



Balanced reporting and expert analysis
Try the FT's comprehensive coverage for just £1.

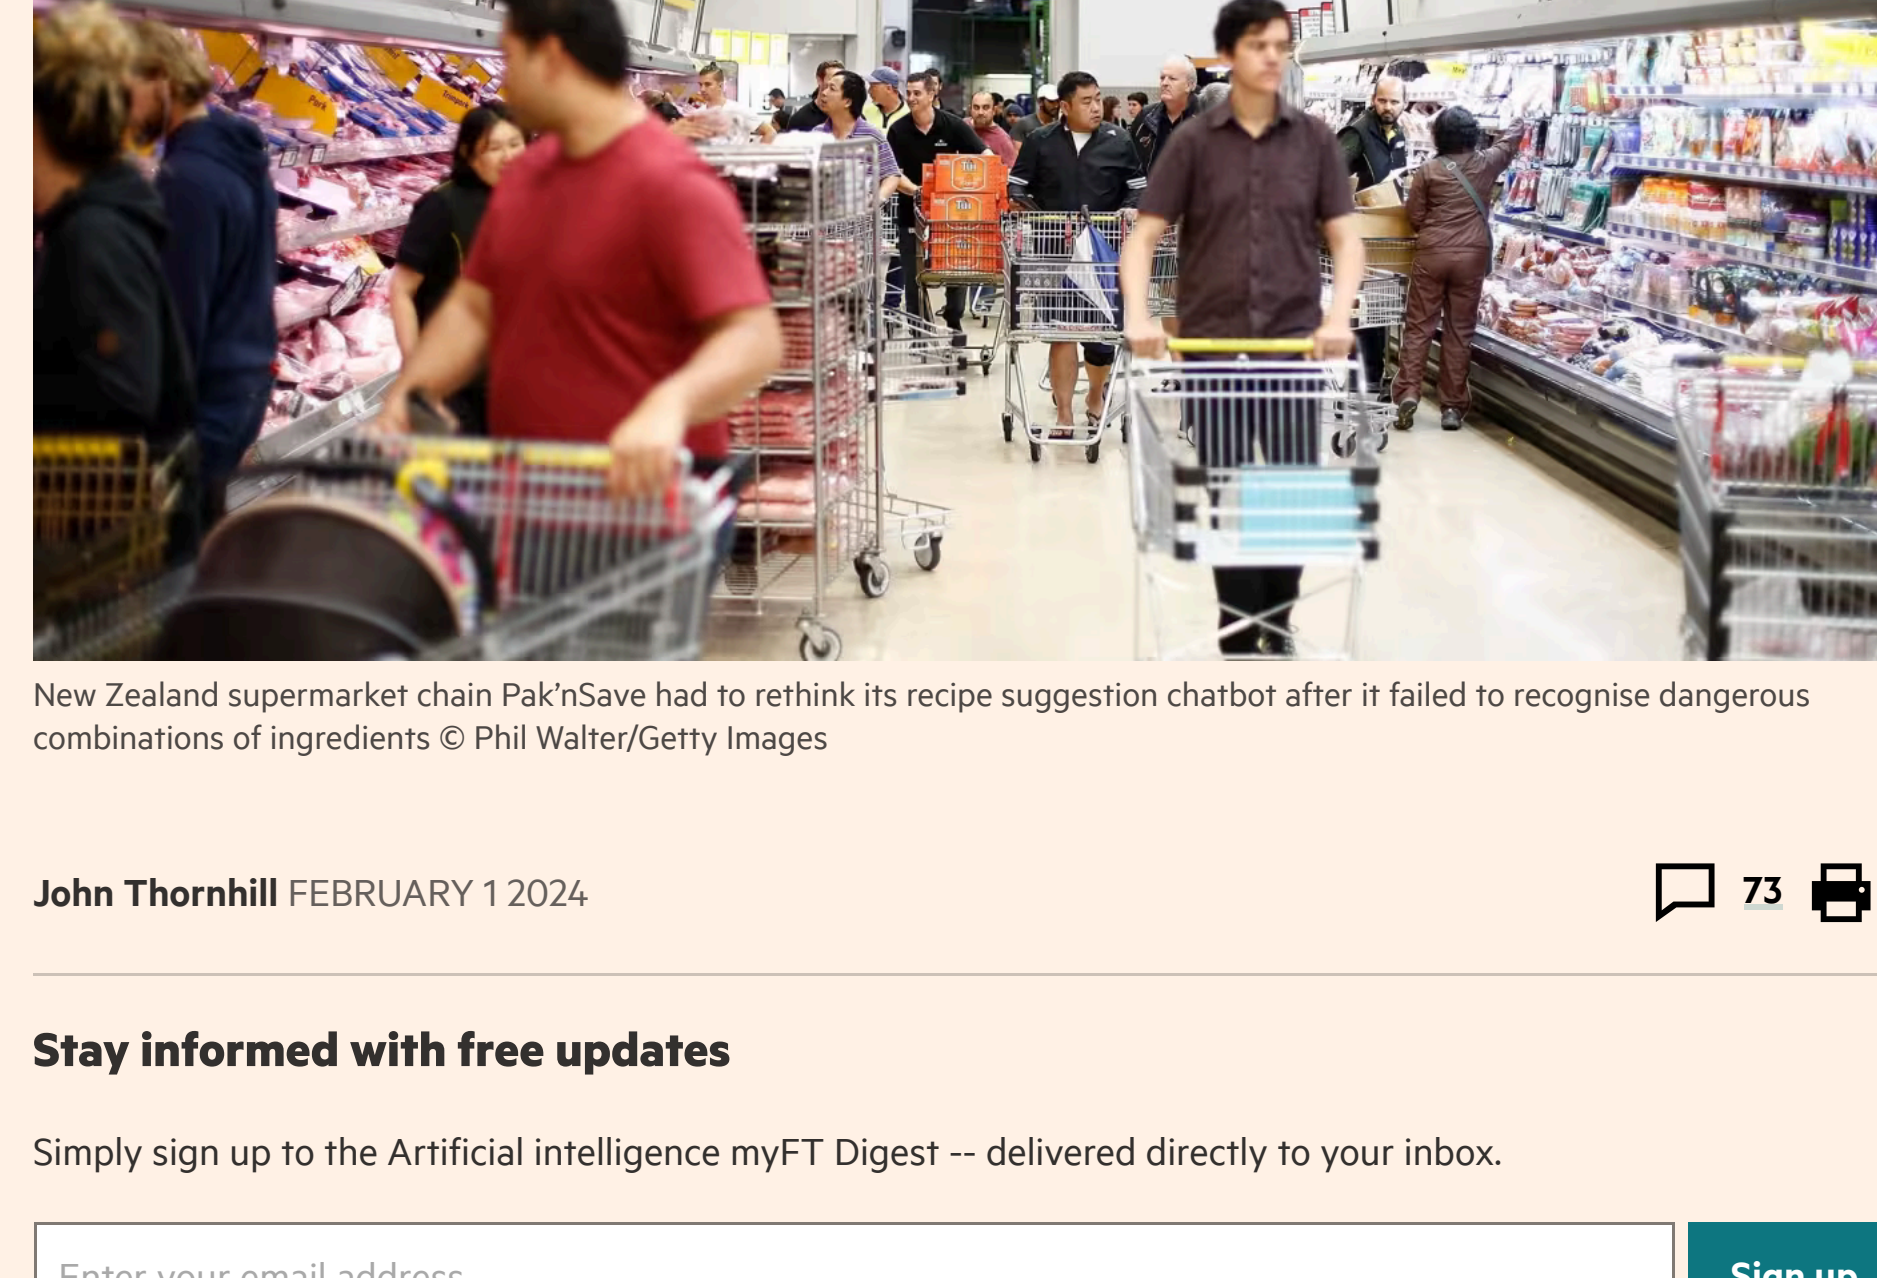

TRY FOR £1

Opinion Artificial intelligence

Why AI hallucinations can be a good thing

Generative artificial intelligence should be welcomed as a giant mash-up machine to enhance creativity

JOHN THORNHILL [+ Add to myFT](#)



New Zealand supermarket chain Pak'nSave had to rethink its recipe suggestion chatbot after it failed to recognise dangerous combinations of ingredients © Phil Walter/Getty Images

John Thornhill FEBRUARY 1 2024

73

Stay informed with free updates

Simply sign up to the Artificial intelligence myFT Digest -- delivered directly to your inbox.

Enter your email address [Sign up](#)

The writer is founder of *Sifted*, an FT-backed site about European start-ups

The tendency of generative artificial intelligence systems to “hallucinate” — or simply make stuff up — can be zany and sometimes scary, as one New Zealand supermarket chain found to its cost. After Pak'nSave released a chatbot last year offering recipe suggestions to thrifty shoppers using leftover ingredients, its *Savey Meal-bot* recommended one customer make an “aromatic water mix” that would have produced chlorine gas.

Lawyers have also learnt to be wary of the [output of generative AI models](#), given their ability to invent wholly fictitious cases. A [recent Stanford University study](#) of the responses generated by three state of the art generative AI models to 200,000 legal queries found hallucinations were “pervasive and disturbing”. When asked specific, verifiable questions about random federal court cases, [OpenAI's ChatGPT 3.5 hallucinated 69 per cent](#) of the time while Meta's Llama 2 model hit 88 per cent.

To many users, generative AI's hallucinations are an irritating bug that they assume the tech companies will fix one day, just like email spam. To what extent companies can do so is currently a [subject of active research](#) and fierce contention. Some researchers argue that hallucinations are inherent to the technology itself. Generative AI models are probabilistic machines trained to give the most statistically likely response. It is hard to code in human traits such as common sense, context, nuance or reasoning.

While technologists try to figure that one out, some users find machine-generated fiction to be an exciting feature and have been exploring its creative possibilities. “I love hallucinations,” says Martin Puchner, a Harvard University professor and author. “They are so cool when it comes to artistic creation.”

As the author of *Culture: A New World History*, Puchner has a distinctive view of creativity. In his book, he argues that for millennia humans have been assimilating the inputs of preceding generations and other cultures and mixing them up to generate fresh syncretistic outputs. Consider how much imperial Rome borrowed from ancient Greece, or the Italian renaissance drew from Arabic scholarship, or how Japan adopted Chinese writing and philosophy, or Jamaican Rastafarian culture absorbed Ethiopian traditions.

Every civilisation, Puchner writes, tends to overestimate the originality of their own culture to prop up dubious claims of superiority and ownership. “Such claims conveniently forget that everything comes from somewhere, is dug up, borrowed, moved, purchased, stolen, recorded, copied, and often misunderstood,” he writes. “Culture is a huge recycling project.”

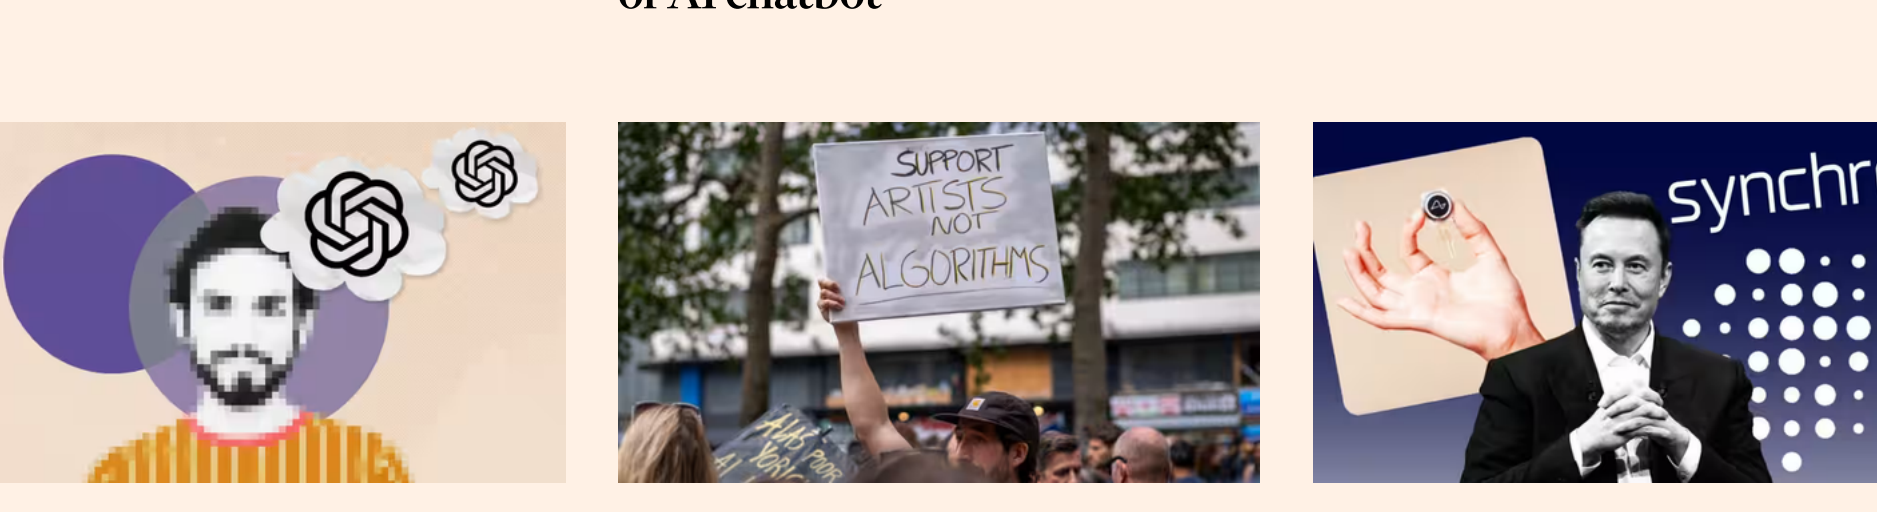
The parallels with generative AI are striking. To some extent, our machines are doing today what humans have been doing forever: mashing up different cultural inputs to generate slightly varied outputs. In that sense, generative AI can be seen as a giant cultural syncretism machine. Reliant on imperfect, partial, data and overconfident in generalising from the particular, machines may be more like fallible humans than we sometimes imagine. Hallucinations may not be so much an algorithmic aberration as a reflection of human culture.

That all sounds abstract. What does it mean in practice? In short, everything depends on the use case. Generative AI models can be a fantastic tool for enhancing human creativity by generating new ideas and content, especially in music, images and video.

If prompted in the right way, these models can act as a useful sparring partner and a tireless source of inspiration for creative projects. They can be the algorithmic equivalent of thinking outside the box. Puchner has himself been [experimenting with customised chatbots](#) to converse with historical figures, such as Aristotle, Confucius and Jesus, based on their words and ideas. “Prompt engineering should be part of Harvard's syllabus,” Puchner says.

The founder of one generative AI company tells me we are rapidly entering the age of “generative media”. “The internet and social media brought the cost of distribution down to zero. Now, thanks to gen AI, the cost of creation is also going close to zero,” he says. That trend could lead to an explosion of creativity, for both good and bad. It may also deepen concerns over disinformation and intellectual property theft.

Impressive as it is as a fiction writer, generative AI still has a long way to go when it comes to reliable non-fiction. That is fine so long as we use it as a tool to enhance human capabilities rather than imagine it as an agent that can replace all human processing power. To hallucinate a phrase: caveat promp



Copyright The Financial Times Limited 2024. All rights reserved.

[Reuse this content](#) [Comments](#)

Latest on Artificial Intelligence

Nvidia
'Sell Nvidia'
1 HOUR AGO

Microsoft Corp
Microsoft in deal with Semafor to create news stories with aid of AI chatbot

NEWS BRIEFING
FT News Briefing podcast 11 min listen
Turkey's central bank chief steps down

PARTNER CONTENT
by UNICEF
How an executive chief went from leader to philanthropic investor

Work Watch Andrew Hill
Can AI make brainstorming less mind-numbing?

Artificial Intelligence
UK shelves proposed AI copyright code in blow to creative industries

News in-depth Elon Musk
Musk's Neuralink hype irks rivals yet shines a light on brain implants

Janan Ganesh
The digital age hasn't gone far enough

Comments

Comment guidelines
Please keep comments respectful. Use plain English for our global readership and avoid using phrasing that could be misinterpreted as offensive. By commenting, you agree to abide by our [community guidelines](#) and [these terms and conditions](#). We encourage you to report inappropriate comments.

Post a comment

[B](#) [I](#) [L](#) [99](#) [≡](#)

[Sign in and Join the Conversation](#)

All Comments [12](#) Sort by [Newest](#)

Tommy J 1 DAY AGO
Good to hear that we're not yet at the matrix phase of AI I guess...

[Recommend](#) [Reply](#) [Share](#) [Report](#)