



MMC/Mammoth-2025 Problem

Plague 2.0

Nowadays, more and more texts, videos, music, and other content on the internet is generated using various artificial intelligence (AI) systems, including large language models (LLMs). A certain portion of the generated material is the result of so-called AI “hallucinations”, i.e. it contains false information presented in a serious tone and with a claim to authority. The online community has dubbed such materials “AI slop”. The reasons for the spread of this “slop” are clear: sufficiently long attention of the audience to materials created without any human effort generates “views” of advertisements associated with this content, and the content authors receive advertising payments.

However, this phenomenon has the flip side: AI developers do not stop, they create new, larger AI models, which need more and more information for training. But in practice, there is only one source of data large enough to satisfy the growing needs of AI training: the internet content. This means that new AI models can find “significant” materials (where “significance” is measured by the larger number of “views”), which are, in fact, “AI slop” themselves, consider them to be true by definition, and use them for their training. As a result, new AI models may stray further and further from the facts of the real world.

What are the prospects of this process? How long will it take for the internet to be “flooded” with “AI slop”? What factors could influence this process?

Tasks

- Using open sources and statistical data, estimate:
 - The rate at which new AI models are contaminated by false facts (“factoids”) from the internet: how long might it take from a factoid appearing in online materials to its inclusion into widely-used AI models?
 - What proportion of “popular” online publications is generated by AI systems, and what proportion of those is likely to be false information resulting from AI “hallucinations”? Please explain which materials you consider to be “popular”.
- Assuming that the proportion of false facts in the total amount of internet information is initially small and remains relatively small for a certain period of time, propose a formula that approximately describes the proportion of false facts as a function of time (within this “certain period”). How long will it take to reach a proportion of 10% and 20% if the process starts at 0.1%?
- Using the findings from tasks 1 and 2, build a mathematical model that predicts the dynamics of the proportion of false facts in online information regardless of the proportion being small or not. Using the model you have built, predict how the situation will develop: if the behaviour of internet users and creators of new AI systems does not change, will the internet eventually become filled with false information almost to 100%, or will the proportion of false information stabilise at some lower value? How long might this take?

4. Do you expect any change in the behaviour of internet users and creators of AI systems in response to the obviously high level of “AI slop” on the internet? How will your model's predictions change if these behavioural changes are incorporated into it? Predict the actual behaviour of the proportion of false information on the internet over the next few decades.

It should be noted that, obviously, there are other sources of information for training AI systems besides online materials, such as scientific publications (articles) or fiction literature. However, these areas are now also heavily influenced by AI-generated materials. While the authors of materials in these fields are guided by considerations that are largely different from those of the authors of content intended specifically for publication on the internet, the end result is qualitatively similar. In addition, the total amount of information in scientific and fiction literature is much smaller than that of the “entire internet”.

Paper formatting requirements

- A. The paper should explicitly identify the sections or passages that contain the answers to tasks 1–4. For example, you may write “(Task 3)” at the end of the title of the section with an answer to this task.
- B. The paper should be no more than 15 A4 pages, in 12 pt font, line spacing at least 1.5 and margins of 2 cm on all sides. The reference list and appendices are not included into the page count; however, they must not contain information essential for understanding and evaluation of your paper. Please review the full requirements in the Mammoth-2025 Instructions, which you will receive along with this problem (a link to the instructions can also be found on the MMT-2025 main webpage)!
- C. In the appendix at the end of the paper, there should be a separate section titled “Use of Artificial Intelligence (AI)”. It should state whether you have used AI in your research or writing of the paper. The use of AI is not encouraged, but it is not prohibited. If AI was used, this section must include the following information:
 - what AI tools were used and for what purpose;
 - what specific prompts were used with the AI;
 - how the results of the AI were checked (because AI can “hallucinate”).You can also include a hyperlink to the chat created in GigaChat, ChatGPT, Google Gemini, etc.