

UN-ish speeches cooked by artificial intelligence are quite credible

10 June 2019, by Nancy Cohen



Credit: CC0 Public Domain

Those who worry about artificial intelligence being so good it spins out of control into making humans robo-victims of cooked lies posing as truth had best ignore the recent study which is sure to disturb their sleep. The paper looks at a successful implementation of AI-generated speeches.

Listen, these crazy times will get crazier times. People behind the dotted line will continue yelling [fake news](#) over any item that bears information in conflict with their opinions. People behind the striped line will yell fake news because, well, the words just may be fake news (and photos, and talking head videos) generated through [artificial intelligence](#).

TechCrunch was one of the sites that reported on the research and how it only takes less than ten dollars to fabricate a UN speech using AI. Leave

alone less than ten dollars, the feat was accomplished using even less than that.

"The team got ahold of a readily available language model," said Donovan Alexander in *Interesting Engineering*. He said test components were "a simple deep learning [model](#) and cheap cloud storage."

TechCrunch's Jonathan Shieber said with less than \$7.80 and 13 hours of programming time, "researchers from the United Nations were able to develop a program that could craft realistic-seeming speeches for the United Nations' General Assembly."

Joseph Bullock and Miguel Luengo-Oroz were the two researchers who performed the experiment. Their goal: to train a language model that can generate speech-style text on topics ranging from climate change to terrorism.

They discussed their work in their paper available on arXiv, "Automated Speech Generation from UN General Assembly Statements: Mapping Risks in AI Generated Texts."

Where did the cost and training-time numbers come from? The authors in the paper said that the language model was trained in under 13 hours on NVIDIA K80 GPUs, costing as little as \$7.80 on AWS spot instances.

The two authors of this paper are part of [Global Pulse](#), which is a UN initiative of the United Nations Secretary-General. Its focus is on big data—that is, promoting progress in ways to harness [big data](#) safely and responsibly.

They fed in English language transcripts of speeches between 1970 and 2015 at the UN General Assembly. Results: Their software could generate 50 to 100 words per topic with just the input of one or two sentences relevant to the

headline topic.

Alexander remarked that these speeches "would be [strong](#) enough to engage crowds at the UN today."

Their commonly used model had been trained on Wikipedia, and later fed 40 years of speeches given by [political leaders](#) at the UN General Assembly.

Shieber reported more about the findings. The program In roughly 90 percent of cases was able to generate text credible enough to have come from a General Assembly speaker on a political topic or related to an issue that had been addressed by the secretary general.

Karen Hao, *MIT Technology Review*: The researchers [tested](#) the model on three types of prompts: general topics (e.g. "[climate change](#)"), opening lines from the UN Secretary-General's remarks, and inflammatory phrases (e.g. "immigrants are to blame ...").

Not surprisingly, they found outputs from the first category closely matching real UN speeches roughly 90% of the time. Outputs from the third category, however, required more work to generate, producing convincing outputs about 60% of the time.

Nonetheless, all this left Shieber with a bracing conclusion, that "the age of deepfakes is here and that faked texts could be just as much of a threat as fake videos. Perhaps more, given how cheap they are to [produce](#)."

What about the authors? For sure, their very motivation was, in their own words, "to raise awareness about the dangers of AI text generation to peace and political stability."

Alexander, in particular, called out how easy and inexpensive it was to pull the fake speech text off as credible items. He wrote that "there is a good chance in the near future, with little to no resources, people could create AI speeches and political content to spread misinformation or simply use it for political gains."

Isn't it possible to check with the UN records to

make sure the item in questions is real, not fake? Many people will not bother to check or feel the impulse to do so. The authors wrote that one may generate controversial text for a speech, or even create a 'deep fake' video of the person making the [speech](#), "Given that all of this information can be instantly published via social media, many individuals will not check the original transcript and assume it to be true. "

More effort into detecting and responding to AI generated content in turn will be welcomed.

More information: Automated Speech Generation from UN General Assembly Statements: Mapping Risks in AI Generated Texts, arXiv:1906.01946 [cs.CL] arxiv.org/abs/1906.01946

© 2019 Science X Network

APA citation: UN-ish speeches cooked by artificial intelligence are quite credible (2019, June 10)
retrieved 17 June 2019 from <https://techxplore.com/news/2019-06-un-ish-speeches-cooked-artificial-intelligence.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.