GREGORY BARBER BUSINESS 09.07.2018 01:55 PM

AI Can Recognize Images. But Can It Understand This Headline?

New approaches foster hope that computers can comprehend paragraphs, classify email as spam, or generate a satisfying end to a short story.

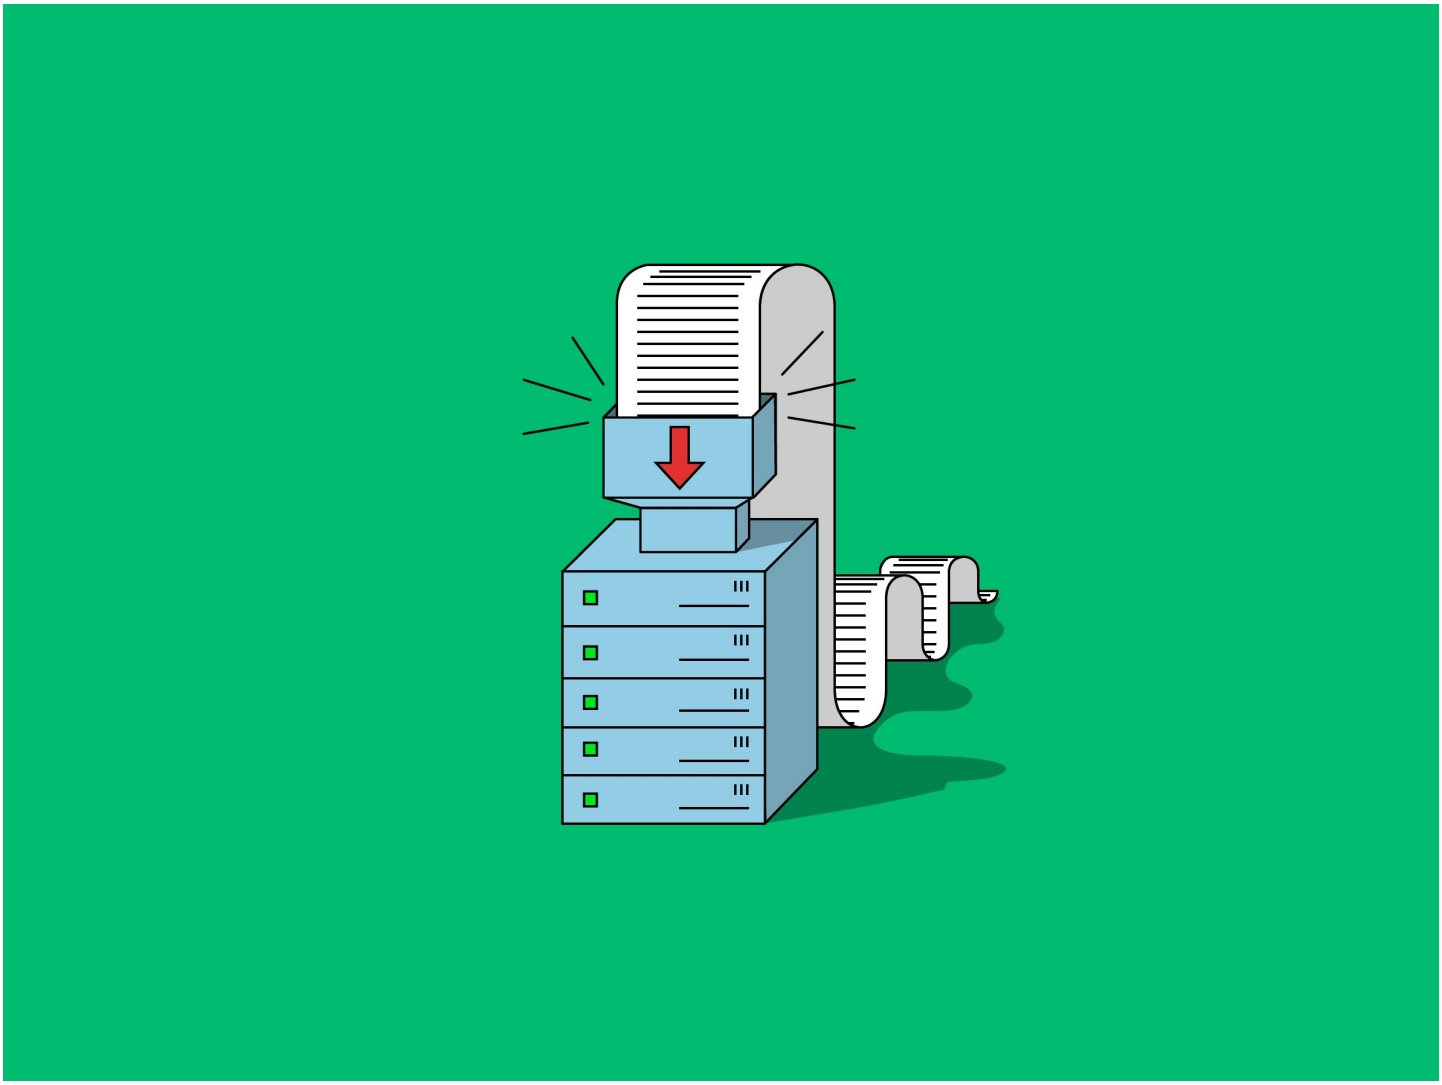

CASEY CHIN

reading a neural network millions of labeled images from a database called ImageNet. It ushered in an exciting phase for computer vision, as it became clear that a model trained using ImageNet could help tackle all sorts of image-recognition problems. Six years later, that's helped pave the way for self-driving cars to navigate city streets and Facebook to automatically tag people in your photos.

In other arenas of AI research, like understanding language, similar models have proved elusive. But recent research from fast.ai, OpenAI, and the Allen Institute for AI suggests a potential breakthrough, with more robust language models that can help researchers tackle a range of unsolved problems. Sebastian Ruder, a researcher behind one of the new models, calls it his field's "ImageNet moment."

The improvements can be dramatic. The most widely tested model, so far, is called Embeddings from Language Models, or ELMo. When it was released by the Allen Institute this spring, ELMo swiftly toppled previous bests on a variety of challenging tasks---like reading comprehension, where an AI answers SAT-style questions about a passage, and sentiment analysis. In a field where progress tends to be incremental, adding ELMo improved results by as much as 25 percent. In June, it was awarded best paper at a major conference.

Dan Klein, a professor of computer science at UC Berkeley, was among the early adopters. He and a student were at work on a constituency parser, a bread-and-butter tool that involves mapping the grammatical structure of a sentence. By adding ELMo, Klein suddenly had the best system in the world, the most accurate by a surprisingly wide margin. "If you'd asked me a few years ago if it was possible to hit a level that high, I wouldn't have been sure," he says.

Models like ELMo address a core issue for AI-wielding linguists: lack of labeled data. In order to train a neural network to make decisions, many language problems require data that's been meticulously labeled by hand. But producing that data takes time and money, and even a lot of it can't capture the unpredictable ways that we speak and write. For languages other than English, researchers often don't have enough labeled data to accomplish even basic tasks.

"We're never going to be able to get enough labeled data," says Matthew Peters, a research scientist at the Allen Institute who led the ELMo team. "We really need to develop models that take messy, unlabeled data and learn as much from it as possible."

Luckily, thanks to the internet, researchers have plenty of messy data from sources like Wikipedia, books, and social media. The strategy is to feed those words to a neural network and allow it to discern patterns on its own, a so-called "unsupervised" approach. The hope is that those patterns will capture some general aspects of language---a sense of what words are, perhaps, or the basic contours of grammar. As with a model trained using ImageNet, such a language model could then be fine-tuned to master more specific tasks---like summarizing a scientific article, classifying an email as spam, or even generating a satisfying end to a short story.

That basic intuition isn't new. In recent years, researchers have delved into unlabeled data using a technique called word embeddings, which maps how words relate to each other based on how they appear in large amounts of text. The new models aim to go deeper than that, capturing information that scales up from words up to higher-level concepts of language. Ruder, who has written about the potential for those deeper models to be useful for a variety of language problems, hopes they will become a simple replacement for word embeddings.

sentences rather than words. That extra context makes the model good at parsing the difference between, say, "May the month and "may" the verb, but also means it learns about syntax. ELMo gets an additional boost by gaining an understanding of subunits of words, like prefixes and suffixes. Feed a neural network a billion words, as Peters' team did, and this approach turns out to be quite effective.

It's still unclear what the model actually learns in the process of analyzing all those words. Because of the opaque ways in which deep neural networks work, it's a tricky question to answer. Researchers still have only a hazy understanding of why image-recognition systems work so well. In a new paper to appear at a conference in October, Peters took an empirical approach, experimenting with ELMo in various software designs and across different linguistic tasks. "We found that these models learn fundamental properties of language," Peters says. But he cautions other researchers will need to test ELMo to determine just how robust the model is across different tasks, and also what hidden surprises it may contain.

One risk: encoding biases from the data used to train them, so doctors are labeled as men, and nurses as women, for example, as word embeddings have previously done. And while the initial results generated by tapping ELMo and other models are exciting, says Klein, it's unclear how far the results can be pushed, perhaps by using more data to train the models, or by adding constraints that force the neural network to learn more effectively. In the long run, AI that reads and talks as fluently as we do may require a new approach entirely.

More Great WIRED Stories

- Kelly Slater's artificial surf pool is really making waves
- This Bugatti goes 18 mph (and is made of Lego)
- How to use Twitter: critical tips for new users
- PHOTO ESSAY: A world without electricity
- Everything you want to know about quantum computing
- Hungry for even more deep dives on your next favorite topic? Sign up for the Backchannel newsletter

Gregory Barber is a staff writer at WIRED who writes about blockchain, AI, and tech policy. He graduated from Columbia University with a bachelor's degree in computer science and English literature and now lives in San Francisco.

STAFF WRITER

TOPICS NEURAL NETWORKS ARTIFICIAL INTELLIGENCE MACHINE LEARNING

MORE FROM WIRED

Facial Recognition Tech Is Coming to a School Near You

WIRED STAFF

A Netflix Series Explores the Brave New World of Crispr

MEGAN MOLTENI

The Best Mobile Games For Your Hectic Commute

AMIT KATWALA, WIRED UK

Rolls-Royce Turns a 747 Into a Flying Lab for New Engines

AARIAN MARSHALL

4 College Admissions Essay Editing Services Reviewed

CHRISTOPHER NULL

This Week's Cartoons: Stock Photos, Robocalls, and Ear Transplants

WIRED CARTOONS

Can Tiny Glass Beads Keep Arctic Ice From Melting? Maaaybe

ERIC NIILER

Juul, Under Heavy Fire, Pulls Fruit-Flavored Pods From US

SARA HARRISON