


NLP VS NLU: WHAT'S THE DIFFERENCE?




A natural language is one that has [evolved over time](#) via use and repetition. It does not involve deliberate planning and strategy. Latin, English, Spanish, and many other spoken languages are all languages that evolved naturally over time.

Natural languages are different from formal or constructed languages, which have a different origin and development path. For example, programming languages including C, [Java](#), [Python](#), and many more were created for a specific reason.

For a machine to be autonomous, a key tenet is to be able to communicate via one of the natural languages known to humans. In the wide world of [Artificial Intelligence](#), one field deals with enabling machines to interact using these languages: Natural Language Processing (NLP).

NLP is an umbrella term which encompasses any and everything related to making machines able to process natural language—be it receiving the input, understanding the input, or generating a response.

In this context, another term which is often used as a synonym is Natural Language Understanding (NLU). Actually, though, NLP and NLU focus on different areas. In this article, we'll look at them to understand the nuances.



What is natural language processing?

From the computer's point of view, any natural language is a [free form text](#). That means there are no set keywords at set positions when providing an input.

Beyond the unstructured nature, there can also be multiple ways to express something using a natural language. For example, consider these three sentences:

- How is the weather today?
- Is it going to rain today?
- Do I need to take my umbrella today?

All these sentences have the same underlying question, which is to enquire about today's weather forecast.

As humans, we can identify such underlying similarities almost effortlessly and respond accordingly. But this is a problem for machines—any algorithm will need the input to be in a set format, and these three sentences vary in their structure and format. And if we decide to code rules for each and every combination of words in any natural language to help a machine understand, then things will get very complicated very quickly.

This is where NLP enters the picture.

NLP is a subset of AI tasked with enabling machines to interact using natural languages. The domain of NLP also ensures that machines can:

- Process large amounts of natural language data
- Derive insights and information

But before any of this natural language processing can happen, the text needs to be standardized.

In [machine learning \(ML\)](#) jargon, the series of steps taken are called [data pre-processing](#). The idea is to break down the natural language text into smaller and more manageable chunks. These can then be analyzed by ML algorithms to find relations, dependencies, and context among various chunks.

Some examples of pre-processing steps are:

- Parsing
- Stop-word removal
- Part-of-speech (POS) tagging
- Tokenization
- Many more

Thus, we can sum up: The aim of NLP is to process the free form natural language text so that it gets transformed into a standardized structure.

What is natural language understanding (NLU)?

Considered [a subtopic](#) of NLP, the main focus of natural language understanding is to make machines:

- Interpret the natural language
- Derive meaning
- Identify context
- Draw insights

For example, in NLU, various ML algorithms are used to [identify the sentiment](#), perform [Name Entity Recognition](#) (NER), process semantics, etc. NLU algorithms often operate on text that has already been standardized by text pre-processing steps.

Going back to our weather enquiry example, it is NLU which enables the machine to understand that those three different questions have the same underlying weather forecast query. After all, different sentences can mean the same thing, and, vice versa, the same words can mean different things depending on how they are used.

Let's take another example:

- The banks will be closed for Thanksgiving.
- The river will overflow the banks during floods.

A task called [word sense disambiguation](#), which sits under the NLU umbrella, makes sure that the machine is able to understand the two different senses that the word "bank" is used.

So, how do NLP & NLU differ?

In natural language, what is expressed (either via speech or text) is not always what is meant. Let's take an example sentence:

- Please crack the windows, the car is getting hot.

NLP focuses on processing the text in a literal sense, like what was said. Conversely, NLU focuses on extracting the context and intent, or in other words, what was meant.

NLP will take the request to crack the windows in the literal sense, but it will be NLU which will help draw the inference that the user may be intending to roll down the windows.



NLP alone could result in literal damage

NLP can process text from grammar, structure, typo, and point of view—but it will be NLU that will help the machine infer the intent behind the language text. So, even though there are many overlaps between NLP and NLU, this differentiation sets them distinctly apart.

Do we need both?

In one word, yes.

On our quest to make more robust autonomous machines, it is imperative that we are able to not only process the input in the form of natural language, but also understand the meaning and context—that's the value of NLU. This enables machines to produce more accurate and appropriate responses during interactions.

Let's take the example of ubiquitous [chatbots](#).

Gone are the days when chatbots could only produce programmed and rule-based interactions with their users. Back then, the moment a user strayed from the set format, the chatbot either made the user start over or made the user wait while they find a human to take over the conversation.

Combining NLU and NLP, today's chatbots are more robust. Using NLU methods, chatbots are able to:

- Be aware of the conversation's context
- Extract the conversation's meaning based on that context
- Guide users on the topic of conversation

Ecommerce websites rely heavily on sentiment analysis of the reviews and feedback from the users—was a review positive, negative, or neutral? Here, they need to know what was said and they also need to understand what was meant.



User reviews aren't always easy to understand

When are machines intelligent?

In the world of AI, for a machine to be considered intelligent, it must pass the [Turing Test](#). A test developed by Alan Turing in the 1950s, which pits humans against the machine.


To pass the test, a human evaluator will interact with a machine and another human at the same time, each in a different room. If the evaluator is not able to reliably tell the difference between the response generated by the machine and the other human, then the machine passes the test and is considered to be exhibiting "intelligent" behavior.

This is a crude gauge of intelligence, albeit an effective one. The first successful attempt came out in 1966 in the form of the famous [ELIZA program](#) which was capable of carrying on a limited form of conversation with a user.

Since then, with the help of progress made in the field of AI and specifically in NLP and NLU, we have come very far in this quest. After all, chatbots are everywhere.

NLP & NLU use cases

According to various industry estimates only about 20% of data collected is structured data. The remaining 80% is [unstructured data](#)—the majority of which is unstructured text data that's unusable for traditional methods. Just think of all the online text you consume daily, social media, news, research, product websites, and more.



NLP and NLU techniques together are ensuring that this huge pile of unstructured data can be processed to draw insights from data in a way that the human eye wouldn't immediately see. Machines can find patterns in numbers and statistics, pick up on subtleties like sarcasm which aren't inherently readable from text, or understand the true purpose of a body of text or a speech.

NLP and NLU are helping to ensure that we are able to process and use this [enormous amount of data](#) being generated. Some common use cases using NLP techniques are:

- Speech recognition (e.g., Siri, Alexa)
- Machine translate (e.g. Google Translate)
- Chatbots
- [Sentiment analysis](#)

The future for language

Thanks to recent advancements, another sub-field of NLP is Natural Language Generation. [NLG](#) has gained a lot of attention.

In addition to processing natural language similarly to a human, NLG-trained machines are now able to generate new natural language text—as if written by another human. All this has sparked a lot of interest both from commercial adoption and academics, making NLP one of the most active research topics in AI today.

(See how [AI language models](#) & [GPT-3](#) actually work.)

Related reading

- [BMC Machine Learning & Big Data Blog](#)
- [Top Machine Learning Algorithms & How To Get Started](#)
- [Interpretability vs Explainability: The Black Box of Machine Learning](#)
- [Tuning Machine Language Models for Accuracy](#)
- [An Alternate Approach To Training Chatbots](#)
- [State of AI in 2021](#)

Learn how to make your own chatbot—see how BMC did it: