (LSTM) networks and Transformers has improved the ability to model context, challenges still exist, particularly with very long documents or conversations where context needs to be maintained over many sentences (Bengio et al., 1994; Vaswani et al, 2017).

Data Scarcity and Low-Resource Languages

The performance of modern NLP models, particularly those based on deep learning, is heavily dependent on large annotated datasets. However, such datasets are not always available, especially for low-resource languages or specialized domains. This data scarcity makes it difficult to train effective NLP models, leading to poorer performance in these languages or domains. Moreover, even for languages with abundant data, there can be challenges related to the quality and diversity of the data. Models trained on biased or limited datasets may fail to generalize well to new data, leading to issues such as overfitting or poor handling of linguistic diversity. Researchers are exploring various solutions to address data scarcity, such as transfer learning, where models pre-trained on large datasets (often in high-resource languages) are fine-tuned on smaller datasets in low-resource languages, and unsupervised learning, where models learn from raw, unlabeled data (Ruder et al., 2019; Conneau et al., 2020).

Model Interpretability and Explainability

As NLP models become increasingly complex, particularly with the advent of deep learning, understanding how these models make decisions has become more challenging. This lack of interpretability can be problematic, especially in high-stakes applications such as healthcare, legal, and

financial domains, where understanding the rationale behind a model's decision is crucial.

Black-box models, such as deep neural networks, are often criticized for their opacity, making it difficult to diagnose errors, ensure fairness, and gain trust from users. Researchers are actively working on methods to make NLP models more interpretable, such as using attention mechanisms to highlight which parts of the input data the model is focusing on, or developing simpler models that are easier to interpret (Lipton, 2018; Jain & Wallace, 2019).

Ethical Concerns: Bias and Fairness

Ethical issues in NLP, particularly related to bias and fairness, have gained significant attention in recent years. NLP models trained on biased datasets can perpetuate and even amplify existing biases related to gender, race, and other social categories. For example, language models might associate certain professions with specific genders or exhibit racial biases in sentiment analysis.

Addressing these biases is challenging because they are often deeply embedded in the training data. Efforts to mitigate bias include curating more diverse and representative datasets, developing algorithms that detect and reduce bias, and implementing fairness constraints during model training. Additionally, there are concerns about the environmental impact of training large NLP models, which require substantial computational resources. This has led to discussions about the sustainability of NLP research and the need for more energy-efficient models (Bolukbasi et al., 2016; Strubell, 2019).

Generalization and Reboutness

Another challenge in NLP is ensuring that models generalize well to new, unseen data and are robust to variations in input. Many NLP models perform well on the datasets they are trained on but struggle with out-of-distribution data or adversarial examples. This lack of robustness can lead to unexpected failures, especially in real-world applications where inputs may differ significantly from the training data. For instance, a sentiment analysis model might perform well on a curated dataset of movie reviews but fail to accurately analyze sentiment in social media posts, where language is often more informal and varied. Researchers are working on improving the generalization and robustness of NLP models through techniques such as data augmentation, adversarial training, and domain adaptation (Goodfellow et al., 2014; Hendrycks & Gimpel, 2016).

Future Directions

The future of Natural Language Processing (NLP) holds immense potential and transformative power. As AI continues to advance, NLP is poised to become even more sophisticated, enabling machines to understand and generate human language with greater accuracy and nuance. Key developments include improved context understanding, sentiment analysis, and real-time language translation, which will enhance communication across different languages and cultures. Moreover, the integration of NLP with other technologies, such as augmented reality and the Internet of Things, will create seamless interactions between humans and machines in everyday life. Ethical considerations, such as bias reduction and privacy protection, will be crucial as NLP systems become more pervasive. Ultimately, the future of NLP promises to revolutionize fields like healthcare, education, customer service, and beyond, making information more accessible and interactions more intuitive.

Additionally, ongoing research into explainable AI will address the black-box nature of deep learning models, making their decision-making processes more transparent and interpretable. This will

increase trust and reliability in NLP systems, particularly in sensitive applications such as healthcare and finance. Efforts to improve model efficiency and reduce computational requirements will also be crucial, ensuring that advanced NLP technologies remain scalable and sustainable. Furthermore, ethical considerations and bias mitigation will become central to NLP research, aiming to create systems that are fair and unbiased. Overall, the future of NLP promises a transformative impact across various industries, leading to more intelligent, responsive, and human-centric technological solutions.

References

Bender, E. M., & Koller, A. (2020). Climbing towards NLU: On meaning, form, and understanding in the age of data. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (pp. 5185-5198). https://aclanthology.org/2020.acl-main.463/

Bengio, Y., Simard, P., & Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. IEEE transactions on neural networks, 5(2), 157-166. doi: 10.1109/72.279181

Bolukbasi, T., Chang, K. W., Zou, J. Y., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. In Advances in neural information processing systems (pp. 4349-4357). https://papers.nips.cc/paper_files/paper/2016/hash/a486cd07e4ac3d270571622f4f316ec5-Abstract.html

Chomsky, N. (1957). Syntactic structures. Mouton. https://tallinzen.net/media/readings/chomsky_syntactic_structures.pdf

Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., & Kuksa, P. (2011, March 1). Natural Language Processing (almost) from Scratch. NASA ADS. https://doi.org/10.48550/arXiv.1103.0398

Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., ... & Stoyanov, V. (2020). Unsupervised cross-lingual representation learning at scale. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (pp. 8440-8451). doi: 10.18653/v1/2020.acl-main.747

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. doi: 10.18653/v1/N19-1423

Goodfellow, I. J., Shlens, J., & Szegedy, C. (2014). Explaining and harnessing adversarial examples. https://doi.org/10.48550/arXiv.1412.6572

Hendrycks, D., & Gimpel, K. (2016). A baseline for detecting misclassified and out-of-distribution examples in neural networks. https://doi.org/10.48550/arXiv.1610.02136

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

Jain, S., & Wallace, B. C. (2019). Attention is not explanation. 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) (pp. 3543-3556). doi: 10.18653/v1/N19-1357

Jurafsky, D., & Martin, J. H. (2009). Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition. Pearson Prentice Hall. https://web.stanford.edu/~jurafsky/slp3/ed3book.pdf

Khurana, D., Koli, A., Khatter, K., & Singh, S. (2022). Natural language processing: State of the art, current trends and challenges. Multimedia Tools and Applications, 82(3), 3713–3744. https://doi.org/10.1007/s11042-022-13428-4

Lipton, Z. C. (2018). The mythos of model interpretability. Communications of the ACM, 61(10), 36-43. https://doi.org/10.1145/3233231

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. https://doi.org/10.48550/arXiv.1301.3781

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving language understanding by generative pre-training. https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf

Ruder, S., Peters, M. E., Swayamdipta, S., & Wolf, T. (2019). Transfer learning in natural language processing. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorials (pp. 15-18). doi: 10.18653/v1/N19-5004

Shannon, C. E. (1948). A mathematical theory of communication. Bell System Technical Journal, 27(3), 379-423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (pp. 3645-3650). doi: 10.18653/v1/P19-1355

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. In Advances in neural information processing systems (pp. 5998-6008). https://doi.org/10.48550/arXiv.1706.03762

Weaver, W. (1955). Translation. Machine translation of languages, 15-23. https://repositorio.ul.pt/bitstream/10451/10945/2/ulfl155512_tm_2.pdf

Wikipedia. (2019, June 28). Natural language processing. Wikipedia; Wikimedia Foundation. https://en.wikipedia.org/wiki/Natural_language_processing

Wilks, Y. (1973). Preference semantics. Formal semantics of natural language, 329-348. https://apps.dtic.mil/sti/citations/AD0764652

Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., & Drame, M. (2020). Transformers: State-of-the-Art Natural Language Processing. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. https://doi.org/10.18653/v1/2020.emnlp-demos.6

Young, T., Hazarika, D., Poria, S., & Cambria, E. (2018). Recent Trends in Deep Learning Based Natural Language Processing [Review Article]. IEEE Computational Intelligence Magazine, 13(3), 55–75. doi: 10.1109/MCI.2018.2840738