International Journal of Trend in Scientific Research and Development (IJTSRD)
Volume 10 Issue 5, Sep-Oct 2026 Available Online: www.ijtsrd.com e-ISSN: 2456 – 6470

Deep Learning for Natural Language Processing

Matthew N. O. Sadiku[1], Matthias Oteniya[2], Janet O. Sadiku[3]

[1,2]Roy G. Perry College of Engineering, Prairie View A&M University, Prairie View, TX, USA
[3]Juliana King University, Houston, TX, USA

ABSTRACT

Natural language processing (NLP) is the branch of artificial intelligence concerned with enabling computers to analyze, understand, and generate human language. Deep learning has transformed NLP by unifying representation learning and task learning. Earlier systems often required separate stages for tokenization, grammatical analysis, feature engineering, and prediction. Modern neural models can instead learn many of these functions jointly. Deep learning changed this approach by allowing multilayer neural networks to learn useful representations directly from data. This shift has produced major improvements in machine translation, information retrieval, question answering, text classification, summarization, and dialogue. The principal importance of deep learning in NLP therefore lies not in a single application, but in a general shift toward learned, reusable representations of language. This paper reviews deep learning and its successful applications to major NLP tasks.

KEYWORDS: *Deep learning, DL, deep machine learning, natural language processing, NLP, deep learning for natural language processing.*

How to cite this paper: Matthew N. O. Sadiku | Matthias Oteniya | Janet O. Sadiku "Deep Learning for Natural Language Processing" Published in International Journal of Trend in Scientific Research and Development (ijtsrd), ISSN: 2456-6470, Volume-10 | Issue-5, October 2026, pp.1719-1728, URL: www.ijtsrd.com/papers/ijtsrd142285.pdf

IJTSRD142285

Copyright © 2026 by author (s) and International Journal of Trend in Scientific Research and Development Journal. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (CC BY 4.0) (http://creativecommons.org/licenses/by/4.0)

INTRODUCTION

Human language is structured, ambiguous, context-dependent, and constantly changing. These qualities make language difficult for computers to process. Natural language processing (NLP) is the field concerned with computational methods for analyzing, understanding, and generating human language. Traditional NLP systems represented words through symbols and manually designed features. Deep learning is a branch of machine learning that uses neural networks with multiple layers to learn useful representations from data. When the two fields are combined, a system can learn linguistic patterns from examples instead of relying entirely on rules written by human experts. The striking success of deep learning in a wide variety of natural language processing (NLP) applications has served as a benchmark for the advances in one of the most important tasks in artificial intelligence. Thanks to the recent advances of deep learning, NLP applications have received an unprecedented boost in performance [1].

NLP using deep learning integrates DL models to better capture the meaning and language, improving performance in complex tasks. This has significantly advanced areas like machine translation, sentiment analysis, chatbots, and summarization. Language unfolds as a sequence, so early deep-learning approaches to NLP relied heavily on recurrent neural networks (RNNs). An RNN processes tokens in order and updates a hidden state that summarizes what it has encountered. More advanced recurrent units, especially long short-term memory networks, were designed to preserve information over longer spans and reduce optimization problems associated with ordinary recurrence.

WHAT IS DEEP LEARNING?

Deep learning is a branch of machine learning that uses artificial neural networks with multiple processing layers to discover patterns in data. Its importance lies not merely in the size of its models, but in their ability to learn useful representations directly from raw or minimally processed information. Earlier machine-learning systems often depended heavily on experts to decide which features of an image, sound, document, or transaction should be measured. Deep-learning models can instead learn

increasingly abstract features during training, enabling one general family of methods to work with images, speech, text, video, biological sequences, sensor readings, and other complex data. This capacity has made deep learning a central technology in contemporary artificial intelligence [2]. Figure 1 shows how deep learning is a subset of machine learning and how machine learning is a subset of artificial intelligence [3].

Deep learning may be regarded as a feature learning technique. Different applications employ different deep learning models. Deep learning has matured from being a special purpose machine learning technique to a general purpose machine learning tool. It has become a big wave of technology trend for big data and artificial intelligence. It has enjoyed success in various applications such as automatic speech recognition, image recognition, computer vision, object detection, bioinformatics, and other areas such as drug discovery and information retrieval.

Fundamentally, deep learning refers to a class of machine learning algorithms in which a hierarchy of layers is used to transform input data into a progressively more abstract and composite representation. The word "deep" in "deep learning" refers to the number of layers through which the data is transformed. Neural network consists of layers of interconnected nodes or neurons that collaborate to process input data. In a fully connected deep neural network data flows through multiple layers where each neuron performs nonlinear transformations, allowing the model to learn intricate representations of the data. In a deep neural network the input layer receives data which passes through hidden layers that transform the data using nonlinear functions. The final output layer generates the model's prediction [4].

Deep learning models are complicated artificial neural networks designed to mimic the human brain. Figure 2 shows a deep learning neural network [5], while Figure 3 shows the key components of deep learning [6]. Deep learning is behind many technologies we use every day like voice assistants and medical tools. It is especially valuable when working with large quantities of unstructured data. Conventional databases organize information into clearly defined rows and columns, but much of the world's information exists as photographs, recordings, documents, and videos. Deep neural networks are well suited to extracting patterns from these formats. Its greatest promise, however, emerges when technical capability is combined with responsible design.

Deep learning is not just a technological advancement. It is a reflection of human intelligence, scaled, accelerated, and embedded into systems around us. It is most commonly used where there is massive data and a need for intelligent interpretation. Hence, its strongest areas of applications include language processing, image recognition, predictive analytics, and automation [7].

FUNDAMENTALS OF NLP

Natural language processing is a subfield of artificial intelligence that empowers computers to understand, interpret, and generate human language. It is a technique where machine can become more human and thereby making human to communicate with the machine easily. NLP seeks to make software intelligent enough to process a natural language as humans. For example, imagine a machine that takes instructions by voice.

NLP analysis generally consists of the following three levels [8]:

- *Syntax,* the study of sentence structure. Syntax deals with the formation of a sentence from individual words. Syntax alone suggests the proper interpretation of "Jimmy loves Lucy."
- *Semantics,* the study of context-independent meaning. This derives the meaning of a sentence based on the meanings of the words/phrases. For example, semantics determines whether the word "bank" refers to a river bank or to a financial institution.
- *Pragmatics,* the study of context-dependent meaning. Pragmatics deals with how meaning changes in the presence of a specific context and how the contexts affect the meaning of the sentences. This level is concerned with the purposeful use of language in situations.

As a foundational pillar of modern artificial intelligence, NLP encompasses a wide array of tasks, including speech recognition, text classification, natural language understanding (NLU), and natural language generation (NLG). NLP encompasses a wide range of tasks, such as information retrieval (IR), named entity recognition (NER), relation extraction, text classification, topic modeling, semantic textual similarity, machine translation, and question answering (QA). Figure 4 shows how NLP transforms raw acoustic data into meaningful interactions [9], while Figure 5 shows different components of NLP [10].

Recently, large language models (LLMs) have shown their ability in learning universal language representations, text understanding and generation. LLMs refer to a model with a large number of

parameters, vast training data, and substantial compute, enabling it to capture complex language patterns. In LLM-based NLP, pre-processing is followed by prompt engineering, which guides LLMs to produce outputs that align with extraction requirements during inference without altering the model's parameters. Models like GPT are pushing the boundaries of language understanding, enabling nuanced and context-aware applications. The GPT (Generative Pretrained Transformer) is a large-scale language model developed by OpenAI that consists of multiple layers of transformer blocks, each with a self-attention mechanism and a forward neural network [11]. GPT-based systems can summarize complex reports or generate creative content like essays, making them versatile in both academic and professional environments. ChatGPT uses NLP techniques to understand prompts. When you enter a prompt, the chatbot comprehends it and provides relevant replies. Figure 6 shows how NLP works [12].

DEEP LEARNING IN NATURAL LANGUAGE PROCESSING

Deep learning has expanded both the accuracy and the scope of language technology. In text classification, neural models identify sentiment, intent, topic, toxicity, or document type. In information extraction, they detect entities, relations, events, and facts in unstructured documents. In semantic search, embeddings match a query with relevant content even when the exact words differ. In machine translation and summarization, sequence models generate new text conditioned on an input. In question answering and dialogue, pretrained models combine retrieval, contextual interpretation, and generation to produce interactive responses. A recent review finds transformer-based methods across these areas and identifies text classification as a particularly common application in the literature it examined [1]. Figure 7 shows deep learning based NLP [13], while Figure 8 shows the Venn diagram of artificial intelligence, machine learning, deep learning, and NLP [14].

APPLICATIONS OF DEEP LEARNING IN NATURAL LANGUAGE PROCESSING

Reliable applications require surrounding systems such as retrieval, citations, constrained decoding, validation rules, human review, and monitoring. The major applications of deep learning in NLP can be organized according to whether a system primarily classifies, extracts, retrieves, or generates language. In practice, modern systems often combine several of these functions. Its applications now include text classification, sentiment analysis, information extraction, translation, search, question answering, summarization, dialogue, and multilingual transfer. Common applications of deep learning in NLP therefore include the following [1]:

- *Text Classification:* Text classification is one of the most widely deployed uses of deep learning in NLP. A model maps a passage to one or more categories, such as spam or not spam, a news topic, a customer-service intent, or a moderation label. Deep models improve this process because they learn distributed representations: semantically related expressions can receive similar internal representations even when they contain different words. Contextual models go further by representing the same word differently according to its surrounding sentence. This capability helps classifiers handle polysemy, paraphrase, and indirect phrasing more effectively than methods based only on isolated keywords.
- *Sentiment Analysis:* Sentiment analysis is a specialized form of classification that estimates whether a text expresses a positive, negative, or neutral attitude. More detailed systems may identify emotions, targets, or aspect-specific opinions. For example, a product review might praise a camera's image quality while criticizing its battery life. A deep model can associate each sentiment with the relevant product attribute rather than assigning one label to the entire review. Businesses use such systems to analyze surveys, reviews, and support conversations, while researchers apply them to social and political discourse. Nevertheless, sarcasm, cultural variation, domain-specific language, and mixed sentiment remain difficult because the intended meaning may depend on knowledge that is not explicit in the text.
- *Text Summarization:* Deep learning supports both extractive summarization, which selects important sentences, and abstractive summarization, which generates a condensed restatement. The latter resembles human summarization because it can combine information, remove repetition, and paraphrase. Applications include news digests, meeting minutes, research reviews, case summaries, and condensed versions of lengthy reports. A summarizer can reduce the time needed to navigate a document collection and can tailor an explanation to a desired length or audience.
- *Information Extraction:* Information extraction converts unstructured language into structured data. Its component tasks include named-entity recognition, relation extraction, event detection, entity linking, and semantic role labeling. A system might identify the names of people,

organizations, medicines, dates, or locations and then determine how those entities are related. Deep contextual representations are particularly useful because an expression such as "Apple" must be interpreted differently when it refers to a company rather than a fruit.

- *Machine Translation:* Machine translation was one of the earliest applications to demonstrate the strength of end-to-end deep learning. Sequence-to-sequence networks learned to map a sentence in one language directly to a sentence in another. Attention improved this process by allowing a model to align target words with relevant parts of the source sentence. Transformers then made it easier to model long-range dependencies and train on large parallel corpora.
- *Multilingual Communication:* Multilingual models can transfer knowledge among languages. XLM-R, for example, was pretrained on one hundred languages and demonstrated gains across cross-lingual inference, question answering, and named-entity recognition, including improvements for some lower-resource languages. Even so, quality remains uneven when languages have limited digital data, distinctive writing systems, or complex morphology. Fluency can also conceal a mistranslation, so professional review remains essential for medical, legal, diplomatic, and other high-stakes material.
- *Question Answering Systems*: Question-answering systems may be extractive, selecting an answer span from a source, or generative, composing a new answer. BERT's results helped establish pretrained contextual encoders as effective foundations for machine reading comprehension. Generative systems provide greater flexibility but create a risk of producing unsupported statements. Retrieval-augmented generation addresses part of this problem by combining a language model's internal parameters with an external document index. Retrieval also permits organizations to update the source collection without retraining the entire model. It does not, however, guarantee correctness: the system can retrieve an irrelevant passage, misinterpret a relevant one, or make claims that the evidence does not support.
- *Multimodal Learning:* Multimodal learning will consequently become central to NLP rather than adjacent to it. Future systems will jointly process language, speech, images, video, interfaces, and sensor data. This integration will improve applications such as accessibility, tutoring, robotics, meeting analysis, and cross-media search. It may also make learning more data-efficient by allowing different signals to constrain one another.

BENEFITS

The main benefits of deep learning in NLP are improved accuracy, reduced dependence on manual feature engineering, effective use of context, transfer across tasks, support for generation and multiple data types, and a more unified approach to building language technologies. Deep learning has given NLP unprecedented fluency, adaptability, and practical reach. Other benefits of deep learning in NLP include the following [1]:

- *Representation Learning*: The first major advantage of deep learning is representation learning. A system can learn task-relevant features from examples rather than depending entirely on manually specified rules. A traditional NLP system may require experts to specify which words, grammatical categories, dictionary entries, or sentence patterns should influence a prediction. A deep neural network instead learns internal numerical representations from examples. Its layers can encode information at several levels of abstraction, from character and word patterns to sentence meaning and broader context. This general ability to learn layered representations is a defining feature of deep learning. Representation learning also reduces the need for extensive manual feature engineering.
- *End-to-End Learning*: Deep learning enables many NLP applications to be trained end to end. In an end-to-end system, a model learns a direct mapping from an input, such as a sentence, to a desired output, such as a translation or label. An end-to-end neural model can optimize its internal components together for the final task. In machine translation, for example, the model learns from pairs of source and target sentences rather than relying on a large collection of separately engineered linguistic rules. This simplification does not eliminate all preprocessing or human oversight, but it aligns training more closely with the result that users actually need. It can also reduce maintenance costs because fewer task-specific modules must be coordinated.
- *Transfer Learning*: A central practical benefit is transfer learning. Pretrained models can be adapted to new tasks with less labeled data than would be required to train a comparable model from scratch. Multilingual pretraining extends this benefit across languages. A model can first be pretrained on a large collection of unlabeled text and later adapted to a narrower task using a smaller labeled dataset. This process separates

general language learning from task-specific learning. The model acquires broad knowledge of vocabulary, syntax, and recurring semantic relationships during pretraining, while fine-tuning teaches it how to apply that knowledge to a particular objective. Transfer learning makes high-quality NLP more accessible in domains where labeled data are expensive.

- *Computation:* As technology continues to evolve, computation will remain a cornerstone of applied mathematics, providing essential tools for solving problems that are beyond the reach of traditional methods. Deep learning changes the interface to computation. A sufficiently capable language model can convert free-form requests into structured outputs, explanations, database queries, or software instructions. Yet fluency must not be confused with verified knowledge. A model generates outputs by learning statistical patterns; it does not guarantee that a statement is true, current, or supported by evidence.
- *Evaluation*: Responsible NLP begins by defining success more carefully. A benchmark score captures performance on a particular dataset, not general language understanding or social acceptability. Evaluation should therefore examine accuracy, robustness, calibration, subgroup performance, privacy, latency, resource use, and the cost of errors. High-stakes applications in medicine, law, finance, education, or public administration require domain experts and meaningful human oversight. NLP should not be evaluated only by whether a machine can produce human-like text. It should be judged by whether a system helps people accomplish legitimate goals accurately, safely, and fairly. Deep learning will remain central, but its success should be measured not only by increasingly fluent machines, but also by the reliability and human value of the systems built with them.

CHALLENGES

The same flexibility that makes deep learning powerful also creates risk. The central challenge is not simply to make deep NLP models larger or more accurate. It is to make them meaning-sensitive, robust, interpretable, inclusive, efficient, secure, and accountable in the environments where language is actually used. The challenges of deep learning for NLP are interconnected. Other challenges of deep learning in NLP include the following [1]:

- *Cost*: Computational cost has both economic and environmental consequences. Training a large model requires specialized hardware, energy, engineering expertise, and extensive data infrastructure. Inference can also become expensive when a model serves millions of requests. Efficient architectures, compression, quantization, distillation, sparse computation, and smaller domain-specific models are therefore not merely technical conveniences. They are important tools for accessibility and sustainability.
- *Data Quality*: Deep learning is data-intensive, so the quality of an NLP system is constrained by the quality and composition of its training corpus. Web-scale datasets contain valuable linguistic diversity, but they also contain misinformation, duplicated material, stereotypes, harassment, private information, and text produced under unequal social conditions. A model trained to reproduce the statistical structure of such data may also reproduce its prejudices. The problem is not limited to offensive vocabulary. Bias can appear in associations between occupations and gender, differences in performance across dialects, assumptions about names or nationalities, and uneven rates of harmful classification.
- *Data Practices*: Data practices are equally important. Developers should document where data came from, what populations they represent, how they were filtered, and what uses are appropriate. Model documentation should state capabilities, known limitations, test conditions, and failure modes. After deployment, systems need logging, incident response, user feedback channels, and periodic reassessment because language, users, and operating environments change.
- *Bias:* Bias is another fundamental concern. Deep-learning systems learn from data rather than from an explicit, complete model of truth. Consequently, their behavior reflects the quality, coverage, and biases of their training material. Training corpora contain stereotypes, exclusions, abusive language, and historical inequalities. A model may reproduce these patterns or perform less effectively for communities that are poorly represented in the data. Dataset scale does not remove this problem. In some cases, scale makes the provenance and composition of training material harder to audit.
- *Privacy:* Privacy and intellectual-property questions arise when models are trained on massive collections whose contents, permissions, and personal information may be unclear. Security risks include prompt injection, data leakage, malicious automation, and adversarial

manipulation. Concentration of model development within a small number of well-resourced institutions may also limit scientific reproducibility and distribute benefits unevenly.

- *Security*: Security risks extend beyond memorization. Attackers can manipulate inputs, attempt to override system instructions, poison public data sources, or use generated language for fraud and influence operations. When an NLP model is connected to retrieval systems, databases, or software tools, a linguistic attack may become an operational attack. Security must therefore cover the entire application: data collection, model training, interfaces, permissions, external tools, monitoring, and incident response.
- *Hallucination:* The strongest NLP models remain vulnerable to hallucination, the generation of plausible but unsupported content. This problem follows partly from the training objective: next-token prediction rewards likely continuation, not truth. Models may also fail under distribution shift, respond inconsistently to small changes in wording, or exploit superficial patterns in benchmarks. High average accuracy can conceal severe errors for rare cases or marginalized groups.
- *Linguistic Inequality:* Linguistic inequality adds another dimension. Although thousands of languages are spoken worldwide, modern NLP resources and research remain concentrated in a small group of high-resource languages. A multilingual architecture cannot compensate fully for the absence of digitized text, annotated examples, suitable evaluation sets, or local expertise. Even within a well-resourced language, standardized written forms may dominate over regional dialects, code-switching, informal speech, and specialized communities.
- *Robustness*: Another technical challenge is robustness, meaning that a system should preserve correct behavior when an input changes in a way that does not change its meaning. Deep NLP models can perform well on examples drawn from the same distribution as their training data while failing under new topics, writing styles, demographic groups, or carefully constructed distractions. This occurs because optimization rewards any predictive pattern, including superficial shortcuts that happen to work in a benchmark.
- *Interpretability*: Interpretability will also become more practical and role-specific. Researchers may need methods for identifying internal circuits and causal mechanisms. Auditors may need reproducible evidence about data, evaluations, and known limitations. End users may need citations, confidence indicators, and concise explanations of why a system took an action. No single explanation method will satisfy every audience. Meaningful transparency will depend on giving each stakeholder the information required to challenge or verify a result.

FUTURE OF DEEP LEARNING IN NATURAL LANGUAGE PROCESSING

The future of deep learning in NLP is likely to emphasize models that are not only larger, but also more efficient, grounded, controllable, and inclusive. The next phase of deep learning for NLP will likely focus less exclusively on scale and more on efficiency, grounding, controllability, and evaluation. Smaller models trained on carefully selected data can outperform larger ones in specialized contexts. Retrieval-augmented systems can separate stored knowledge from linguistic generation. Tool-using models can consult databases, execute calculations, and verify intermediate results. Multimodal systems can connect language with images, audio, video, and action. Future NLP systems will likely combine learned intuition with explicit procedures. Future NLP development will incorporate safety throughout the system lifecycle.

Another important direction is the design of hybrid systems. Neural models are strong at recognizing patterns and generating flexible language, while databases, symbolic rules, calculators, and workflow engines are stronger at preserving exact facts and enforcing constraints. Combining these components can produce systems that are both capable and accountable. The future of deep learning for natural language processing will be neither a simple continuation of scaling nor a rejection of large models. It will be a reorganization of the field around efficient and accountable language systems [1].

CONCLUSION

Deep learning has transformed natural language processing (NLP), the field concerned with enabling computers to analyze, generate, and interact through human language. It has reshaped NLP by replacing many hand-engineered pipelines with models that learn rich representations from data. It has moved NLP from manually engineered pipelines toward general systems that learn representations and tasks from data. The enduring contribution of deep learning is its ability to learn flexible linguistic representations at scale. This progression has yielded extraordinary practical capabilities, but it has not eliminated the

fundamental difficulties of language or the responsibilities attached to automated communication.

Research must improve linguistic and cultural coverage. Many languages lack large digital corpora, standardized benchmarks, and commercial incentives. Progress will require community participation, multilingual data governance, and methods that learn effectively from limited resources. Deep learning's lasting contribution to NLP will depend not only on how well machines can process language, but also on how responsibly that ability is applied. More information about deep learning in natural language processing can be found in the books in [15-24].

REFERENCES

[1] https://manus.im

[2] M. N.O. Sadiku, M. Tembely, and S.M. Musa, "Deep learning," *International Research Journal of Advanced Engineering and Science,* vol. 2, no. 1, 2017, pp. 77,78.

[3] "Deep learning," *Wikipedia,* https://en.wikipedia.org/wiki/Deep_learning

[4] "Introduction to deep learning," May 2026, https://www.geeksforgeeks.org/deep-learning/introduction-deep-learning/

[5] N. Somani, "What are deep learning models?" August 2026, https://www.igmguru.com/blog/what-are-deep-learning-models

[6] "25 Deep learning applications driving AI innovation," https://www.intellspot.com/deep-learning-applications/

[7] A. Kumar, "Applications of deep learning: Real-world examples and uses," May 2026, https://www.theiotacademy.co/blog/applications-of-deep-learning/

[8] J. Hirschberg, B. W. Ballard, and D. Hindle, "Natural language processing," *AT&T Technical Journal*, Jan./Feb. 1988, vol. 67, no. 1, 1988.

[9] A. Jain, "What is the role of NLP in voice assistants?" August 2024, https://www.analyticsinsight.net/nlp/what-is-the-role-of-nlp-in-voice-assistants#google_vignette

[10] "How Google uses NLP to improve SERPs, featured Snippets & UX," https://digitalguider.com/blog/what-is-google-nlp/

[11] X. Jiang et al., "Applications of natural language processing and large language models in materials discovery," *NPJ Computational Materials*, vol. 11, no.79, 2025.

[12] "What is NLP? How it works, benefits, challenges, examples," June 2025, https://www.shaip.com/blog/what-is-nlp-how-it-works-benefits-challenges-examples/

[13] M. Han et al., "Perspectives of machine learning and natural language processing on characterizing positive energy districts," *Buildings*, vol. 12, no.2, January 2024.

[14] M. Surdeanu and M. A. Valenzuela-Escárcega, *Deep Learning for Natural Language Processing: A Gentle Introduction.* Cambridge University Press, 2024.

[15] P. Goyal, S. Pandey, and K. Jain, *Deep Learning for Natural Language Processing: Creating Neural Networks with Python.* Apress, 2018.

[16] M. Ekman, *Learning Deep Learning: Theory and Practice of Neural Networks, Computer Vision, Natural Language Processing, and Transformers Using TensorFlow*. Addison-Wesley Professional, 2021.

[17] K. R. Bokka et al., *Deep Learning for Natural Language Processing: Solve Your Natural Language Processing Problems With Smart Deep Neural Networks*. Packt Publishing, 2019.

[18] A. S. Patel, *Deep Learning for Natural Language Processing: Deep Learning With NLP*. Notion Press, 2024.

[19] J. Brownlee, *Deep Learning for Natural Language Processing: Develop Deep Learning Models for your Natural Language Problems.* Machine Learning Mastery, 2017.

[20] L. Deng and Y. Liu (eds.), *Deep Learning in Natural Language Processing.* Springer, 2018.

[21] K. Huang et al. (eds), *Deep Learning: Fundamentals, Theory and Applications: Cognitive Computation Trends.* Springer, 2019.

[22] D. Rivera, *Deep Learning for Natural Language Processing.* Chapman Press, 2025.

[23] S. Raaijmakers, *Deep Learning for Natural Language Processing.* Manning, 2022.

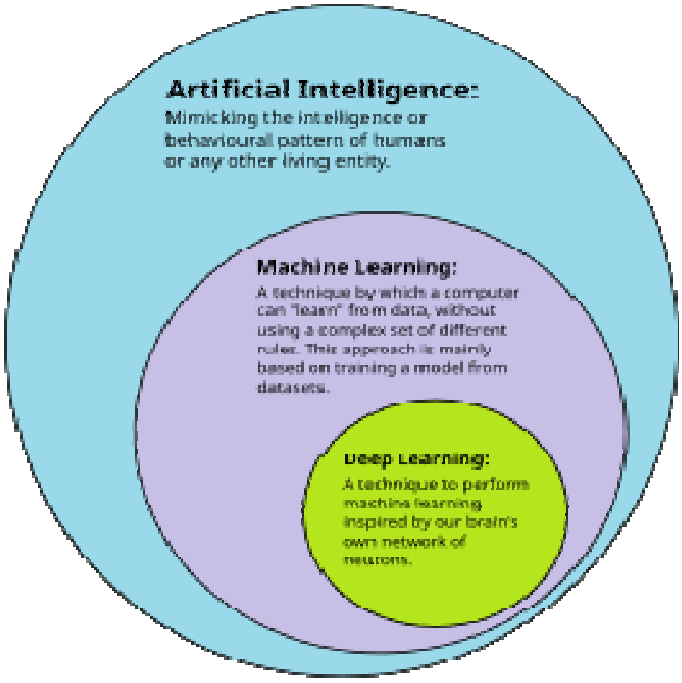


Figure 1 Deep learning is a subset of machine learning and machine learning is a subset of artificial intelligence [2].

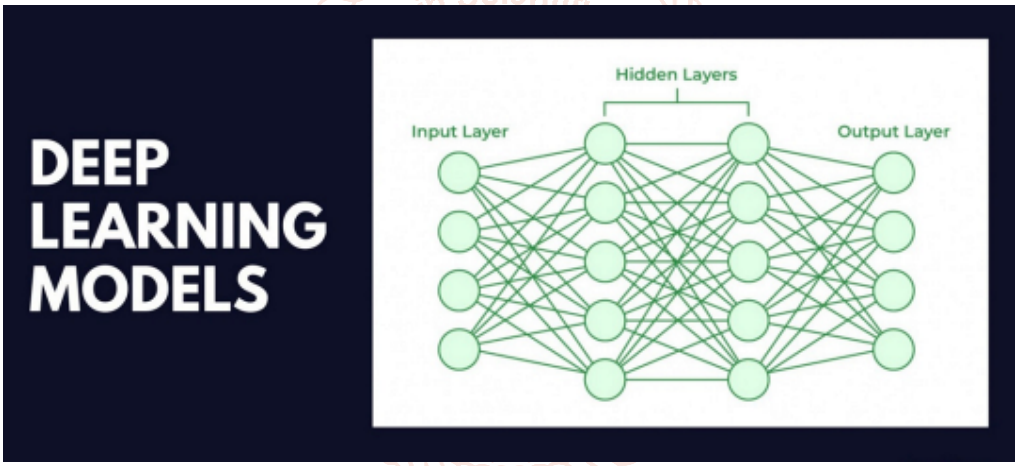


Figure 2 A deep learning neural network [5].

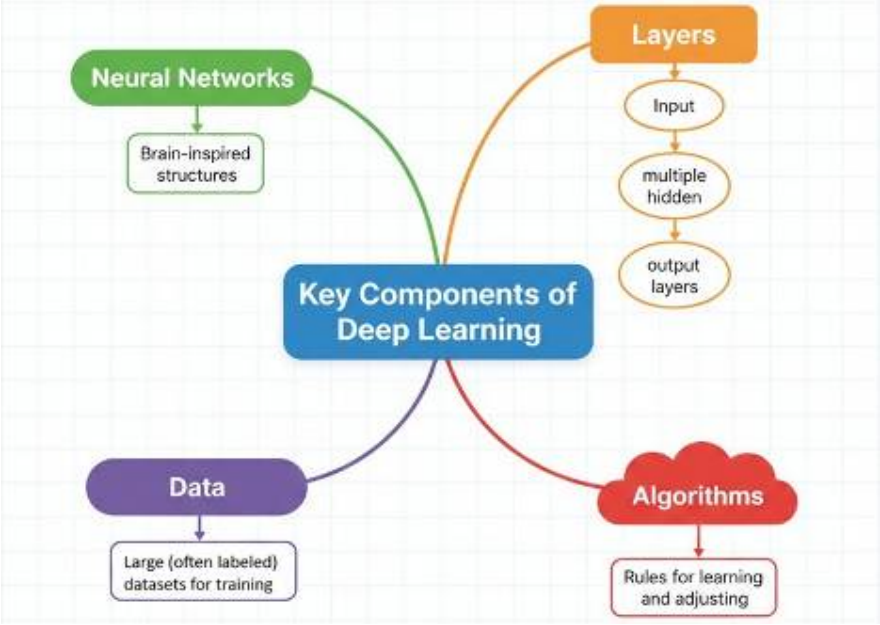


Figure 3 The key components of deep learning [6].

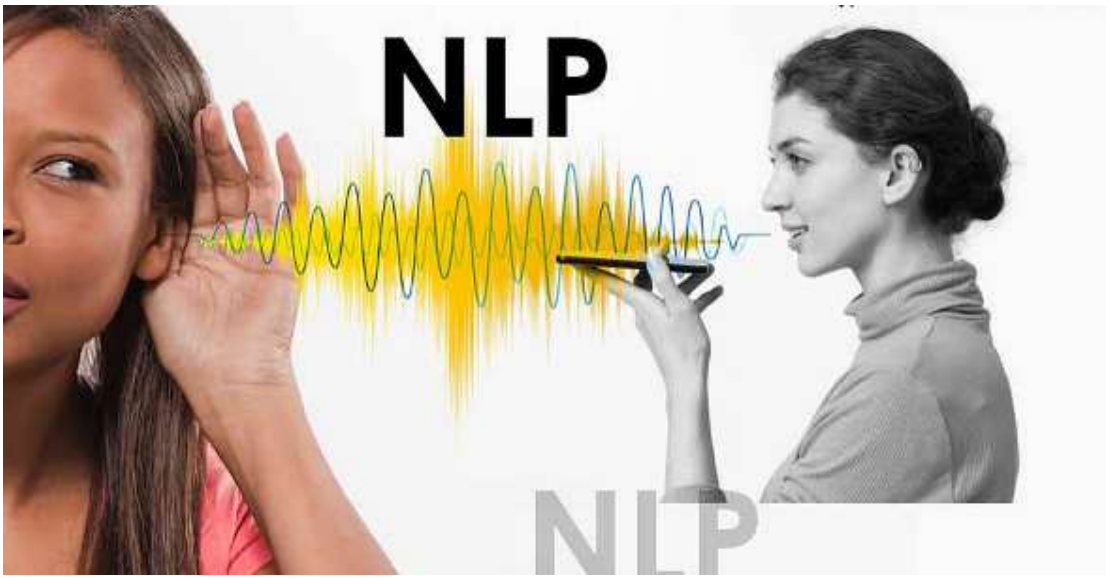


Figure 4 A representation of NLP [9].

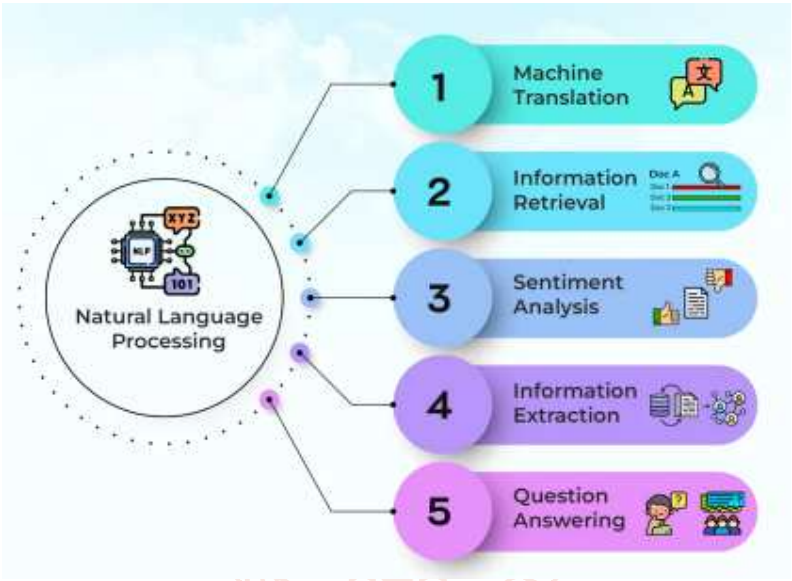


Figure 5 Different components of NLP [10].

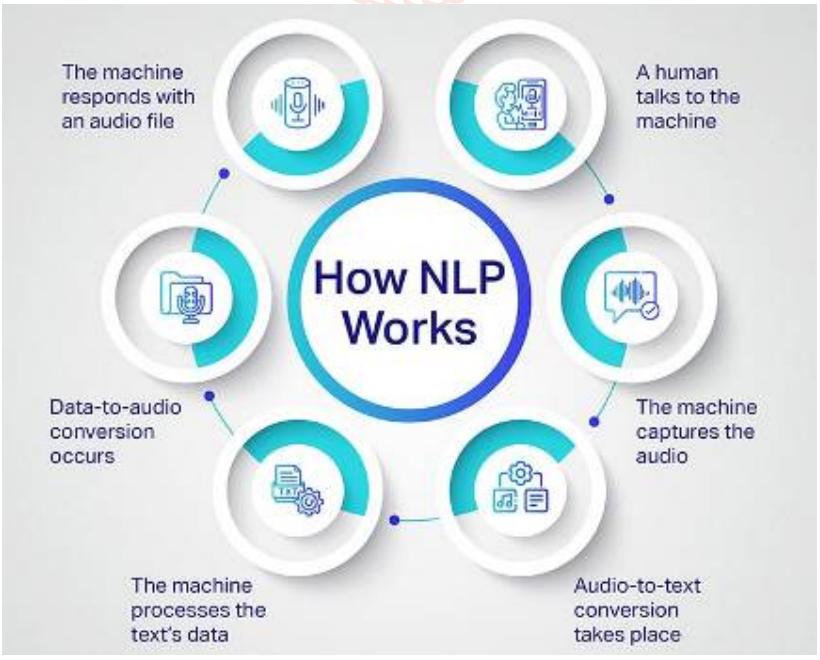


Figure 6 How NLP works [12].

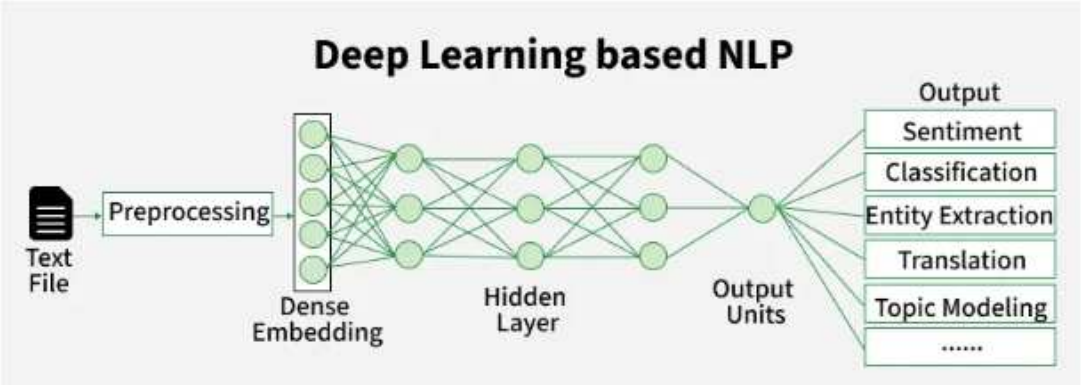


Figure 7 Deep learning based NLP [13].

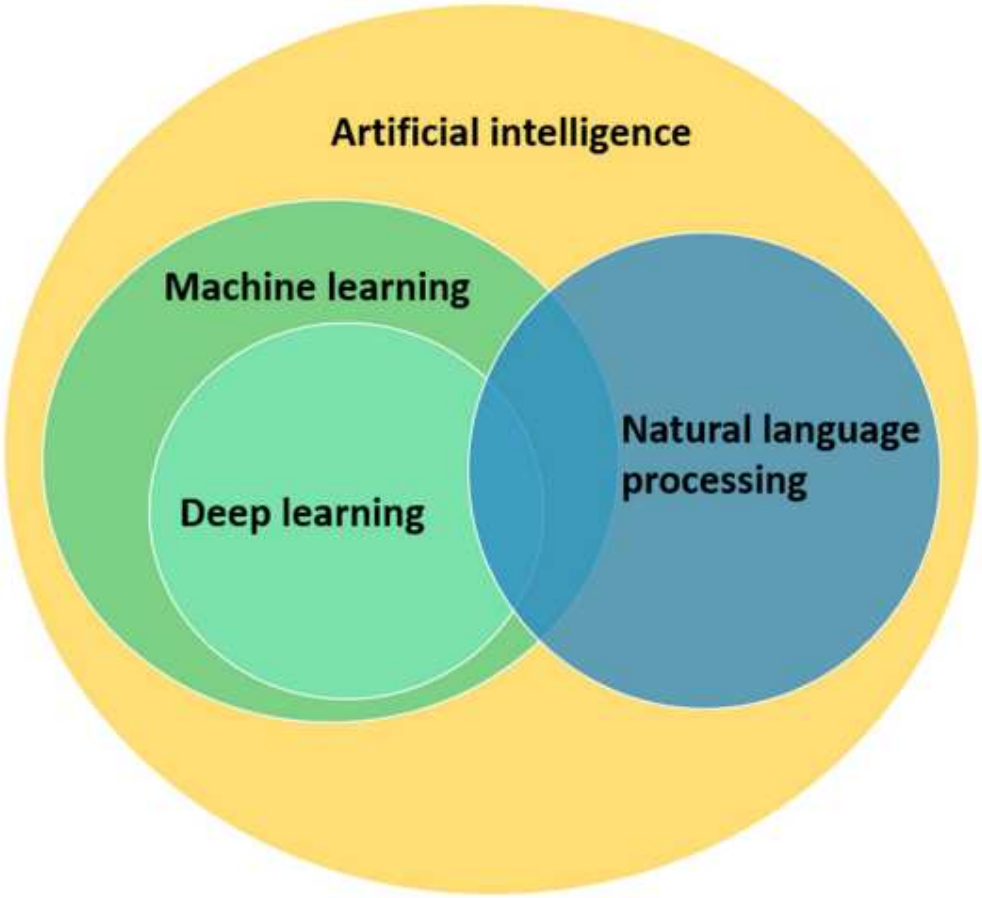


Figure 8 Venn diagram of artificial intelligence, machine learning, deep learning, and NLP [14].