Time will tell just how much of an impact AI-generated content had on the outcome of the election

Picture credit: Getty

# AI AND THE US ELECTION

**Christopher Thissen** *considers the part that artificial intelligence played in the recent race for the Presidency*

**This Presidential election campaign has seen myriad stories and comments published online by supporters on both sides looking to influence voters. While many were written by humans, an increasing number were generated by AI. Advances in machine learning mean AI generated text is now almost indistinguishable from anything written by people.**

While AI created content has many legitimate uses, such as helping overworked journalists write articles, there are instances where it can be used to promote fake news or a disinformation agenda. This can be through making up stories to influence public opinion, burying real news in a deluge of bogus reportage or increasing the number of comments a genuine article receives to make issues seem more popular than they actually are.

AI-generated text is used extensively by businesses to respond to customers that need help on their websites or via apps. Customers input a question into a chat box and the chatbot will respond with further questions to point them in the right direction, such as a specific page on the website or handing them over to a human operator.

Similar technology is now being used to create original content, such as news articles, blog posts and even creative writing. The latest advancement in AI language modeling technology is GPT-3, developed by OpenAI, which creates articles so realistic they are difficult to distinguish from those written by humans.

AI is now used regularly by big news outlets to write articles that are heavily based on statistical information, such as financial reports or data from sporting events. In fact, around a third of the content published by Bloomberg News is at least in part written by a machine learning program called Cyborg.

In the age of 24-hour, instant-access news, AI is an incredibly useful tool for generating reports quickly. In the world of finance, news outlets are often presented with lengthy company reports, sometimes hundreds of pages long, containing thousands of statistics. Going through this much information manually to find a good news angle can take days. On top of this, there is also the race against competitors to be the first to break the news. Instead, the reports can be uploaded into an AI program that extracts much of the relevant information and creates an article based on this data within minutes. The article is then checked by an editor to ensure the right story has been created from the extracted information. In this way, financial stories can be created at a much faster rate than using traditional methods. A similar process is used in sports journalism to rapidly produce match reports.

## KEEPING IT REAL

The likes of GPT-3 are taking this a step further. With GPT-3, AI-generated articles are no longer limited to the structured, data rich formats common in sports and finance. Instead, these next-generation AIs can create realistic articles about nearly any topic, often without requiring additional human intervention or editing.

Troublingly, AI also excels at creating fake news. Such stories are very easy to create – seeded with only a title, subtitle and perhaps a political slant, the AI can generate the rest of the article. Of course, you don't need a sophisticated AI to generate false and outlandish stories that might impact an election – plenty of human-generated fiction went viral during the 2016 election. Rather, the concern is the sheer volume and rapidity of original, convincing content that the AI can create.

With the latest AI, it would be straightforward to devise a list of topics, quickly generate hundreds or thousands of articles with a specific political bent or conspiracy theory in mind, and automatically share the stories on social media or blogs. A study conducted by Oxford University found that roughly half of all news shared on Twitter in Michigan during the 2016 election came from untrustworthy sources. By generating realistic content that is difficult to distinguish from professional journalism, this percentage could rise much higher, burying true news stories in vast amounts of AI-generated fake articles.

The speed and scale at which realistic content may be generated is particularly concerning for issues surrounding election results. As recently as September, the FBI and the Department of Homeland Security issued a joint announcement warning that delays in final tallies may provide an opportunity for foreign actors and cybercriminals to "create or share corresponding social media content… in an attempt to discredit the electoral process and undermine confidence in US democratic institutions." AI's may facilitate rapid generation of articles about voter suppression, cyber attacks on election infrastructure, voter fraud and others in attempt to make the election results appear illegitimate.

Social media platforms – where users can post articles regardless of veracity – are a common target for disinformation campaigns and citizens that get their news from social media sites are especially vulnerable to disinformation. A study by the Pew Research Center found that one in five US citizens primarily get their daily political news from social media, while a quarter find out what is happening in the world via news websites. This means that a significant proportion of the US population are likely to have viewed news stories that have been partly or wholly generated by AI. In most cases, the consumers would have been unaware of this.

> THE CONCERN IS THE SHEER VOLUME OF ORIGINAL, CONVINCING CONTENT AI CAN CREATE

The research also found that those who used the likes of Facebook and Twitter as their main or only source of news were less knowledgeable about current events. This means that not only are these users more likely to be exposed to disinformation campaigns, but also they are less likely to doubt what they are reading or to find other sources to confirm or deny the news stories that appear in their streams.

AI may also be used to generate the appearance of a political movement, further influencing voters. A study by the Pew Research Center of the FCC's 2014 open comment period on net neutrality found that the vast majority of the 21.7-million online comments likely resulted from organised campaigns designed to influence the outcome of the FCC's decision by flooding the board with fake messages.

Most of these fake comments (which comprise 94 percent of all comments) were identified by being copies of other comments. The more realistic comments generated by AI would be more difficult to detect. Comment bots may already be leveraged on social media sites and forums to make specific opinions appear widely held. Rather than influence specific government organisations, these bots would seek to influence public opinion. Research demonstrates that an individual is more likely to adopt a viewpoint when it is held by others in their social network, even if the majority opinion is only an illusion.

Government bodies in the US were clearly concerned about the influence disinformation, including that generated by AI, might have on the election. Earlier this year, those departments and agencies responsible for protecting the nation, including the Departments of Defense, Justice and Homeland Security as well as the FBI and NSA, released a joint statement about disinformation in elections. They warned US citizens to be vigilant about information they consumed and that a well-informed public was the best defence against disinformation.

People power can potentially be a critical weapon in the fight against the spread of disinformation and fake news. Indeed, Facebook is using crowdsourcing in an attempt to remove fake news from its platform.

Yet spotting disinformation requires those that are reading it to have a decent grasp of current affairs, as well as a good understanding of the language in which it is written when it comes to computer-generated content. OpenAI says that it is possible to tell that an article has been generated by GPT-3 through factual inaccuracies, repetition, non-sequiturs and unusual phrasings, but notes that these indicators may be rather subtle. Given the frequency with which many readers skim articles, such indicators are likely to be missed.

> AI ARTICLES ARE NOW SO REALISTIC THEY ARE DIFFICULT TO DISTINGUISH FROM THOSE BY HUMANS

There are others who are taking the approach of fighting AI with AI. For instance, the US Defense Advanced Research Projects Agency (DARPA) is developing the Semantics Forensics (SemaFor) program. DARPA says the project aims to: "develop technologies to automatically detect, attribute and characterise falsified, multi-modal media assets (eg text, audio, image, video) to defend against large-scale, automated disinformation attacks". However, the project is not expected to be functional before 2024.

Denying access to the advanced AI used to generate text is another option for limiting the impact of the technology to produce disinformation. OpenAI has pledged to restrict the availability of the GPT-3 AI for ethical uses only, closely monitoring its API (currently in private beta). But API keys are sometimes shared or accidentally exposed in public repositories. Now that the power of these AI capabilities has been proven, no doubt other groups will try to replicate the results for both personal and political gain. Replication efforts are already underway. Whether threat actors ultimately adopt this type of technology depends on whether it improves their workflow, making it easier to achieve their goals.

Disinformation undermines the democratic process. Controlling the proliferation of fake news will only get harder as AI is used more extensively to create it. It is ultimately down to us as individuals to question what we read and find additional sources to back up any claims that seem far fetched or controversial. This applies to all news and commentary, whether written by humans or machines.

Though it is still too early to tell how influential machine generated disinformation has been on the result of this election, my hope is not very ●

**Christopher Thissen** is a senior data scientist at Vectra, harnessing the power of machine learning to detect malicious cyber behaviours. Before joining Vectra, Chris led DARPA-funded machine learning research projects at Boston Fusion Corporation.

**Comment bots typically use social media to make specific opinions appear widely held**