


Opinion **Technology**

## 'Natural language understanding' poised to transform how we work

Written communication will become a less exclusively human task if machines can extract meaning from text

**RICHARD WATERS**



© Alamy

Richard Waters DECEMBER 3 2018

In some ways, the life of an analyst working for one of the US intelligence services is not so different from that of many other knowledge workers.

The day often begins with sifting through information, preparing reports that distil and summarise the most important new events. Intelligence analysts have to work with large amounts of data, analysing and synthesising it to make sense of a complex world.

This is the kind of work where software is starting to play a bigger part. Technology from Primer, a San Francisco [artificial intelligence](#) start-up, is already used by unspecified intelligence services to read through written material in an effort to identify trends and significant events. The results help guide human analysts to focus on what is important.

The same software is used by retailer [Walmart](#), where analysts constantly monitor a large number of product markets to identify opportunities and risks in the company's supply chain.

If language understanding can be automated in a wide range of contexts, it is likely to have a profound effect on many professional jobs. Communication using written words plays a central part in many people's working lives. But it will become a less exclusively human task if machines learn how to extract meaning from text and churn out reports.

In some ways, this computer assistance is badly needed. Many information workers struggle to handle the steadily increasing amount of information they are expected to monitor and process.

“All these companies have data that relates to their own particular world: it keeps increasing and they can’t keep up with it,” says Sean Gourley, Primer’s chief executive.

AI has been slower to have an impact on the world of language than in some other fields. The first big breakthroughs for machine learning, the technology behind recent advances in AI, was in 2012. That was when a deep learning system registered a major advance during the annual ImageNet image recognition competition. It was not long before computer vision surpassed the human variety.

## **Like many other products of machine learning, systems like these probably will not replace humans completely**

Voice recognition followed. But this only identifies the words a person has spoken, not their meaning. Using the same technology to understand so-called “natural language”, which humans use to communicate, has been a much harder nut to crack.

Part of the problem is the sheer complexity and ambiguity of language. The neural networks in the most advanced machine learning systems use a form of pattern recognition, relying on past examples to make sense of new information. In the world of language, where meaning depends heavily on context and often the relationship between the people communicating, such techniques are less effective.

Another hindrance in training neural networks to do the type of work analysts face — distilling information from several sources — is the scarcity of appropriate data to train the systems. It would require public data sets that include both source documents and a final synthesis, giving a complete picture that the system could learn from.

Despite challenges such as these, recent leaps in natural language understanding (NLU) have made the systems more effective and brought the technology to a point where it is starting to find its way into many more business applications. New techniques such as transfer learning have helped (see box).

If the basic technology problems can be cracked, then the advantages that machines would have over people in processing language are clear. They can handle information far faster than humans and in much greater volume. They also have greater “dimensionality”: they can analyse things on many different parameters at once.

One result, says Sean Gourley, is that computers should be able to detect far more significant facts in a large body of information than people can. Humans are at their most effective using what he calls precision: identifying an important event or occurrence and homing in on it.

But while finding one thing worth examining more closely, they probably miss many others that might be equally worthy of attention. This is where machines could come in, throwing up many extra possibilities that have been overlooked, whether those are potential areas of demand for a new product or sources of terrorist threats.

Like many other products of machine learning, systems like these probably will not replace humans completely. But they will change the nature of work in fundamental ways and require people with new skills and a different outlook. What it will do to the overall number of analyst jobs that are left, and how exactly the nature of the work will change, are open questions.

One possibility is that demand for human input will fall. With more preprocessed information at their fingertips, fewer, smarter analysts will be needed to make sense of it.

But it is also possible that the impact of language systems will be felt not in terms of the number of analysts who are needed, but rather in the number of things that get analysed.

As in other areas of information processing, automating some of the work will greatly reduce the costs of two tasks that were previously the preserve of humans: reading and writing. When some of the work does not require a highly paid analyst, many new things will be susceptible to analysis where this was previously uneconomic.


This could be the big dividend of the AI era. As the costs of generating automated intelligence fall, results spewed out by machine learning systems like this should become available as a common input for many types of human decision-making.

Meanwhile, with falling costs, the sheer volume of language being processed and the volume of new texts being produced by language generation engines could explode. Today's information overload would seem paltry by comparison.

In tomorrow's world of work, that could make computers as essential to processing information in language form as they are to handling other forms of data today. But flesh-and-blood analysts should still have the final say: only humans, after all, are in a position to pass judgment on what really matters in the human world.

## How Wikipedia is helping computers learn about the world

---



Making connections: Digital-language interpretation systems mimic human neural networks

Teaching a digital “neural network” to solve a particular task is usually a painstaking process. It involves training the system from scratch using a large amount of relevant data that has been labelled, usually by humans.

Transfer learning — taking a general-purpose model trained on one data set and fine-tuning it to work on a different problem — is one of the most promising techniques for making the technology more flexible and cost effective. One approach that could produce significant results was outlined this year by natural language researchers Jeremy Howard and Sebastian Ruder.

They began with a model that was trained on a large subset of [Wikipedia](#) — essentially a general-purpose language model.

They then fine-tuned it for different purposes, using as few as 100 labelled examples. It would have taken 10,000 labelled examples to train a neural network for the same task from scratch.

The researchers say that in the field of computer vision this technique has already produced a general-purpose model capable of doing everything from “improving crop yields in Africa to building robots that sort Lego bricks”. A similar approach in language processing could open up many specialised uses.

[Copyright](#) The Financial Times Limited 2019. All rights reserved.