Write your own words

Against AI for writing

Willem L. Middelkoop
Nov. 11, 2024

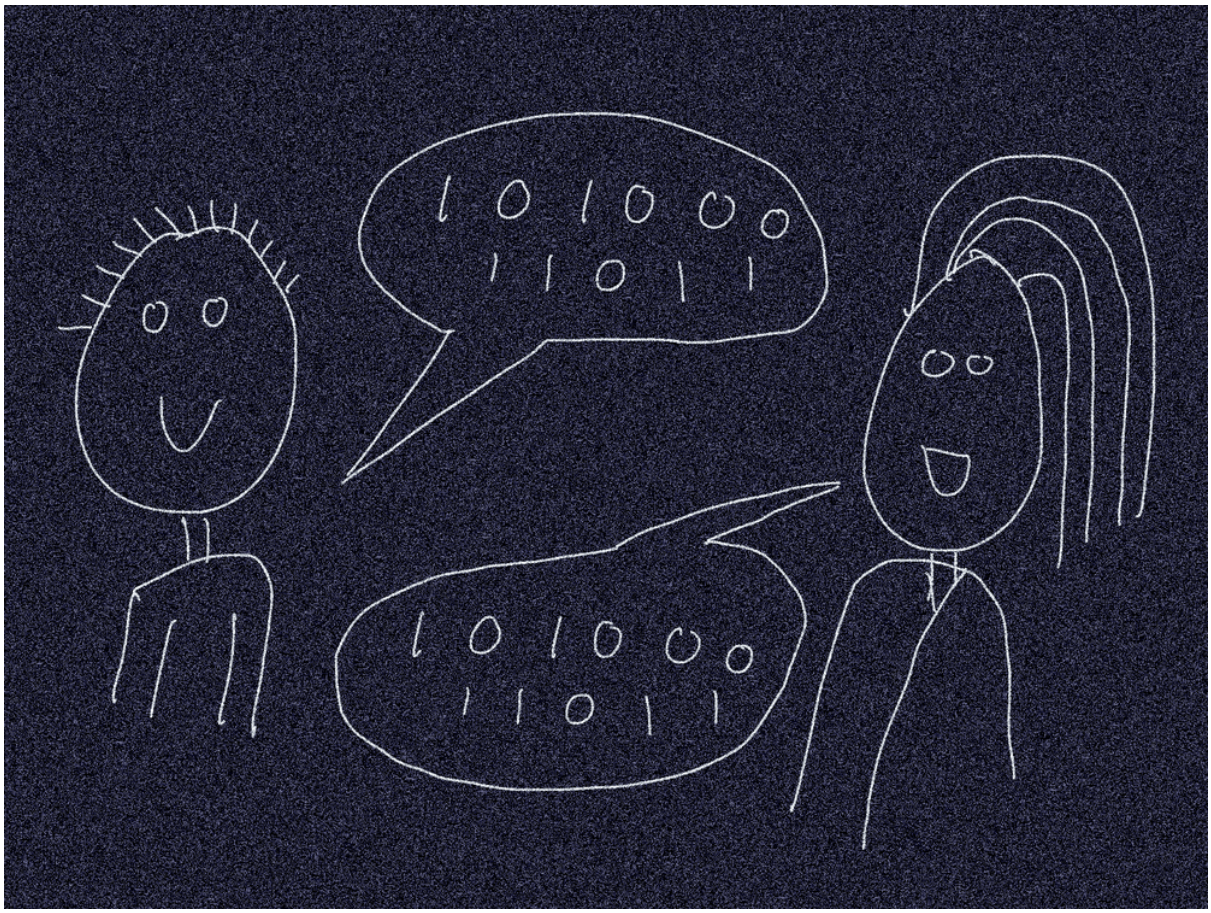

Two years ago ChatGPT launched, it popularised AI as tool to create things. Large language models generate text based on prompts, outputting words that are hard to distinguish from being written by a person. The short term benefits are clear, but I suspect many people underestimate the long term costs.

Ever since OpenAI launched ChatGPT, I have been actively working with the new technology to understand its capabilities. I have programmed with AI, I learned with AI and I taught people to work with AI. I went back to the University to brush up my skills for technologies underlying AI. I wrote a real looking blog post with AI. But moreover, I have been thinking about the impact of AI on our society.

It is a powerful tool, capable of doing amazing things like detecting cancer from MRI images or structuring and documenting computer code for reusability. Everybody can use AI, and therein lies a potential problem, specifically for content creation.

Uniquely human

If we all use the same tools to *generate* content, everything will start to look the same eventually. We risk losing diversity. Especially as new computer models are being trained on newly generated content. It is a recursive loop causing new content to appear more similar to existing content.

You can already see this problem if you look at Google search results. People have noticed that results are getting worse. This is often attributed to an increasing amount of search-engine-optimised but low-quality content. You should check out this paper by Bevendorff Et Al. for a longitudinal investigation of SEO spam in search engines. Low-quality, but *good looking* content is pushing the actual valuable content away.

What can we do about it? For sure the folks that make the computer models will come up with intelligent filters designed to exclude low-quality material used for training. **I expect that there will be a premium on content that is not generated by AI**. Just like how old low-background steel is valued because it is not contaminated with traces of nuclear fallout present in modern steel.

Steel in old shipwrecks is valued for its purity - HMS Royal Oak Seabed (Scapa Flow)

Detecting AI Content

Did AI write this? How differently does ChatGPT write when compared to humans? Researchers from Preply looked for signs to detect AI written texts to determine the origin of student essays. They analysed 12,000 essays and created an analyses of the most frequent words, phrases and errors in both human-written and computer-generated works. These are some of their key takeaways:

- Human-written texts were far more likely to feature mistakes, with 78% of texts containing at least one error, compared to only 13% of AI essays
- The most common punctuation errors in human texts included missing commas, hyphen errors and missing spaces
- ChatGPT uses 35% less unique vocabulary than humans, with only 7,308 unique words overall, compared to 11,248 unique words from human-written texts

- Humans are more likely to use common nouns and verbs like 'one' or 'people', while ChatGPT favours complex and specialised terms like 'cultural' or 'economic'.

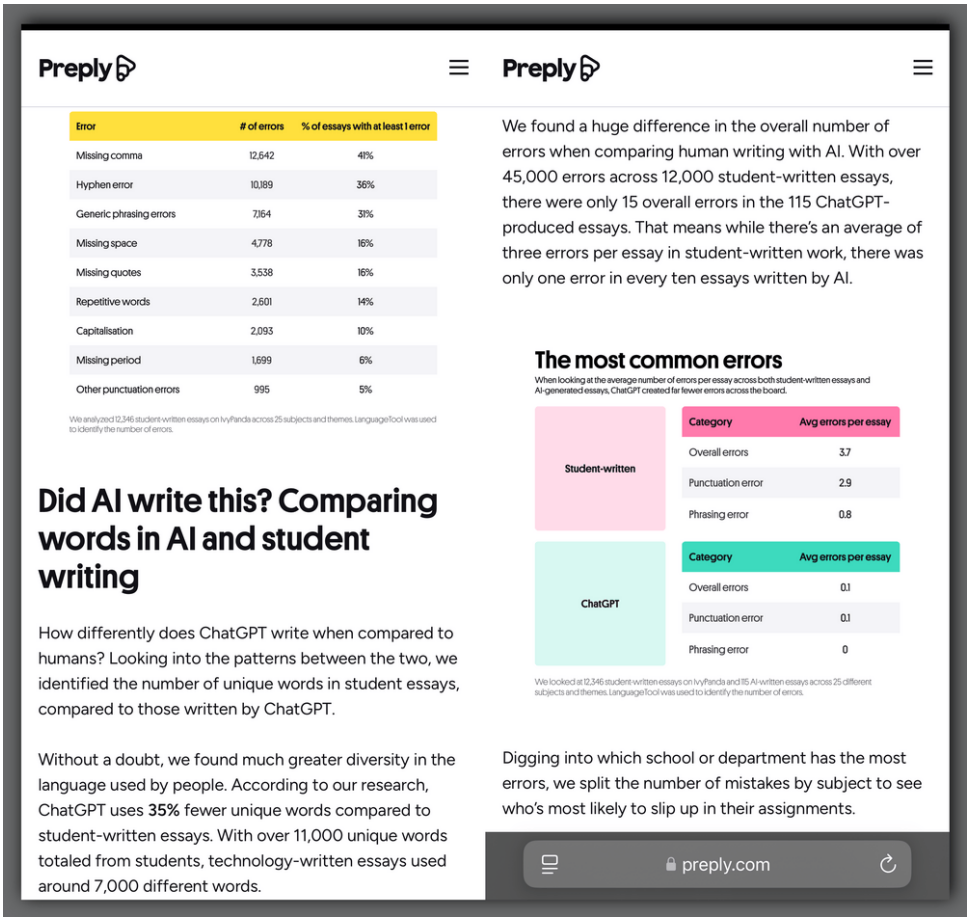

Research findings by Preply: Distinguishing AI-written texts from Human-written texts

Art instead of Artificial

If you care about content you should consider seeing it as art, there can be a human quality to it that determines its meaning for other humans. In a detailed comparison Gold Penguin compares human-created art with AI-generated art. Their key findings in spotting AI-created content are:

- **Overuse of patterns:** Machines don't do well with entropy or random data, making AI generators more likely to repeat certain patterns that are prominent in other images
- **Complexity**: AI struggles to replicate the unique or unusual elements that human artists are more likely to experiment with.
- **Lack of emotion**: While they could sometimes be visually better, AI art lacks the emotional and conceptual depth of human art.
- **Mistakes**: Have you ever seen a piece of art and you just know it's fake even though you can't put a finger on why it is? That's the uncanny valley. It usually manifests in small instances rather than something right in your face. Six fingers, bleeding pupils, improper holding of objects — the list goes on.

Spot the real Monet (the Artist's garden at Vétheuil) - you're probably able to easily find it, despite both images being pretty

Spot the real Van Gogh - an easy task if you look at the repeating patterns

Now I understand that not all texts are equally important and some may not be worth it to be elevated as pieces of art. There are very good usecases for AI-generators, including translations or summaries. I am not against large language models in general, I only argue that they should be seen as a tool with accommodating pluses and minuses.

Conclusion

The more powerful a tool is, the greater the harm it *can do* - it is up to the person handling the tool to use it well. While large language models are powerful in their capabilities to generate good looking texts, its value is ultimately determined by the meaning humans attach to it. If museums full of art are any indication, I expect there will always be appreciation for things created by hand. We should start by simply caring a little more about the things we create, by writing our own words for example.