


Robots are creating images and telling jokes: Five things to know about foundation models and the next generation of AI

April 13 2022, by Aaron J. Snowwell and Dan Hunter



An image created by DALL-E 2 in response to the prompt 'a robot hand drawing'. Credit: OpenAI

If you've seen photos of [a teapot shaped like an avocado](#) or read a well-written article that [veers off on slightly weird tangents](#), you may have been exposed to a new trend in artificial intelligence (AI).

Machine learning systems called [DALL-E](#), [GPT](#) and [PaLM](#) are making a splash with their incredible ability to generate creative work.

These systems are known as "foundation models" and are not all hype and party tricks. So how does this new approach to AI work? And will it be the end of human creativity and the start of a deep-fake nightmare?

DALL-E 2 is here! It can generate images from text, like "teddy bears working on new AI research on the moon in the 1980s."

It's so fun, and sometimes beautiful. <https://t.co/XZmh6WkMASpic.twitter.com/3zOu30IqCZ>

— Sam Altman (@sama) [April 6, 2022](#)

1. What are foundation models?

[Foundation models](#) work by training a single huge system on large amounts of general data, then adapting the system to new problems. Earlier models tended to start from scratch for each new problem.

DALL-E 2, for example, was trained to match pictures (such as a photo

of a pet cat) with the caption ("Mr. Fuzzyboots the tabby cat is relaxing in the sun") by scanning hundreds of millions of examples. Once trained, this model knows what cats (and other things) look like in pictures.

But the model can also be used for many other interesting AI tasks, such as generating new images from a caption alone ("Show me a koala dunking a basketball") or editing images based on written instructions ("Make it look like this monkey is paying taxes").

Our newest system DALL·E 2 can create realistic images and art from a description in natural language. See it here:

<https://t.co/Kmjko82YO5> pic.twitter.com/QEh9kWUE8A

— OpenAI (@OpenAI) [April 6, 2022](#)

2. How do they work?

Foundation models run on "[deep neural networks](#)," which are loosely inspired by how the brain works. These involve sophisticated mathematics and a huge amount of computing power, but they boil down to a very sophisticated type of pattern matching.

For example, by looking at millions of example images, a deep neural network can associate the word "cat" with patterns of pixels that often appear in images of cats—like soft, fuzzy, hairy blobs of texture. The more examples the model sees (the more data it is shown), and the bigger the model (the more "layers" or "depth" it has), the more complex these patterns and correlations can be.

Foundation models are, in one sense, just an extension of the "[deep learning](#)" paradigm that has dominated AI research for the past decade. However, they exhibit un-programmed or "emergent" behaviors that can be both surprising and novel.

For example, Google's PaLM language model seems to be able to produce explanations for complicated metaphors and jokes. This goes beyond simply [imitating the types of data it was originally trained to process](#).

3. Access is limited, for now

The sheer scale of these AI systems is difficult to think about. PaLM has *540 billion* parameters, meaning even if everyone on the planet memorized 50 numbers, we still wouldn't have enough storage to reproduce the model.

The models are so enormous that training them requires massive amounts of computational and other resources. One estimate put the cost of training OpenAI's language model GPT-3 at [around US\\$5 million](#).

As a result, only huge tech companies such as OpenAI, Google and Baidu can afford to build foundation models at the moment. These companies limit who can access the systems, which makes economic sense.

Usage restrictions may give us some comfort these systems won't be used for nefarious purposes (such as generating fake news or defamatory content) any time soon. But this also means independent researchers are unable to interrogate these systems and share the results in an open and accountable way. So we don't yet know the full implications of their use.

4. What will these models mean for 'creative' industries?

More foundation models will be produced in coming years. Smaller models are already being published in [open-source forms](#), [tech](#)

[companies](#) are starting to [experiment with licensing and commercializing these tools](#) and AI researchers are working hard to make the technology more efficient and accessible.

The remarkable creativity shown by models such as PaLM and DALL-E 2 demonstrates that creative professional jobs could be impacted by this technology sooner than initially expected.

Traditional wisdom always said robots would displace "blue collar" jobs first. "White collar" work was meant to be relatively safe from automation—especially professional work that required creativity and training.

Deep learning AI models already exhibit super-human accuracy in tasks like reviewing X-rays and [detecting the eye condition macular degeneration](#). Foundation models may soon provide cheap, "good enough" creativity in fields such as advertising, copywriting, stock imagery or graphic design.

The future of professional and [creative work](#) could look a little different than we expected.

5. What this means for legal evidence, news and media

Foundation models will inevitably [affect the law](#) in areas such as intellectual property and evidence, because we won't be able to assume [creative content is the result of human activity](#).

We will also have to confront the challenge of disinformation and misinformation generated by these systems. We already face enormous problems with disinformation, as we are seeing in the unfolding Russian

invasion of Ukraine and the nascent problem of deep fake images and video, but foundation models are poised to super-charge these challenges.

Time to prepare

As researchers who [study the the effects of AI on society](#), we think foundation models will bring about huge transformations. They are tightly controlled (for now), so we probably have a little time to understand their implications before they become a huge issue.

The genie isn't quite out of the bottle yet, but foundation models are a very big bottle—and inside there is a very clever genie.

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation

Citation: Robots are creating images and telling jokes: Five things to know about foundation models and the next generation of AI (2022, April 13) retrieved 16 April 2024 from <https://techxplore.com/news/2022-04-robots-images-foundation-ai.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--