

3 Issues with nonfake news generation

Radar and Monok are benefiting from the coming of age of language generation and summarisation technologies, and utilise them to address specific needs in the news production sector, in the former case enabling regional newspapers to survive in an age of falling ad revenues, and in the latter case providing a cost-effective way to create high-quality factual content that does well in SEO terms. The underlying methods start from structured data or existing news articles, and are not easily adapted for fake news generation.

This doesn't mean that nonfake news generation doesn't come with its own risks. There is potential for intentional misuse: automatically generated news (AGN) can be arbitrarily selective in content, slanted to further specific agendas, and shot off to any number of destinations at the touch of a button, potentially overwhelming news output, and drowning out more balanced accounts. As with fake news, automatic detection would offer some leverage (based on spread patterns and text properties, in the absence of fake content that can be fact-checked). Unfortunately GLTR-like automatic detection doesn't work for summarisation and template generation, because these incorporate (paraphrased) word strings written by humans. Figure 3 shows GLTR results for a Monok article deemed slightly more likely to be human-authored than the similar, actually human-authored, article above it (more red and purple, more human).

Humans and search engines don't do much better. Both Monok's and Radar's articles are published without human postprocessing (although in the latter case, editors add local information), indicating a level of quality that readers find hard to distinguish from human-written. Monok escapes all Google penalties for automatically generated and non-unique content, meaning it passes for human-written and new, despite being neither.

Even with bona fide use, bias and misrepresentation can creep into generated articles from source data, and automatic selection of input risks omission of contextualising information resulting in misrepresentations or even untrue entailments.

But perhaps the elephant in the room is transparency. Automation has some hand in producing most of the news we read now, but you would never know it from attributions. Newspapers are running full AGN under human bylines, some editors believing readers "couldn't care less who

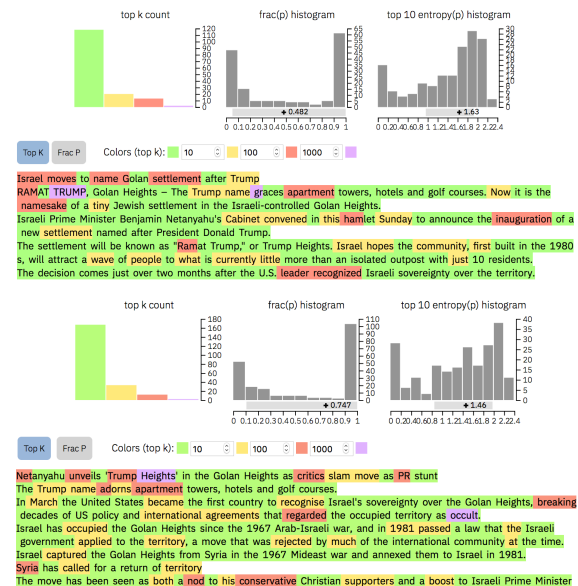


Figure 3: GLTR-2 verdict on AGN (bottom, Monok) vs. human-authored news (top, Fox News).

wrote it" (Nilsson, 2019). Even automation pioneer AP has nothing on AGN, its use or attribution, in the 2019 edition of its iconic AP Stylebook, just a chapter on (human-authored) data journalism.

The industry still assumes "no one wants AI to fully replace humans in the newsroom, and no one is actively working towards that end" (Chiusi and Beckett, 2019), but start-ups like Monok and Radar show that full AGN is entering general news production. Developing tools for detecting AGN, addressing data biases, and contextualising reported facts can mitigate against some of the adherent risks. More importantly, journalism needs to develop guidance and best practice to ensure the transparent, ethical and responsible use of AGN, if further erosion of concepts like truth, authorship, and responsibility for content is to be avoided.

References

- F. Chiusi and C. Beckett. 2019. [The journalism AI global survey: What we've learned so far](#).
- S. Gehrmann, H. Strobelt, and A. M. Rush. 2019. GLTR: Statistical detection and visualization of generated text. *ArXiv*, abs/1906.04043.
- P. Nilsson. 2019. [Robo-reporter writes front-page news](#). *Financial Times*, 4 June 2019.
- A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8).
- R. Zellers, A. Holtzman, H. Rashkin, Y. Bisk, A. Farhadi, F. Roesner, and Y. Choi. 2019. Defending against neural fake news. *ArXiv*, abs/1905.12616.