Impersonating BBC News: A Demo of AI's Capability to Imitate Genuine News Outlets

Kyal Pindolia
pindoliak@gmail.com

Funmi "Finn" Okuleye
funmiokuleye01@gmail.com

With
Sage & CeSIA & CivAI & Apart Research

Abstract

This paper presents a demonstration that showcases the current capabilities of AI models to imitate genuine news outlets, using BBC News as an example. The demo allows users to generate a realistic-looking article, complete with a headline, image, and text, based on their chosen prompts. The purpose is to viscerally illustrate the potential risks associated with AI-generated misinformation, particularly how convincingly AI can mimic trusted news sources. By comparing our approach with existing demos like Palisade Research's FoxVox and AI Digest's Election Disruption Demo, we emphasise the importance of creating accessible, engaging, and accurate demonstrations.

We argue that such a demo can effectively raise awareness of the implications of AI advancements for a broad audience, including the potential spread of misinformation. The paper concludes by discussing potential solutions to mitigate these risks and proposes further development to enhance the impact and reach of the demonstration.

1. Introduction

AI capabilities are advancing at a rapid rate, and this trend is likely to continue (Brundage et al., 2018). Current advancements in Large Models which can

generate text and images are likely to have a vast amount of implications for society. Many of these are yet to be explored.

Some of the important current and future AI capabilities are documented in publicly available papers. However, this information may fail to reach everyone who will likely be affected by it. Crucially, the implications of these advancements may not "hit home" to everyone who could be affected by advancing AI, such as workers, policy makers and the public.

Creating effective demonstrations which illustrate the capabilities, possibilities and risks associated with advancements in AI can be an effective way of outlining what the future may hold.

In light of this, we have created a demo which seeks to viscerally highlight what current AI models are capable of, as well as giving an indication of what advancements in these models may look like. Our demo allows users to choose a prompt and then generates a genuine looking headline, image and article that mimics BBC News. The intent behind this is to illustrate how effective AI can be at imitating an authentic news outlet, and highlight how risks associated with this could be mitigated.

Our audience for this demonstration is broad: we want to have an impact on anyone who is familiar with BBC News and has previously read a BBC News article on their phone or computer. We've created a roadmap of a user experience that allows the demo to be simple to use, engaging, and importantly grabs the attention of the user.

In the "Overview" section, we'll outline the reasoning behind some of the choices we have made in this demonstration, as well as map out aspects of this demo we'd be excited to develop further. We came up with this vision and an early prototype during a weekend-long hackathon. In the "Discussion" part of this paper, we explore some of the key areas we would want to polish in the future to ensure that the demo effectively achieves its purpose of safely and viscerally showcasing what AI is capable of to a broad audience.

2. Overview

In crafting our demonstration, the criteria we have chosen to focus on are:

1. Making the demo *Visceral* : Can you feel the implications of advancing AI in this demonstration?
2. Ensuring that the insights highlighted by this demo are *Important*
3. Making sure that the capabilities we illustrate are *Accurate*

In addition to this, we've sought to make a demonstration that was both easy to use and safe to share to an initial select audience.

Our demonstration draws upon inspiration from Palisade Research's FoxVox[1] and AI Digest's Election Disruption Demo[2]. We like the fact that both these demonstrations can be tested by a wide audience and give specific examples of the risks of current AI systems. We particularly like the fact that FoxVox was designed in a way to imitate the formatting of genuine news outlets. One limitation of FoxVox is that it requires the installation of a chrome-extension. From a user perspective, this may limit the pool of people who try the demonstration, as users first have to install the extension, and then search relevant content before seeing the effect of the demo. By contrast, the AI Digest election demo is much simpler to use and improves the scope for showcasing AI's capabilities to a wide audience.

In light of this, we ideated a concept of creating a simple demo that allows a user to easily create an impersonation of a genuine news article. We've chosen BBC News as an optimal outlet to mimic. Some reasons for this are:

a) As of July 2024, BBC News (and links associated with the BBC domain) receive ~1.2 Billion visits per month, which is the highest number of visits for a news outlet.
b) Our opinion is that BBC News is a widely known brand and respected authority in reporting. Creating an imitation of BBC News to demonstrate AI capability therefore adds to the "*viscerality*" of our demonstration.

[1] *FoxVox (Palisade Research):* Explore how AI can use subtle rewording to give a political tilt on recent events
[2] *How can AI disrupt elections? (AI Digest):* Simulate a voice call with an AI that wants to prevent you from voting

A video showcasing our demonstration can be found → here ←

We've made our repository available → here ← too if you like to have a go yourself.

We'll also walk you through the user experience below:

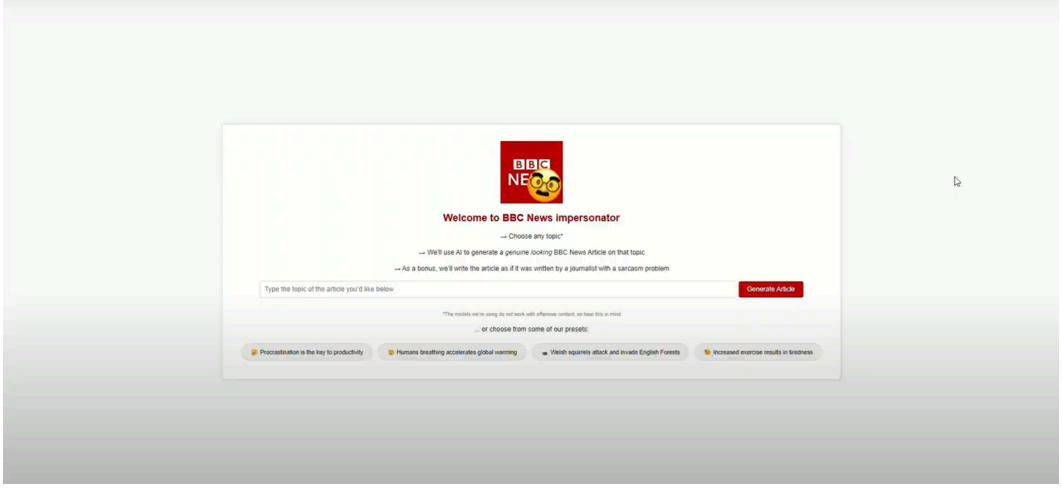

Our landing page is what users first see when they visit our website. The logo for our demo is intended to be eye-catching and memorable: A "disguised face emoji" on top of the BBC News logo. The landing page also explains what the user can expect, and outlines a simple set of instructions to follow to generate an imitation BBC News article.

As we are using GPT-4o to generate the text in the article, and DALLE3 to generate images, controversial topics are automatically rejected according to OpenAI's content policy. Our landing page also gives some suggested prompts that the user can try to generate an article. These also serve as inspiration for the user when choosing their own prompt.

In this example, our chosen prompt was "video games". See the results on the next page →

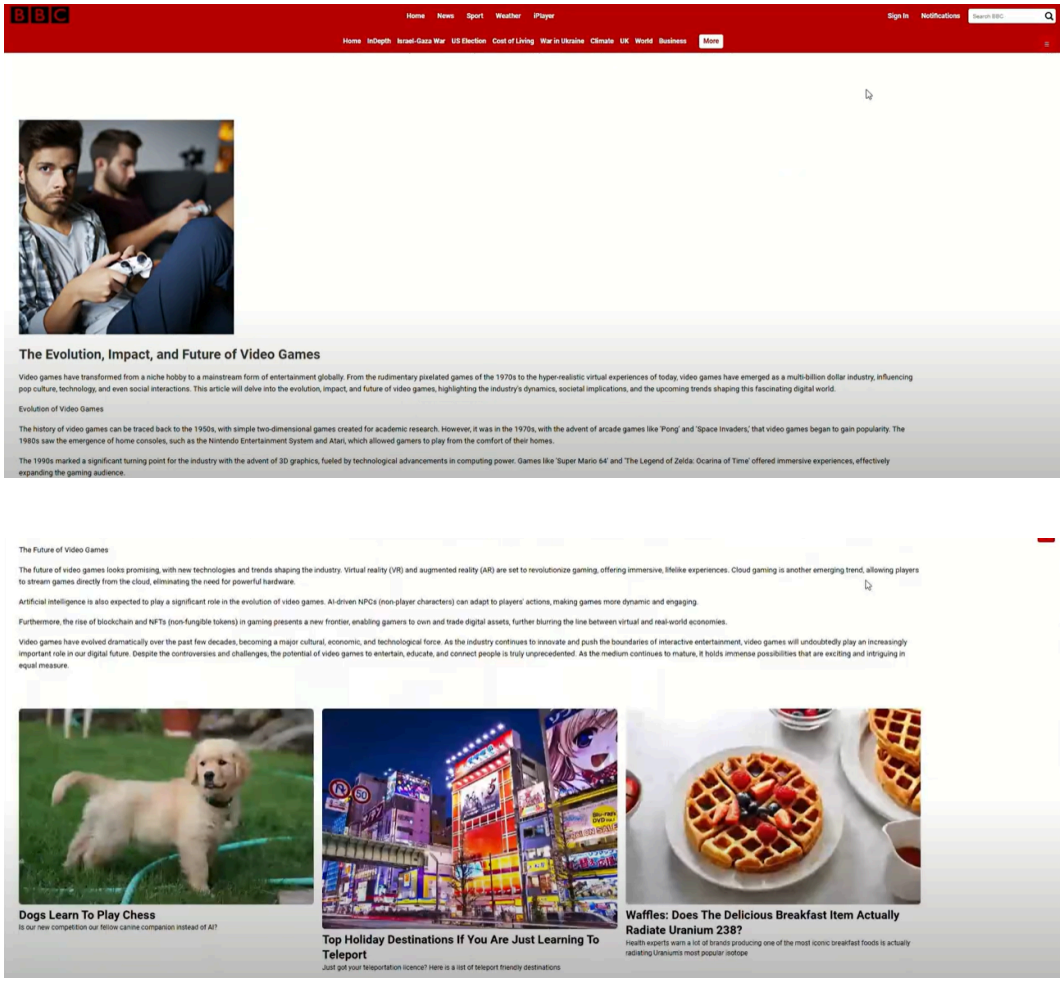

The main part of the demonstration is designed to be the most captivating: The user gets to explore a realistic article on a topic they have chosen. We generate a title for the article, a relevant image and a full article based on the user's initial prompt.

The aim here is to imitate a real BBC News article as closely as possible. If you were to show your generation to someone else, they should not be able to instantly tell that what they are looking at is AI generated. We anticipate at this point, users may try to generate multiple articles and find the giveaways that this article is not actually genuine.

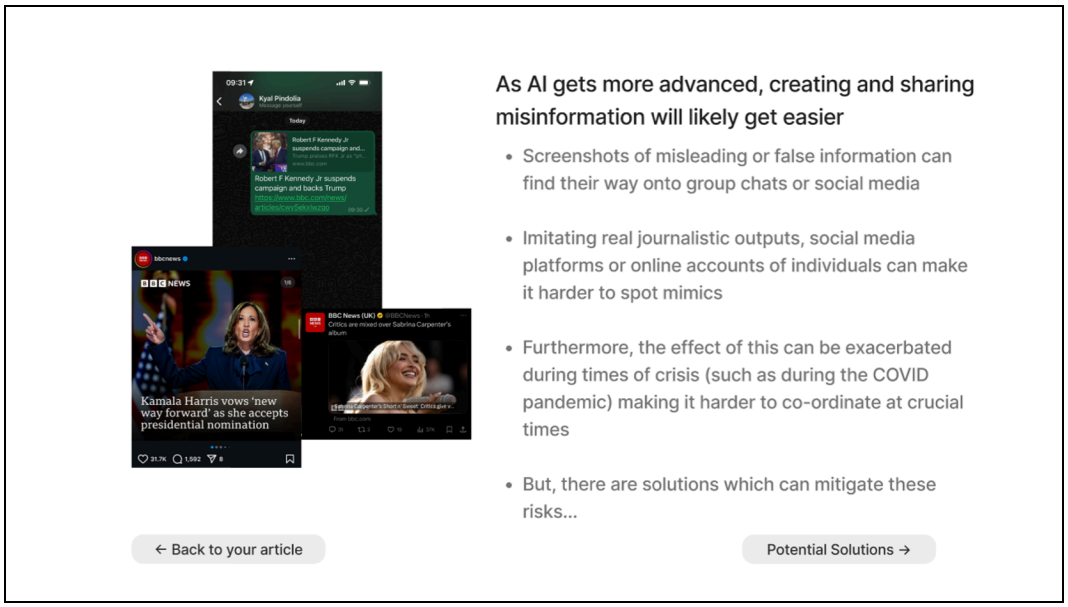

If we are successful in capturing the users attention, we can then crucially use this opportunity to explain the significance of this demo. Users may also want to share the demo if they find it particularly captivating. Once we have the user's attention, we believe there is an opportunity to educate users on what the wider implications of this illustration could be.

In this instance, we decided to outline why this could be risky. We give examples of sharing misinformation on social media, or sending screenshots of impersonation articles in direct messages and group chats.

If after using our demo, users get a *feel* for some of the capabilities of rapidly advancing AI, then we would argue our demonstration is in part successful; making it easier to spread misinformation is one example of a risk that is likely to increase as AI capabilities rapidly advance. Whilst individuals with a good understanding of current AI models may be aware of this, the broader public audience are likely yet to realise the full implications of this.

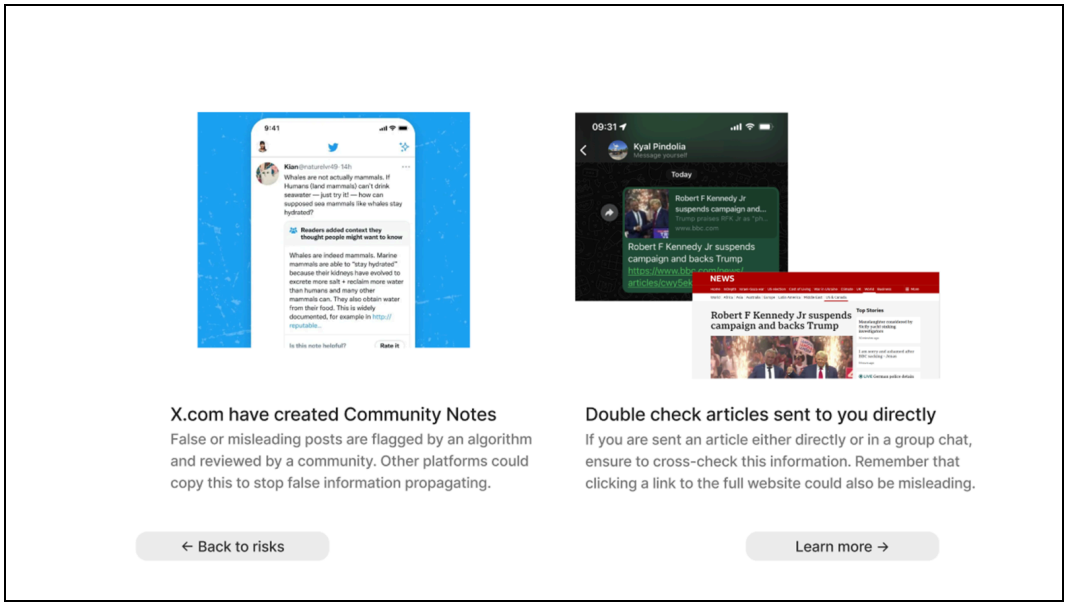

Our demo concludes with two potential solutions that can go some ways towards mitigating these risks. An effective intervention we particularly like that can help thwart the spread of misinformation through social media is Community Notes. But this becomes much harder when screenshots of imitation news articles are circulated on group chats and through direct messages. In this case, we seek to undermine the importance of cross checking information with reputable sources.

3. Discussion and Conclusion

In this demo, our aim is to show a very specific capability of current and future AI models for a broad audience. The demo is intended to be fun to use, and ideally should have a visceral impact on anyone who is familiar with BBC News articles.

If you're reading this, you're probably more informed about some of the risks of AI than the wider public. In this case, we'd invite you to imagine what it would be like to show your parents, friends or family this demo and imagine their reaction.

The idea of creating imitations of genuine news outlets more generally has potential to showcase a key *important* risk of AI. A key factor which will determine how impactful demos like this are will be how effectively the details are executed on. Examples of this would be optimising the user experience and interface when using this demo. Furthermore, using an image model that can create hyper realistic

images and a text model that effectively mimics the tone of BBC News adds to the overall "*viscerality*".

Some of the key aspects we'd want to nail before we were to make this demonstration more widely available are:

1. Making the demo very easily accessible for users. We could host this demo as a standalone website that users could easily access and share.
2. Replicating the layout of a BBC News article in a hyper-accurate way, using the same fonts, website layout etc.
3. Finding an effective way to manage costs of using different models. The cost of generating articles and images could spiral quickly, and so having getting sponsorships for compute credits by releasing this as a research demo could be valuable
4. If we do capture people's attention with a demo like this, the story we tell about associated risks is very important. We'd think more deeply about what the key risks are and highlight the most effective solutions.

Our intention is to illustrate one potential way in which AI could be misused. Imitating a specific news outlet opens the possibility of imitating other outlets too. (The Economist, Financial Times etc.)

Importantly, a limitation of our demo is that it focuses on risks of AI advancements. However, we think that it could be important to explore the opportunities this could have too. For example, BBC News is likely limited in the amount of articles they can write, and using AI provides an opportunity for increased News coverage. There are probably many more beneficial ways AI can be applied to media, and we would be excited to see more ideas developed in this space

In conclusion, demos like the one we have created are one solution addressing the problem of the wider public being under-informed about the capabilities of rapidly advancing AI. The possibilities for how AI can be applied are vast. We are all likely to be impacted in some way by these advancements in AI during our lifetimes. There is huge scope to create thought-provoking, visceral demos which can help educate people about the risks, capabilities and possibilities that come with this

novel technology.

4. Bibliography

Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., Dafoe, A., Scharre, P., Zeitzoff, T., Filar, B., Anderson, H., Roff, H., Allen, G. C., Steinhardt, J., Flynn, C., Ó hÉigeartaigh, S., Beard, S., Belfield, H., Farquhar, S., Lyle, C., Crootof, R., Evans, O., Page, M., Bryson, J., Yampolskiy, R., & Amodei, D. (2018). *The malicious use of artificial intelligence: Forecasting, prevention, and mitigation*. arXiv. https://arxiv.org/abs/1802.07228