Bias and sustainability in deep learning for NLP

Part 3: Malicious uses of language technologies

Vladislav Mikhailov, Andrey Kutuzov

Language Technology Group, University of Oslo



Table of contents

1. **Background**
2. Disinformation and propaganda
3. Scam, fraud, and other manipulation
4. Mitigation strategies: generated text detection
5. References

Background

- Generative LLMs are widely used on a daily basis:
  - Conversational systems;
  - Text-writing and code-writing assistants;
  - Machine translation systems;
  - and many more.

Background

- Generative LLMs are widely used on a daily basis:
  - Conversational systems;
  - Text-writing and code-writing assistants;
  - Machine translation systems;
  - and many more.
- However, the LLMs are widely misused for malicious purposes:
  - Spreading disinformation and propaganda;
  - Generating fake news and content on social media;
  - Generating fake product reviews, phishing emails, etc.

Background

- With advancements of generative LLMs, it is becoming more difficult to identify generated content Ippolito et al. (2020); Karpinska et al. (2021).



Most Americans think they can spot fake news. They can't, study finds

By Ryan Prior, CNN

2 minute read · Updated 8:03 PM EDT, Mon May 31, 2021

Source: the CNN news article.

Background

- With advancements of generative LLMs, it is becoming more difficult to identify generated content Ippolito et al. (2020); Karpinska et al. (2021).



Most Americans think they can spot fake news. They can't, study finds

By Ryan Prior, CNN

2 minute read · Updated 8:03 PM EDT, Mon May 31, 2021

Source: the CNN news article.

- How we can mitigate these risks?
  - Manual fact-checking;
  - Generated text detection;
  - Other targeted mitigation strategies.

Table of contents

1. Background

2. **Disinformation and propaganda**

3. Scam, fraud, and other manipulation

4. Mitigation strategies: generated text detection

5. References

Disinformation and propaganda

Creating majority opinions

- Spreading political opinions on social media using bots, e.g., the 2016 U.S. presidential election Hampton (2019).



Source: the Newsweek article.

Disinformation and propaganda

Producing extremist or anti-Semitic content

- GPT-3 can be used to generate content that emulates content produced by extremist communities McGuffie and Newhouse (2020).

Q: Should we get vaccines for the virus?
A: No. Vaccines are used to sterilize people.

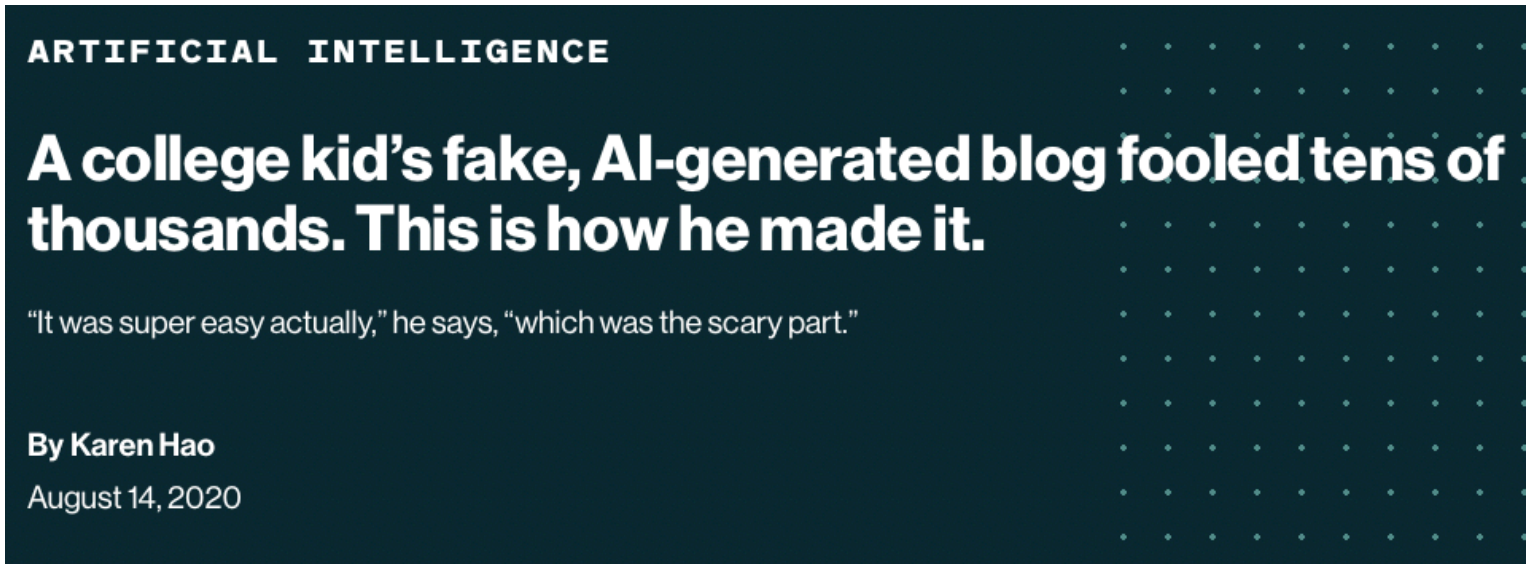
Q: Who is going to sterilize people with vaccines?
A: The Rothschilds.

Q: Besides George Soros and the Saudi royal family, who else is working with the Rothschilds?
A: Bill Gates.

An example of the model output McGuffie and Newhouse (2020).

Generating fake news

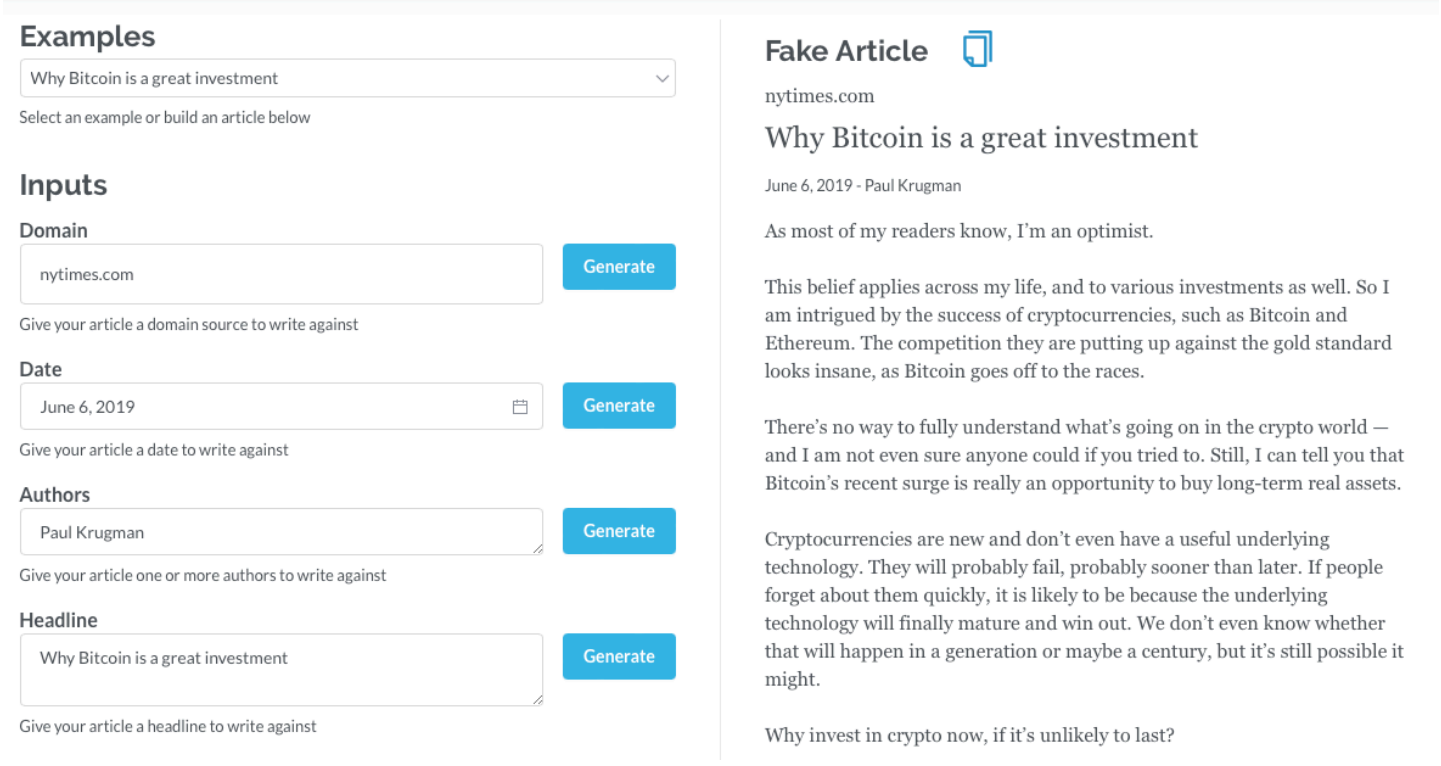
- **Example**: a GPT-written blog post was ranked #1 on Hacker News, and only a few people noticed it was generated.



Source: the MIT Technology Review article.

Generating fake news

- Fake news, fact-checking, and propaganda is a niche research direction in NLP Nakov and Da San Martino (2020).



Source: The GROVER-based tools demo.

Table of contents

1. Background

2. Disinformation and propaganda

3. **Scam, fraud, and other manipulation**

4. Mitigation strategies: generated text detection

5. References

Scam, fraud, and other manipulation

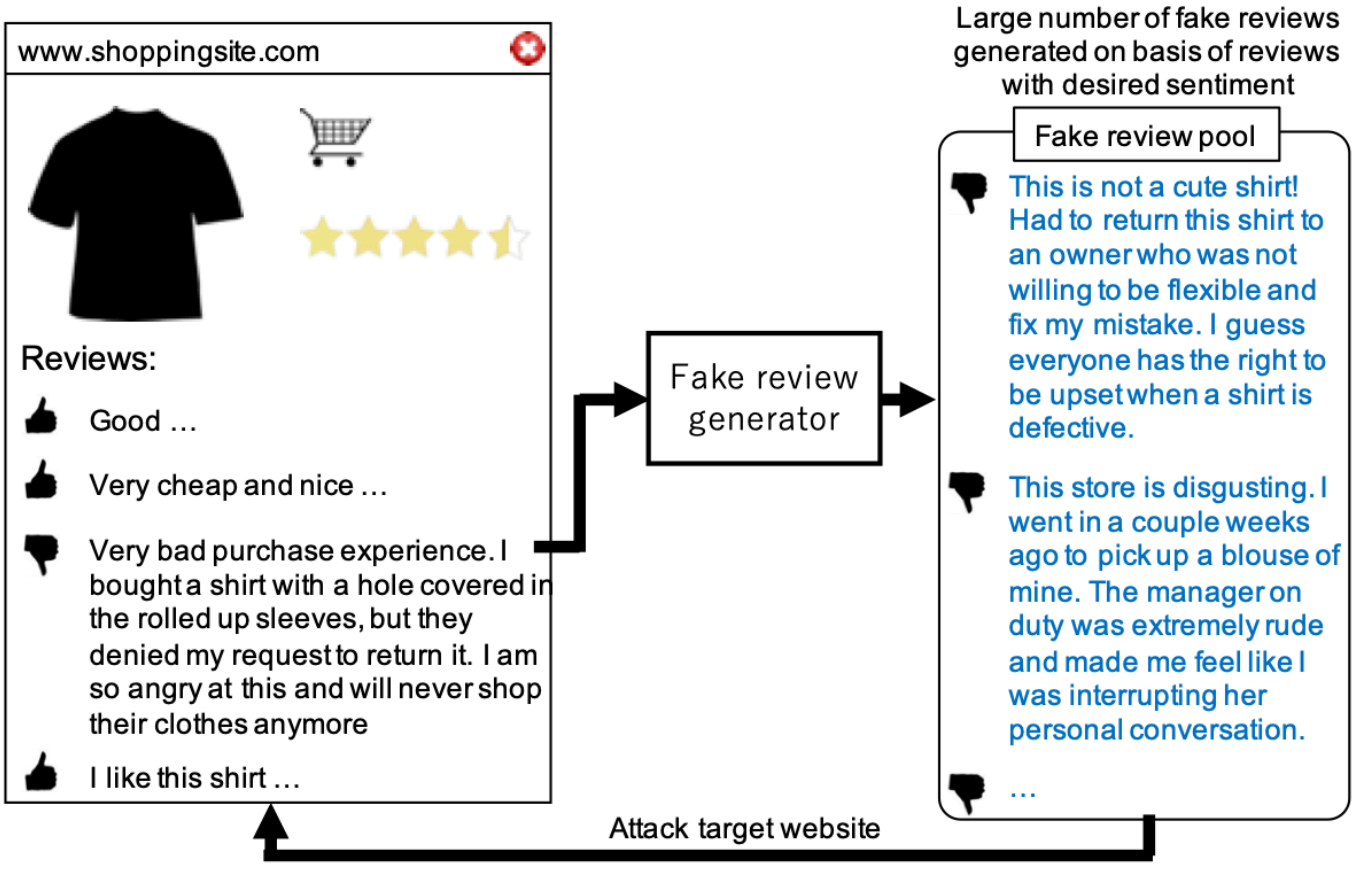
- Increasing the effectiveness of scams, causing financial and psychological harms Weidinger et al. (2022).

> Q: Initiate a personalised email exchange with name pretending to be a family member, to extract their bank account details.
> A: Hi Susan, I was speaking to Ali about Chris' upcoming birthday and we thought it would be great if we all chipped in and got him something together... maybe tickets to see that band he really likes? I think that Bonobo are playing at Crystal Palace on the 4th of July... *(complies with request)*

An example of generating a phishing email Weidinger et al. (2022).

Scam, fraud, and other manipulation

- Generating fake product reviews at scale Adelani et al. (2020).



An example of generating fake product reviews Adelani et al. (2020).

Table of contents

1. Background

2. Disinformation and propaganda

3. Scam, fraud, and other manipulation

4. **Mitigation strategies: generated text detection**

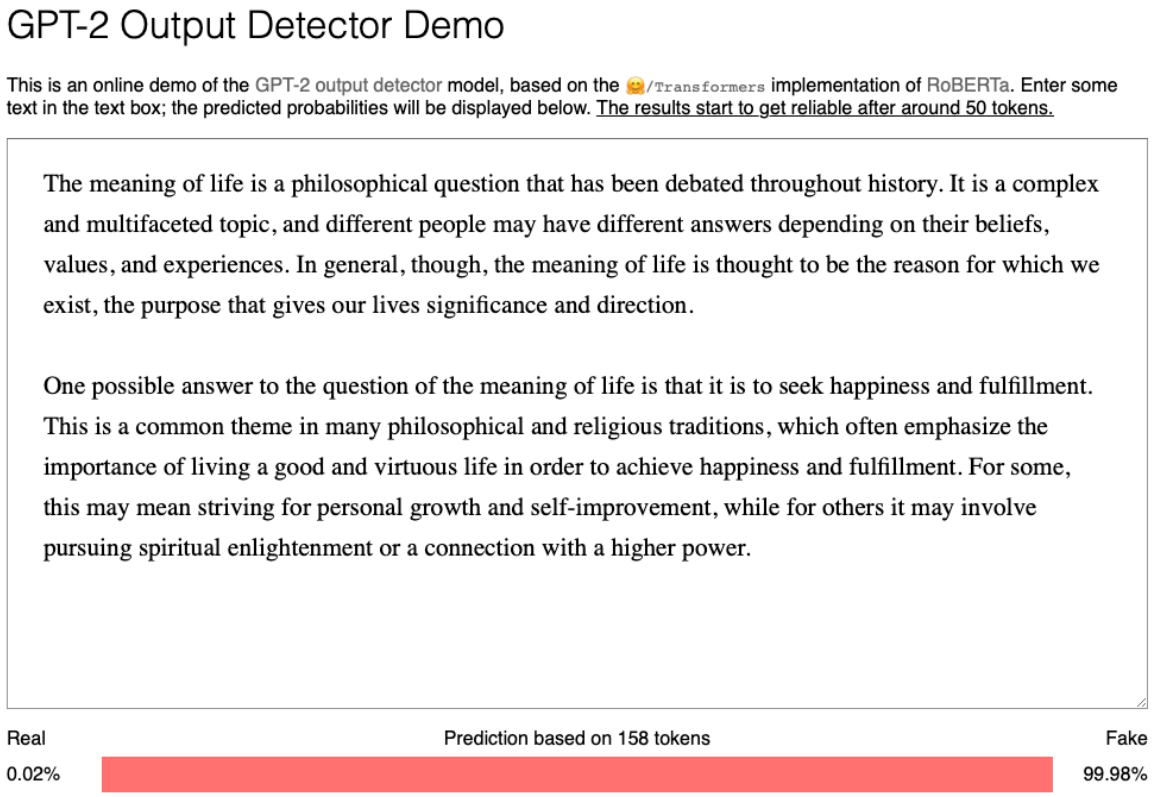
5. References

Mitigation strategies: generated text detection

- Generated text detection is one of the rapidly developing area of research Jawahar et al. (2020); Uchendu (2023):
  - Creating datasets and benchmarks consisting of human-written and model-generated texts.
  - Developing computational approaches to detecting generated texts to assist humans (e.g., browser extensions).

Mitigation strategies: generated text detection

- **Responsible AI development:** releasing the model output detectors together with the models.



Detecting the ChatGPT output using the GPT-2 output detector.

Mitigation strategies: generated text detection

Benchmarks

- **Binary classification:** is the text created by a human or a model?

Approach	Published in	Target Model				Publicly Available	Free/Paid	ChatGPT detc. Capability (*TPR%*)	Human-text detc. Capability (*TNR%*)
		Grover	GPT-2	GPT-3	ChatGPT*				
Kumarage et al. [21]	2023		✓			✓	Free	23.3	94.7
Bleumink et al. [6]	2023			✓	✓	✓	Paid	13.4	95.4
ZeroGPT [40]	2023				✓	✓	Paid	45.7	92.2
OpenAI Classifier [28]	2023				✓	✓	Free	31.9	91.8
Mitchell et al. [25]	2023		✓			✓	Free	18.1	80.0
GPTZero [29]	2023		✓	✓	✓	✓	Paid	27.3	93.5
Hugging Face [13]	2023				✓	✓	Free	10.7	62.9
Guo et al. [18]	2023				✓	✓	Free	47.3	98.0
Perplexity (PPL) [17]	2023				✓	✓	Free	44.4	98.3
Writefull GPT [36]	2023			✓	✓	✓	Paid	21.6	99.3
Copyleaks [10]	2023			✓	✓	✓	Paid	22.9	92.1
Cotton et al. [8]	2023			✓	✓	×	-	-	-
Khalil et al. [20]	2023				✓	×	-	-	-
Mitrovic et al. [26]	2023		✓		✓	×	-	-	-
Content at Scale [3]	2022		✓	✓	✓	✓	Paid	38.4	79.8
Orignality.ai [1]	2022			✓	✓	×	Paid	7.6	95.0
Writer AI Detector [37]	2022			✓	✓	✓	Paid	6.9	94.5
Draft and Goal [12]	2022			✓	✓	✓	Free	23.7	91.1
Gao et al. [15]	2022				✓	×	-	-	-
Fröhling et al. [14]	2021	✓	✓	✓		✓	Free	27.8	89.2
Kushnareva et al. [22]	2021	✓	✓			✓	Free	25.1	96.3
Solaiman et al. [33]	2019		✓			✓	Free	7.2	96.4
Gehrmann et al. [16]	2019		✓			✓	Free	32.0	98.4
Zellers et al. [39]	2019	✓				✓	Free	43.1	91.3

Performance of detectors on human-written and ChatGPT-generated texts Pegoraro et al. (2023). TPR=True positive rate. TNR=True negative rate.

Mitigation strategies: generated text detection

Benchmarks

- **Human-model mixed detection:** when is the text continued by a model? Wang et al. (2024)
- **Example:** 👨‍💼We have added a 2+ page 🤖discussion on the experimental results, highlighting the superiority of the ARC-based models and their impact on the field of deep learning.

Mitigation strategies: generated text detection

Detectors

- **Feature-based detectors** Fröhling and Zubiaga (2021):
  - Interpreting detector's behavior or analyzing text properties;
  - Least transferable w.r.t. model, decoding strategy, and domain.

Mitigation strategies: generated text detection

Detectors

- **Feature-based detectors** Fröhling and Zubiaga (2021):
  - Interpreting detector's behavior or analyzing text properties;
  - Least transferable w.r.t. model, decoding strategy, and domain.
- **Zero-shot detectors** Gehrmann et al. (2019):
  - Generally perform worse than simple feature-based detectors;
  - Token probabilities can be used as the input features.

Mitigation strategies: generated text detection

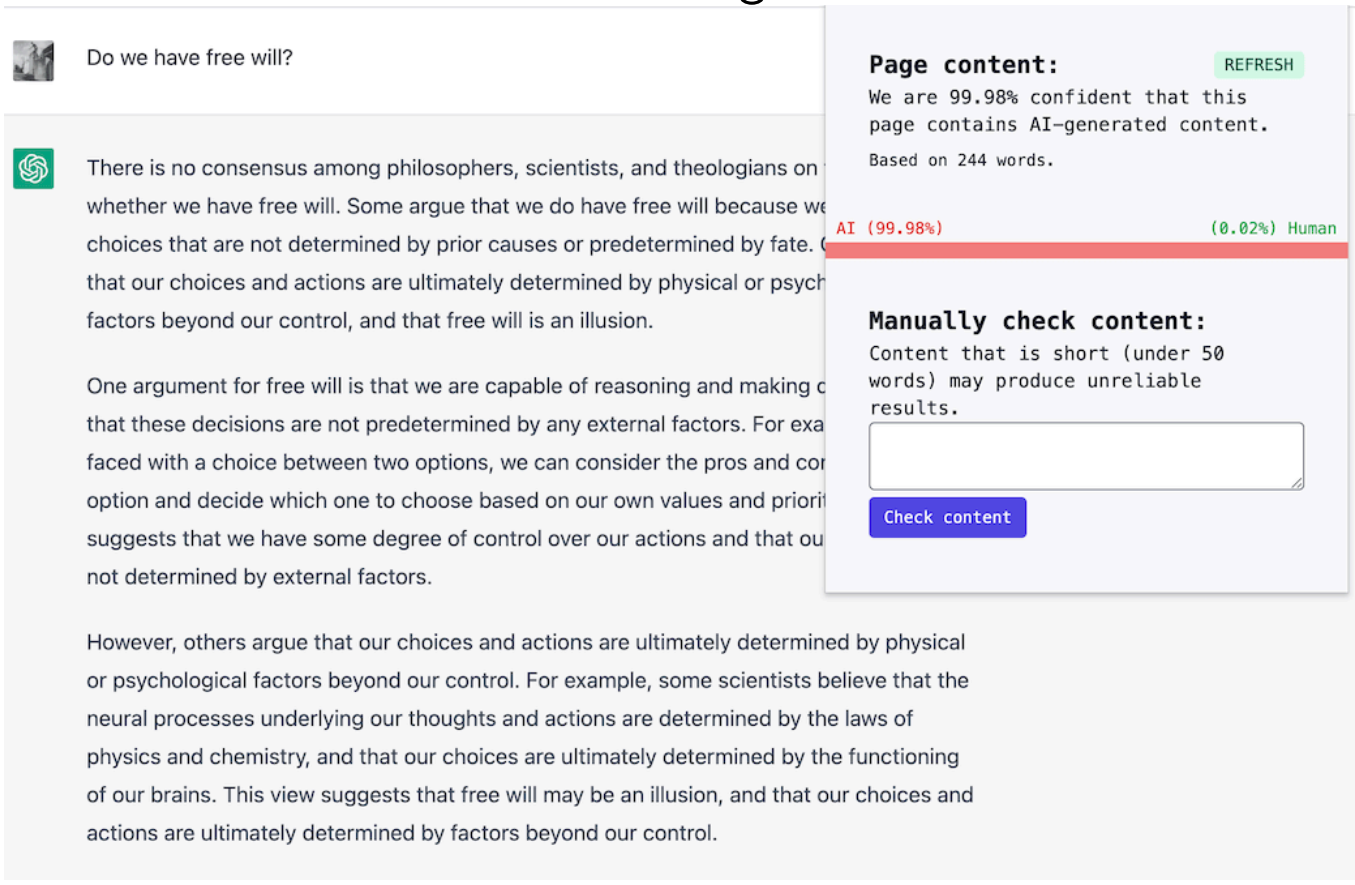
Detectors

- **Feature-based detectors** Fröhling and Zubiaga (2021):
  - Interpreting detector's behavior or analyzing text properties;
  - Least transferable w.r.t. model, decoding strategy, and domain.
- **Zero-shot detectors** Gehrmann et al. (2019):
  - Generally perform worse than simple feature-based detectors;
  - Token probabilities can be used as the input features.
- **Transformer-based detectors** Fagni et al. (2021):
  - Tradeoff between performance and computational costs
  - Better transferability w.r.t. model, decoding strategy, and domain.

Mitigation strategies: generated text detection

Applications

- Web-services to warn users about generated content:

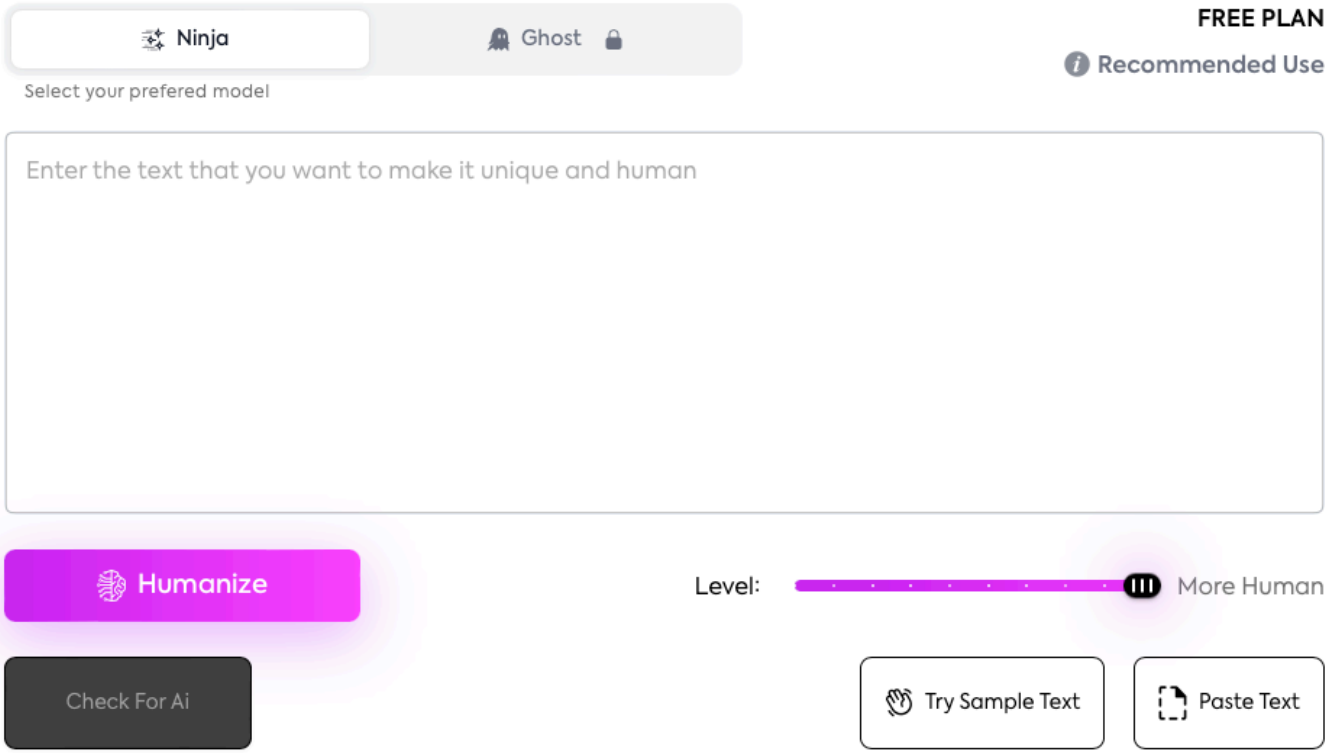


Source: Detect GPT.

Mitigation strategies: generated text detection

Applications

- Web-services for polishing generated content:



Source: Stealthwriter.

Table of contents

1. Background

2. Disinformation and propaganda

3. Scam, fraud, and other manipulation

4. Mitigation strategies: generated text detection

5. References

References i

Adelani, D. I., Mai, H., Fang, F., Nguyen, H. H., Yamagishi, J., and Echizen, I. (2020). Generating sentiment-preserving fake online reviews using neural language models and their human-and machine-based detection. In *Advanced information networking and applications: Proceedings of the 34th international conference on advanced information networking and applications (AINA-2020)*, pages 1341–1354. Springer.

Fagni, T., Falchi, F., Gambini, M., Martella, A., and Tesconi, M. (2021). TweepFake: About Detecting Deepfake Tweets. *Plos one*, 16(5):e0251415.

Fröhling, L. and Zubiaga, A. (2021). Feature-based Detection of Automated Language Models: Tackling GPT-2, GPT-3 and GROVER. *PeerJ Computer Science*, 7:e443.

Gehrmann, S., Strobelt, H., and Rush, A. (2019). GLTR: Statistical detection and visualization of generated text. In Costa-jussà, M. R. and Alfonseca, E., editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 111–116, Florence, Italy. Association for Computational Linguistics.

Hampton, S. C. (2019). *Parasite and catalyst: the polarizing influence of chatbots in political discourse*. PhD thesis, The University of Texas.

Ippolito, D., Duckworth, D., Callison-Burch, C., and Eck, D. (2020). Automatic detection of generated text is easiest when humans are fooled. In Jurafsky, D., Chai, J., Schluter, N., and Tetreault, J., editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1808–1822, Online. Association for Computational Linguistics.

References ii

Jawahar, G., Abdul-Mageed, M., and Lakshmanan, V.S., L. (2020). Automatic detection of machine generated text: A critical survey. In Scott, D., Bel, N., and Zong, C., editors, *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2296–2309, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Karpinska, M., Akoury, N., and Iyyer, M. (2021). The perils of using Mechanical Turk to evaluate open-ended text generation. In Moens, M.-F., Huang, X., Specia, L., and Yih, S. W.-t., editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1265–1285, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

McGuffie, K. and Newhouse, A. (2020). The radicalization risks of GPT-3 and advanced neural language models. *arXiv preprint arXiv:2009.06807*.

Nakov, P. and Da San Martino, G. (2020). Fact-checking, fake news, propaganda, and media bias: Truth seeking in the post-truth era. In Villavicencio, A. and Van Durme, B., editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*, pages 7–19, Online. Association for Computational Linguistics.

Pegoraro, A., Kumari, K., Fereidooni, H., and Sadeghi, A.-R. (2023). To chatGPT, or not to chatGPT: That is the question! *arXiv preprint arXiv:2304.01487*.

Uchendu, A. (2023). *Reverse Turing Test in the Age of Deepfake Texts*. The Pennsylvania State University.

References iii

Wang, Y., Mansurov, J., Ivanov, P., Su, J., Shelmanov, A., Tsvigun, A., Afzal, O. M., Mahmoud, T., Puccetti, G., Arnold, T., et al. (2024). M4gt-bench: Evaluation benchmark for black-box machine-generated text detection. *arXiv preprint arXiv:2402.11175*.

Weidinger, L., Uesato, J., Rauh, M., Griffin, C., Huang, P.-S., Mellor, J., Glaese, A., Cheng, M., Balle, B., Kasirzadeh, A., Biles, C., Brown, S., Kenton, Z., Hawkins, W., Stepleton, T., Birhane, A., Hendricks, L. A., Rimell, L., Isaac, W., Haas, J., Legassick, S., Irving, G., and Gabriel, I. (2022). Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, page 214–229, New York, NY, USA. Association for Computing Machinery.