


# WHAT IS A LANGUAGE MODEL?



A language model is a statistical tool to predict words. Where weather models predict the 7-day forecast, language models try to find patterns in the [human language](#). They are used to predict the spoken word in an audio recording, the next word in a sentence, and which email is spam. Grammarly recently released a new feature to detect how confident a user's email might sound.

Images have dominated the field of A.I. for the past four years; language models are the next frontier for [artificial intelligence](#). Images were more accessible for modelling tasks because they are easier to label, and they have computer-interpretable information already encoded in them, pixels.

Language models are more challenging for modelers to use because language is complex. Language has layers of meaning. Language changes, sentence to sentence, author to author. Sentences are also harder to label than images. Where images can be labeled widely on general tasks, language models must be developed for very specific tasks. The meaning of language depends largely on context, and small variations in the words of sentences can vastly change the sentence's meaning. Only within the past year or two have labelling tools become available to the public. SpaCy's [Prodigy tool](#) is excellent for labeling language models.

## How language models work

Computers can only read numbers. So, in order for a language model to be created, all words must be converted to a sequence of numbers for the computer to read. For modelers, these are called encodings.

Encodings can range in complexity. At their basic level, every word can get its own number. This is called label-encoding. Label encoding assigns one number for each unique word in a sentence. If we had the sentence, "The girl sat by the tree", we could represent the sentence by the numbers, . This is the most basic encoding one could give to a set of words.

## 2019: The transformer language model

In 2014, Word2Vec was discovered and became a much better encoding method. Since then, encodings have greatly evolved from label-encoding methods, to one-hot encoding methods, to word-embeddings, and to 2019's best encoding method, the transformer.

Language modelling got a big popularity boost in 2019 with the development of the transformer. Last year, there were four major releases of very sophisticated transformers provided by Google, Open AI, and Salesforce. These transformer models given by Open AI and Google cost those companies \$60,000+ of computer resources to build. Another large contributor was the research team at [Hugging Face](#) who packaged these often hard-to-use language models in a user-friendly framework for machine learning engineers.

## Potential challenges and benefits of language models

In modelling, if we tried to predict the grade on an exam based on how many hours we studied and whether or not we ate breakfast, we would say that model has two parameters, time spent studying and whether we ate breakfast. The transformer models these companies have released have been trained on 1.5 billion parameters, and are often so big they can't work on most computers or mobile devices.

When openAI released its GPT-2 model, it came with some caveats. The company said it was a new technology and they wanted to see how their release of the language model would be used by developers first:

*"Due to our concerns about malicious applications of the technology, we are not releasing the trained model. As an experiment in responsible disclosure, we are instead releasing a much [smaller model](#) for researchers to experiment with, as well as a [technical paper](#)." - openAI*

Language models pose a threat to disinformation. Where deep fakes are putting out fake videos of people, language models can be used by bad parties to create junk news, or increase cat-phishing scams through fake text messages, emails, and social media status updates. Articles generated by a machine continue to be released that look like other people actually wrote them.

Here is an example of the auto-generated texts:

SYSTEM PROMPT (HUMAN-WRITTEN)	<i>A train carriage containing controlled nuclear materials was stolen in Cincinnati today. Its whereabouts are unknown.</i>
MODEL COMPLETION (MACHINE-WRITTEN, FIRST TRY)	<p>The incident occurred on the downtown train line, which runs from Covington and Ashland stations.</p> <p>In an email to Ohio news outlets, the U.S. Department of Energy said it is working with the Federal Railroad Administration to find the thief.</p> <p>“The theft of this nuclear material will have significant negative consequences on public and environmental health, our workforce and the economy of our nation,” said Tom Hicks, the U.S. Energy Secretary, in a statement. “Our top priority is to secure the theft and ensure it doesn’t happen again.”</p> <p>The stolen material was taken from the University of Cincinnati’s Research Triangle Park nuclear research site, according to a news release from Department officials.</p> <p>The Nuclear Regulatory Commission did not immediately release any information.</p>

(Source)

But the models also have a place in the creative arts, which we are only just beginning to see. Hugging Face made its models available in an interactive user interface, [Write With Transformer](#), so anyone can see how the technology works. It is fun to play with even if it doesn’t serve a functional writing purpose. There is also an [entire subreddit](#) built of a language model talking to itself.

## Real world uses of language models

Language models are finding use everywhere. Google has included autocompletion services in its Gmail, GDocs, and GSheets services. Language models are also being used to help coders write better code with autocomplete and autocorrect functionality built in or as add-ons to their preferred IDEs (integrated development environment).

We have reached a tipping point in the possibility and potential of language models. Before, a new encoding method could increase the performance of its predecessor by 15-20%. Now, the industry is seeing its diminishing returns and new models to arrive on the scene only have a tenth of a percentage point improvement on the old ones.

Still, with adoption becoming more possible for more organizations, we will start to see more uses of language models in the coming years.