nature portfolio

Corresponding author(s): Daniel. Alber (daniel.alber@nyulangone.org)

Last updated by author(s): November 21, 2024

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed
☐	☒ The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒ The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☒	☐ A description of all covariates tested
☒	☐ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒ For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☒	☐ Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	No software was used for data collection.
Data analysis	We used Python 3.10 and 3.11 as well as many open-source libraries, including ChromaDB 0.4.18, Flash Attention 2.0.1, matplotlib 3.8.2, numpy 1.26.2, pandas 2.1.3, PyTorch 2.0.1 and 2.1.1, scikit-learn 1.3.2, seaborn 0.13.0, spaCy 3.7.2, Hugging Face Transformers 4.31.0 and 4.35.2, and wandb 0.13.7. The LLM training code was modified from the Dao AI Lab Flash-Attention GitHub repository (https://github.com/Dao-AILab/flash-attention).

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

Every pre-training dataset and benchmark used in this paper was available as an open-source download at the time of this work. The Pile is no longer available for

public download. Due to security concerns, we do not plan to release our AI-generated poisoned medical articles nor any outputs from our poisoned LLM. The BIOS biomedical knowledge graph is available for public download (https://bios.idea.edu.cn/), and the UMLS can be accessed through an institutional or personal account (https://www.nlm.nih.gov/research/umls). Icons were sourced from the Noun Project (https://thenounproject.com/).

Research involving human participants, their data, or biological material

Policy information about studies with human participants or human data. See also policy information about sex, gender (identity/presentation), and sexual orientation and race, ethnicity and racism.

Reporting on sex and gender	N/A
Reporting on race, ethnicity, or other socially relevant groupings	N/A
Population characteristics	N/A
Recruitment	N/A
Ethics oversight	N/A

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	The study evaluated three sets of LLM outputs: 5,400 passages for the multi-concept 1.3B parameter models, 500 passages for each single-concept configuration of the 1.3B models, and 500 passages per configuration for the 4B parameter models. While these sample sizes detected statistically significant differences ($p<0.05$) between poisoned and baseline models, no formal power analysis was performed a priori to justify these specific numbers.
Data exclusions	No data were excluded.
Replication	N/A - This study presents a novel threat assessment of data poisoning in medical LLMs rather than an experimental study requiring replication, though the methodology is described in detail to enable future replication.
Randomization	Medical concepts were randomly selected as attack targets, and training batches were randomly chosen for replacement with malicious articles according to predefined probabilities.
Blinding	Fifteen human reviewers were blinded to both model status (poisoned vs baseline) and concept status (attack target vs control) when evaluating LLM-generated passages for potential medical harm.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern
☒	☐ Plants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Plants

Seed stocks	N/A
Novel plant genotypes	N/A
Authentication	N/A