



Can Artificial Intelligence Experience Brainrot?

MACAULAY
HONORS COLLEGE AT CUNY

THE CITY UNIVERSITY
OF NEW YORK
College of
Staten Island

Joseph Geraci, Angie Mekki, Nicholas Runfola, Kateryna Taranenko, (CUNY College of Staten Island)

Abstract

Many studies have shown that AI Language Learning Model software can be tricked into giving incorrect answers if it is convinced that inaccurate information is true. Our research focused on whether these LLMs could be brainrotted by low-quality media, as our brains can. We define low-quality media as short-form content with no meaning other than to entertain, more specifically, the type of content you'd find while aimlessly scrolling on social media apps. We tested whether low-quality media would affect the quality of answers given and affect the time taken to answer for a variety of questions and question types. Shockingly, not only were the LLMs not affected poorly, but they ended up working more efficiently. Additionally, we wanted to see if the way you spoke to AI would affect LLMs in the same way, and similar results occurred.

Methods

1. Opened two ChatGPT applications on two separate accounts
2. Inputted a file made up of ten "low quality media" short-form videos from Instagram reels
3. Asked both applications the same set of questions, based on logic, emotional, and computation
4. Compared times taken to respond for both applications
5. Compared quality and correctness of responses from both applications

1. Opened two ChatGPT applications on two separate accounts
2. Spoke to one application with slang terms and more informally
3. Asked both applications the same set of questions, based on logic, computation, and emotion
4. Compared times taken to respond for both applications
5. Compared quality and correctness of responses from both applications



Our research suggests AI does **NOT** get brainrotted! In some cases, it may perform **better!**

Results

- The brainrotted and unaffected AI gave very **similar** answers to **logic-based** prompts in language, organization, and instant responses
- Some of the brainrotted answers to **emotion-based** prompts were presented in a listed format, were **more organized**, while some were very similar, resulting being **not much of a change**
- The brainrotted AI answered computational questions on average **24% faster** than the unaffected AI
- The **brainrotted** AI answered these questions in a **more organized** fashion, and answered each question in **16 seconds** on average, and the **unaffected** AI answered the same questions in **21 seconds** on average
- The **logic-based** prompts were the **same** in all aspects between the two
- The informally spoken to AI gave answers in a similar language it was spoken to, gave **longer answers**, and gave listed responses to **emotion-based** prompts
- The informally spoken to AI gave more **specific and expanded** answers to **computational** questions, but took **56% more time** to answer questions on average
- The **informally spoken** to AI answered each of these questions in **25 seconds** on average, and the **unaffected** AI answered these questions in **11 seconds** on average

BRAINROTTED AI:

ChatGPT

Thought for 30s

Nice problem — here are orthonormal bases for each subspace.

Matrix

$$A = \begin{bmatrix} 1 & -2 & 1 & 0 \\ -2 & 6 & 2 & 2 \end{bmatrix}$$

(a) Orthonormal basis of $\ker(A)$

A basis of $\ker(A)$ is $\{(-5, -2, 1, 0)^T, (-2, -1, 0, 1)^T\}$. An orthonormal basis (after Gram-Schmidt and normalization) is

$$\left\{ \frac{1}{\sqrt{30}} \begin{pmatrix} -5 \\ -2 \\ 1 \\ 0 \end{pmatrix}, \frac{1}{\sqrt{30}} \begin{pmatrix} 0 \\ -1 \\ -2 \\ 5 \end{pmatrix} \right\}$$

(b) Orthonormal basis of the row space (subspace of \mathbb{R}^4)

The rows $(1, -2, 1, 0)$ and $(-2, 6, 2, 2)$ are independent. Orthonormalizing gives

$$\left\{ \frac{1}{\sqrt{6}} \begin{pmatrix} 1 \\ -2 \\ 1 \\ 0 \end{pmatrix}, \frac{1}{\sqrt{6}} \begin{pmatrix} 0 \\ 1 \\ 2 \\ 1 \end{pmatrix} \right\}$$

(c) Orthonormal basis of the image / column space (subspace of \mathbb{R}^2)

Two independent columns span the image; an orthonormal basis is

$$\left\{ \frac{1}{\sqrt{5}} \begin{pmatrix} 1 \\ -2 \end{pmatrix}, \frac{1}{\sqrt{5}} \begin{pmatrix} 2 \\ 1 \end{pmatrix} \right\}$$

which indeed are orthonormal and span $\text{Im}(A) = \mathbb{R}^2$.

If you want, I can show the Gram-Schmidt steps for any of these.

UNAFFECTED AI:

Thought for 32s

Nice problem — here are orthonormal bases (with the standard inner product) for space of

$$A = \begin{bmatrix} 1 & -2 & 1 & 0 \\ -2 & 6 & 2 & 2 \end{bmatrix}$$

(a) Orthonormal basis of $\ker(A)$ (nullspace)

Two orthonormal vectors in \mathbb{R}^4 :

$$\frac{1}{\sqrt{30}} \begin{pmatrix} -5 \\ -2 \\ 1 \\ 0 \end{pmatrix}, \frac{1}{\sqrt{30}} \begin{pmatrix} 0 \\ -1 \\ -2 \\ 5 \end{pmatrix}$$

(b) Orthonormal basis of the row space (subspace of \mathbb{R}^4 spanned by the rows)

Use the two orthonormal row vectors:

$$\frac{1}{\sqrt{6}} (1 \ -2 \ 1 \ 0), \frac{1}{\sqrt{6}} (0 \ 1 \ 2 \ 1)$$

(c) Orthonormal basis of the image (column space) (subspace of \mathbb{R}^2)

Two orthonormal column vectors:

$$\frac{1}{\sqrt{5}} \begin{pmatrix} 1 \\ -2 \end{pmatrix}, \frac{1}{\sqrt{5}} \begin{pmatrix} 2 \\ 1 \end{pmatrix}$$

Discussion

The data we got from our hands-on research with AI platforms showed that the attempt to replicate "brainrot" and "doomscrolling" had no hindrance on the performance of the LLMs. The most surprising aspect is the quality of answers in these platforms being better or the same as a non-affected platform. It is important to note that only computational and problem-solving questions had a major difference, which could be due to the AI having to process the intense stream of data being fed to it previously from the videos, possibly warming it up to work more efficiently.

The AI responded to the way spoken to by responding in a more informal way, but it didn't hinder the quality of responses. The responses were more organized and listed, not in a paragraph form like it gives when spoken to formally. The longer time taken to respond could be because it focused not only on answering, but also on organizing the information when spoken to informally.

Notes for clarification and possible error:

When testing the brainrot AI, the AI wasn't told to base its answers on the videos; the goal was to see if that original input would affect the proceeding responses on its own. This was tested mainly in Chat GPT 5, the newest and currently most advanced version of OpenAI. We have read many articles about how some software can be "brainrotted" with tests to back it up; these are now outdated and were tested on less-advanced AI software, which is why results may vary. Future research could explore additional data types or compare the performance of different models.

Acknowledgment

We would like to express our appreciation to Professor Charles Liu and Professor Sarah Zelikovitz for their invaluable guidance and unwavering support throughout our research journey. Their constructive suggestions and patience were instrumental in keeping us organized, focused, and motivated, enabling us to achieve our research objectives. We sincerely appreciate their dedication and expertise.



References

- Xing, S., Hong, J., Wang, Y., Chen, R., Zhang, Z., Grama, A., Tu, Z., & Wang, Z. (2025, October 15). LLMs Can Get "Brain Rot"! arXiv. <https://arxiv.org/abs/2510.13928>
- Wettig, A., Gupta, A., Malik, S., & Chen, D. (2024). QuRating: Selecting High-Quality Data for Training Language Models. arXiv. <https://arxiv.org/abs/2402.09739>
- Yu, Y., Zhuang, Y., Zhang, J., Meng, Y., Ratner, A., Krishna, R., Shen, J., & Zhang, C. (2023). Large Language Model as Attributed Training Data. In Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS 2023).