

Do you trust AI to write the news? It already is—and not without issues

November 6 2023, by Rob Nicholls



Credit: ThisIsEngineering from Pexels

Businesses are increasingly using artificial intelligence (AI) to generate media content, including news, to engage their customers. Now, we're even seeing AI used for the "gamification" of news—that is, to create

interactivity associated with news content.

For better or worse, AI is changing the nature of news media. And we'll have to wise up if we want to protect the integrity of this institution.

How did she die?

Imagine you're reading a tragic article about the [death of a young](#) sports coach at a prestigious Sydney school.

In a box to the right is a poll asking you to speculate about the cause of death. The poll is AI-generated. It's designed to keep you engaged with the story, as this will make you more likely to respond to advertisements provided by the poll's operator.

This scenario isn't hypothetical. It was played out in [The Guardian's recent reporting](#) on the death of Lilie James.

Under a [licensing agreement](#), Microsoft [republished The Guardian's story](#) on its news app and website Microsoft Start. The poll was based on the content of the article and displayed alongside it, but [The Guardian had no involvement](#) or control over it.

If the article had been about an upcoming sports fixture, a poll on the likely outcome would have been harmless. Yet this example shows how problematic it can be when AI starts to mingle with news pages, a product traditionally curated by experts.

The incident led to reasonable anger. In a letter to Microsoft president Brad Smith, Guardian Media Group chief executive Anna Bateson said it was "an inappropriate use of genAI [generative AI]", which caused "significant reputational damage" to *The Guardian* and the journalist who wrote the story.

Naturally, the poll was removed. But it raises the question: why did Microsoft let it happen in the first place?

The consequence of omitting common sense

The first part of the answer is that supplementary news products such as polls and quizzes [actually do engage](#) readers, as [research](#) by the Center for Media Engagement at the University of Texas has found.

Given how cheap it is to use AI for this purpose, it seems likely news businesses (and businesses displaying others' news) will continue to do so.

The second part of the answer is there was no "human in the loop", or limited human involvement, in the Microsoft incident.

The major providers of large language models—the models that underpin various AI programs—have a financial and reputational incentive to make sure their programs don't cause harm. Open AI with its [GPT- models and DAll-E](#), Google with PaLM 2 (used in [Bard](#)), and Meta with its [downloadable Llama 2](#) have all made significant efforts to ensure their models don't generate harmful content.

They often do this through a process called "reinforcement learning", where humans curate responses to questions that might lead to harm. But this doesn't always prevent the models from producing inappropriate content.

It's likely Microsoft was relying on the low-harm aspects of its AI, rather than considering how to minimize harm that may arise through the actual use of the model. The latter requires common sense—a trait that can't be programmed into large language models.

Thousands of AI-generated articles per week

Generative AI is becoming accessible and affordable. This makes it attractive to commercial news businesses, which have been reeling from [losses of revenue](#). As such, we're now seeing AI "write" news stories, saving companies from having to pay journalist salaries.

In June, News Corp executive chair [Michael Miller revealed](#) the company had a small team that produced about [3,000 articles a week](#) using AI.

Essentially, the team of four ensures the content makes sense and doesn't include "hallucinations": false information made up by a model when it can't predict a suitable response to an input.

While this news is likely to be accurate, the same tools can be used to generate potentially misleading content parading as news, and nearly indistinguishable from articles written by professional journalists.

Since April, a NewsGuard investigation [has found hundreds](#) of websites, written in several languages, that are mostly or entirely generated by AI to mimic real news sites. Some of these included harmful misinformation, such as the claim that US President Joe Biden [had died](#).

It's thought the sites, which were teeming with ads, were likely generated to get ad revenue.

As technology advances, so does risk

Generally, many large language models have been limited by their underlying training data. For instance, models trained on data up to 2021 will not provide accurate "news" about the world's events in 2022.

However, this is changing, as models can now be fine-tuned to respond to particular sources. In recent months, the use of an AI framework called "[retrieval augmented generation](#)" has evolved to allow models to use very recent data.

With this method, it would certainly be possible to use licensed content from a small number of news wires to create a news website.

While this may be convenient from a business standpoint, it's yet one more potential way that AI could push humans out of the loop in the process of news creation and dissemination.

An editorially curated news page is a valuable and well-thought-out product. Leaving AI to do this work could expose us to all kinds of misinformation and bias (especially without human oversight), or result in a lack of important localized coverage.

Cutting corners could make us all losers

Australia's News Media Bargaining Code was designed to "level the playing field" between big tech and media businesses. Since the code came into effect, a secondary change is now flowing in from the use of generative AI.

Putting aside click-worthiness, there's currently no comparison between the quality of news a journalist can produce and what AI can produce.

While generative AI could help augment the work of journalists, such as by helping them sort through large amounts of content, we have a lot to lose if we start to view it as a replacement.

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

Provided by The Conversation

Citation: Do you trust AI to write the news? It already is—and not without issues (2023, November 6) retrieved 8 May 2024 from <https://techxplore.com/news/2023-11-ai-news-isand-issues.html>

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.