

<https://doi.org/10.1038/s41746-025-01926-2>

Robustness tests for biomedical foundation models should tailor to specifications

R. Patrick Xian, Noah R. Baker, Tom David, Qiming Cui, A. Jay Holmgren, Stefan Bauer, Madhumita Sushil & Reza Abbasi-Asl

Check for updates

The rise of biomedical foundation models creates new hurdles in model testing and authorization, given their broad capabilities and susceptibility to complex distribution shifts. We suggest tailoring robustness tests according to task-dependent priorities and propose to integrate granular notions of robustness in a predefined specification to guide implementation. Our approach facilitates the standardization of robustness assessments in the model lifecycle and connects abstract AI regulatory frameworks with concrete testing procedures.

The growing presence of biomedical foundation models (BFMs), including large language models (LLMs), vision-language models (VLMs), and others, trained using biomedical or de-identified healthcare data, suggests they will eventually become integral to healthcare automation. Discussions on the risks of deploying algorithmic decision-making and generative AI in medicine have focused on bias and fairness. Robustness¹ is an equally important topic, which generally refers to the consistency of model prediction to distribution shifts. It is quantified using aggregated performance metrics, stratified comparisons across subsets of data, and worst-case performance. Robustness failures are an origin of the performance gap between model development and deployment, performance degradation over time, and, more alarmingly, the generation of misleading or harmful content by imperfect users or bad actors². The robustness of software also affects the legal responsibilities of providers³ because the software may cause harm (e.g., misinformation, financial loss, or injury) to users or third parties or require regulatory body authorization for deployment (e.g., medical devices)^{4,5}.

We examined over 50 existing BFMs covering different biomedical domains (see Fig. 1 and Supplementary Data 1). About 31.4% of them contain no robustness assessments at all. The most commonly presented evidence of model robustness is consistent performance across multiple datasets, which is adopted in 33.3% of the selected BFMs. Despite being a convenient proxy, consistent performance is not equivalent to a rigorous robustness guarantee because the relationships between datasets are generally unknown. Evaluations on shifted (5.9%) or synthetic data (3.9%), or data from external sites (9.8%) can be more effective but are not yet popular. To ensure the constructive and beneficial use of BFMs, we need to consider robustness evaluation across the model lifecycle and in intended application settings⁶. In biomedical domains, the various robustness concepts that warrant consideration (see Box 1 and the [repo](#)) motivate test customization. Inspired by test case prioritization in software engineering⁷, which improves

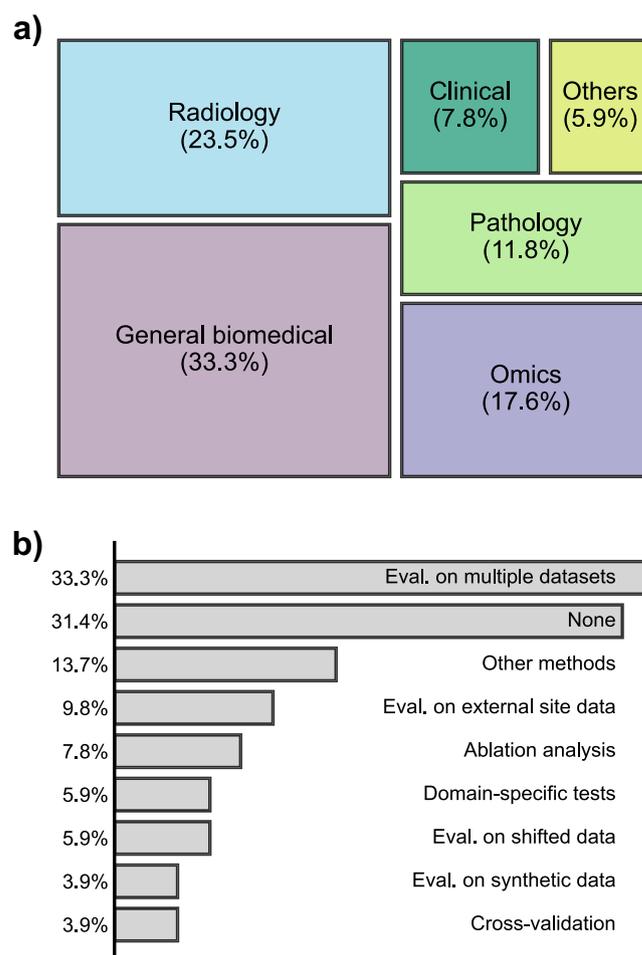


Fig. 1 | Existing robustness tests used for biomedical foundation models. The treemap in **a** illustrates the topical areas of the BFMs looked at for this study. “General biomedical” indicates that the model is trained on general-purpose biomedical datasets, and no domain specialization is emphasized in the model description. **b** Shows the distributions of robustness tests (eval. = evaluation). Because multiple tests were conducted for some models, the total proportion in **b** is larger than unity.

the cost-effectiveness of software testing by focusing on important test scenarios, we suggest designing effective robustness tests according to task-dependent robustness specifications constructed from priority scenarios (or priorities, see Fig. 2c) to facilitate test standardization while utilizing existing

Box 1 | Varieties of robustness concepts

Theoretical perspective

Natural robustness refers to the resistance to natural variations of data distribution, also called non-adversarial robustness to emphasize the distinction.

Adversarial robustness is a model's resistance to accidental or deliberate malicious data manipulations such as adversarial attacks and jailbreaks.

Interventional robustness refers to the invariance of model behavior to causally-defined interventions (e.g., image background or contrast shift, patient age shift) that characterize distribution shifts.

Out-of-distribution robustness is a model's resistance to semantic or covariate shift of test examples from the training data distribution.

Average-case & Worst-case robustness define the robustness concepts based on either the average or worst-case performance metric.

Feature & Data robustness define the robustness concepts based on the model performance change by removing or adding specific features or data in model development.

Engineering perspective

Development-stage robustness resists corruption during model development by noise in labels or data, or by adversarial attacks with poisoning and back-door attacks as the most common examples.

Deployment-stage robustness resists corruption at inference time due to transformed inputs, distracting information, or ill-intentioned behavior manipulations such as prompt injection and jailbreaks.

Component & Aggregated robustness refer to the individual-level robustness or combined robustness of the addressable components in a multi-expert or multiagent system.

Application perspective

Group robustness refers to the resistance of model performance against subpopulation distribution shifts. It is also called subgroup robustness.

Instance robustness refers to the consistency of model performance to label-preserving changes at the level of a particular instance or the worst-case instance.

Longitudinal robustness refers to the consistency of model performance against temporal shifts or nonstationary effects in the data distribution.

Knowledge robustness refers to the consistency of model knowledge against perturbations at the level of entities and their relations, against missing or nonstationary elements in a knowledge graph used for development or inference.

Vendor robustness refers to the consistency of model prediction to instrument vendors or the associated acquisition protocols that generate data. It is also known as acquisition-shift robustness.

Uncertainty-aware robustness defines a model's ability to account for uncertainties in the data distribution or in the model's internal mechanisms.

Prevalence-shift robustness refers to the consistency of model performance to changes in disease prevalence from epidemiological transitions.

Behavioral robustness refers to the consistency of model behavior in continuous interactions with humans or other models in a changing environment.

specialized tests as building blocks. Next, we introduce our proposal along with the background and motivation.

The robustness evaluation challenges

Foundation model characteristics. The versatility of use cases and exposure to complex distribution shifts are two major challenges of robustness evaluation (or testing)⁸ for foundation models that differentiate from prior generations of predictive algorithms. The versatility comes from foundation models' increased capabilities at inference time with knowledge injection through in-context learning, instruction following, the use of external tools (e.g., function calling) and data sources (e.g., retrieval augmentation), and with user steering of model behavior using specially designed prompts. These new learning paradigms blur the line between development and deployment stages and open up more avenues where models are exploited for their design imperfections.

Distribution shifts arise from natural changes in the data or intentional and sometimes malicious data manipulation (i.e., adversarial distribution shift)⁹. However, their distinction is increasingly nuanced in the era of foundation models⁹ due to the growing number of use cases. Natural distribution shifts can manifest biomedically in changing disease symptomatology, divergent population structure, etc. Inadvertent text deletion or image cropping results in data manipulations, potentially leading to adversarial examples that alter model behavior.

More elaborate shifts have been designed by targeted manipulation in model development and deployment^{10,11} through the cybersecurity lens.

Poisoning attacks involve stealthy modification of training data, while in backdoor attacks, a specific token sequence (called a trigger) is inserted during model training and activated during inference time¹². Distribution shifts in the deployment stage result in the majority of failure modes, including input transforms applied to text (deletion, substitution, and addition, including prompt injection, jailbreaks, etc) or images (noising, rotation, cropping, etc). Both natural distribution shifts and data manipulation yield out-of-distribution data¹³. They can have high domain-specificity or be created to target specific aspects of the model lifecycle, resulting in complex origins that are hard to trace exactly.

Robustness framework limitations. Aside from the challenges in scope, how to generate appropriate test examples for robustness evaluation is not often discussed. Two important robustness frameworks in ML, adversarial and interventional robustness, come from the security and causality viewpoints, respectively. The adversarial framework typically requires a guided search of test examples within a distance-bounded constraint, such as the bounds established by edit distance for text and by Euclidean distance for image in Fig. 2c, yet there is no practical guarantee that the test examples are sufficiently naturalistic to reflect reality. The interventional framework requires predefined interventions and a corresponding causal graph, which is not immediately available for every task. Theoretical guarantees provided by these frameworks generally require justifications in the asymptotic limit and don't necessarily translate into effective robustness in diverse yet highly contextualized deployment settings of specialized

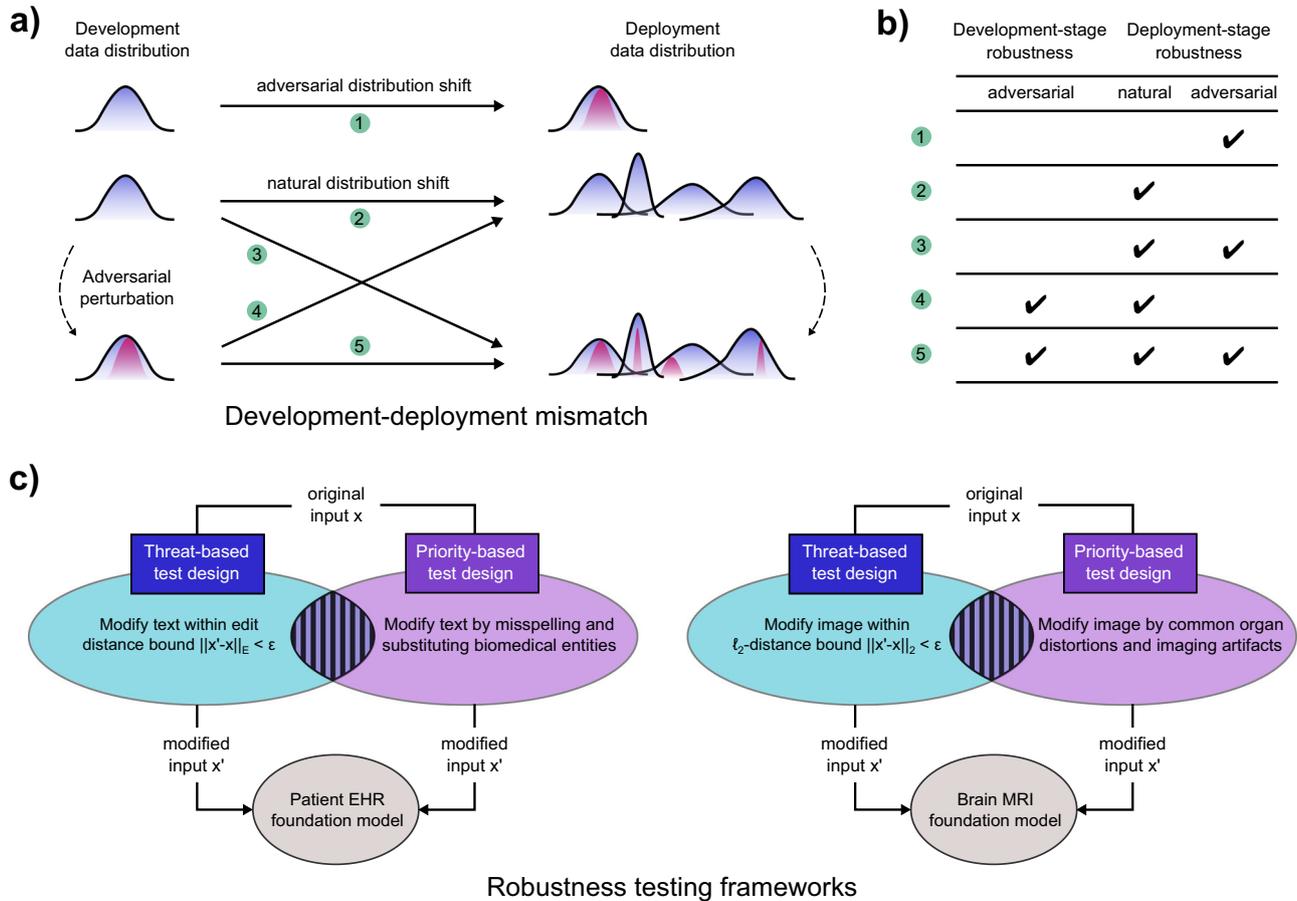


Fig. 2 | Settings and designs of robustness tests. Visualization in **a** illustrates the potential settings of development and deployment mismatches, which are represented in **b** according to the types of distribution shifts. Setting 1 indicates an adversarial contribution shift. Setting 2 refers to the natural distribution shift. In setting 3, adversarial perturbations are introduced in deployment, while in setting 4, they are applied to the training data. Setting 5 contains adversarial perturbations both during model development and deployment, such as in backdoor attacks.

c Specification of robustness by a simplified threat model (defined by a distance bound) or priority (defined by realistic artifacts) in the task domain, shown with two examples. The threat-based robustness tests use the error bound from edit distance for the EHR foundation model (left) and the Euclidean distance for the MRI foundation model (right). An overlap exists between these two approaches to generating test examples.

domains^{14,15}. Because robustness testing (and hence its associated guarantee) is critically dependent on the robustness framework of choice, we should design robustness tests that are more aligned with naturalistic settings and reflective of the priorities in corresponding domains.

Specifying robustness by priorities

Effective robustness evaluations require a pragmatic framework. The two aspects central to its specification are: (i) the degradation mechanism behind a distribution shift, and (ii) the task performance metric that requires protection against the shift. Mechanistically understanding a robustness failure mode requires establishing a connection between (i) and (ii), which is costly when accounting for every type of user interaction or impractical when the users have insufficient information on model development history or blackbox access. Moreover, multiple degradation mechanisms can simultaneously affect a particular downstream task.

Technical robustness evaluations in ML have generally tackled simplified threats for obtaining statistical guarantees, where a specific degradation mechanism guides the creation of test examples. Most adversarial

and interventional robustness tests fit into this category⁹, which often targets a considerably broader set of scenarios than those that are meaningful in reality. From the efficiency perspective, taking a priority-based viewpoint⁷ and focusing on retaining task performance under commonly anticipated degradation mechanisms in deployment settings is sufficient. Robustness tests based on simplified threat models and priorities are not mutually exclusive because accounting for realistic and meaningful perturbations (priority-based) has certain overlap with distance-bounded perturbations (threat-based), while the outcomes of priority-based tests should directly inform model quality. Figure 2c contains two examples comparing threat- and priority-based robustness tests for text and image data inputs. It illustrates the relationship between these two approaches for designing robustness tests.

We refer to the collection of priorities that demand testing for an individual task as a *robustness specification*. To contextualize it in naturalistic settings, we constructed two examples in Box 2 for an LLM-based pharmacy chatbot for over-the-counter (OTC) medicine and a VLM-based radiology report copilot for magnetic resonance imaging (MRI), both of

Box 2 | Examples of robustness specifications

Pharmacy chatbot for OTC medication

The chatbot, which can generate a medication request for patients, has been extensively tested for robustness under the following scenarios:

1. Up to 12 turns of conversation on topics related to primary care medication
2. Knowledge of specialized vocabulary, including products of major pharmaceutical vendors
3. Partly missing patient information (may refuse the task if the patient is unwilling to provide sufficient information)
4. Drug combinations, recommended intake quantities, and dosage limits
5. Adverse drug interactions and patient history
6. Stock availability and provider preference
7. Limited prescription authority (will refuse the task if asked for non-OTC or prescription drugs)
8. Paraphrases and typos in the instruction prompt
9. Off-topic requests and non-existent drugs

MRI radiology report copilot

The copilot that generates a report from one or a series of images has been extensively tested for robustness under the following scenarios:

1. Up to 15 turns of conversation, including comparison of up to five images.
2. Data from top-four vendors of MRI scanners.
3. Variability of MRI acquisition mode, image contrast, cropping or masking, and resolution.
4. Common imaging artifacts (may refuse the task if image quality is too low).
5. Major human organs and body sections (may refuse the task if images given don't resemble human anatomy).
6. Knowledge of specialized vocabulary.
7. Spatial relations of human anatomy.
8. Temporal ordering of distinguishing disease manifestations.
9. Paraphrases and typos in the instruction prompt.
10. Off-topic requests and non-MRI images.

which are attainable with existing research in BFM development. The specification contains a mixture of domain-specific (e.g., drug interaction, scanner information) and general aspects (e.g., paraphrasing, off-topic requests) that can induce model failures. The specification breaks down robustness evaluation into operationalizable units such that each is convertible into a small number of quantitative tests with guarantees. In reality, the test examples may come from augmenting or modifying the specified information in an existing data record^{14,15}, such as a clinical vignette or case report. The specification can accommodate the future capability expansion of models and risk assessment updates accordingly. We discuss below the feasibility of our proposal using existing and potential realizations of major types of robustness tests for BFMs in application settings (see Box 1).

Knowledge integrity. BFMs are knowledge models, and the knowledge acquisition process in the model lifecycle can be tempered to compromise knowledge robustness. Demonstrated examples for BFMs include a poisoning attack on biomedical entities, which have been shown to affect an entire knowledge graph in LLM-based biomedical reasoning¹⁰ and a backdoor attack using noise as the trigger for model failures in MedCLIP¹¹. Testing knowledge robustness should focus on knowledge integrity checks using realistic transforms. For text inputs, one may prioritize typos and distracting domain-specific information involving biomedical entities over random string perturbation under an edit-distance limit (see Fig. 2b). Existing examples include deliberately misinforming the model about the patient history¹⁶, negating scientific findings¹⁷, and substituting biomedical entities¹⁸ to induce erroneous model behaviors. For image inputs, one may prioritize the effects of common imaging and scanner artifacts¹⁹, alterations in organ morphology and orientation on model performance (see Fig. 2b).

Population structure. Explicit or implicit group structures are often present in biomedical and healthcare data, including prominent examples such as subpopulations organized by age group, ethnicity, or socio-economic strata, medical study cohorts with specific phenotypic traits, etc. BFM-enabled cross-sectional or longitudinal studies for patient similarity analysis and health trajectory simulation may lead to group or

longitudinal robustness issues when evaluating on incompatible populations. Group robustness assesses the model performance gap between the best- and worst-performing groups, either identifiable through the label or hidden in the dataset. Testing group robustness may modify subpopulation labels in patient descriptions to gauge the change in model performance²⁰. At a finer granularity, instance robustness represents the performance gap between instances that are more prone to robustness failures than others, which are likely corner cases. It is important when the model deployment setting requires a minimum robustness threshold for every instance. Robustness testing in this context may use a balanced metric to reflect the impact of input modifications across individual instances.

Uncertainty awareness. The machine learning community typically distinguishes between aleatoric uncertainty, which comes from inherent data variability, and epistemic uncertainty, which arises from insufficient knowledge of the model in the specific problem context. Robustness tests against aleatoric uncertainty may assess the sensitivity of model output to prompt formatting and paraphrasing, while assessing robustness to epistemic uncertainty may use out-of-context examples²¹ to examine if a model acknowledges the significant missing contextual information in domain-specific cases (e.g., presenting the model with a chest X-ray image and asking for a knee injury diagnosis). Additionally, uncertain information may also be directly verbalized in text prompts, a fitting scenario in biomedicine, to examine its influence on model behavior. Overall, the current generation of robustness evaluations has not yet included realistic uncertain scenarios often encountered in medical decision-making, although robustness against uncertainty is an important topic in practice.

Embracing emerging complexities

Previous scenarios primarily consider assessing a monolithic model using single-criterion robustness tests. Specifying and testing robustness for more complex AI systems should also account for performance tradeoffs, model architectures, and user interactions.

Table 1 | Robustness tests in the adaptation and update of BFM-based devices and services

Adaptation / Update	Execution	Cost	Potential usage of robustness tests
Model weights	Model provider	high	Helps choose and optimize the training method that yields models with a better robustness-accuracy tradeoff
Model architecture	Model provider	high	Helps choose architectural parameters that are easier to attain a robust model than others when training on the same dataset
User prompt	Model user	low	Helps select more robust prompt templates and improve the reliability of outputs
User confidence	Model user	medium	Helps adjust and calibrate user confidence to deploy the model for appropriate tasks

Metrics and stakeholders. Evaluating tradeoffs between various robustness metrics and criteria offers a balanced view of a model's robustness across different dimensions and through metric aggregation. These more comprehensive robustness tests are essential in assessing whether the model's behavior reaches an optimal balance or is suitable for applications with distinct risk levels or stakeholders (see SI section 1). When models are integrated into a healthcare workflow, they can affect downstream biomedical outcomes. For example, using LLMs to summarize or VLMs to generate case reports may influence clinician decisions by emphasizing certain conditions or sentiments, affecting diagnoses or procedures. This highlights the need for considering robustness tests with the relevant stakeholder(s) in the loop and behavioral robustness across diverse interaction settings to assess the model's impact on the care journey.

Compound systems. As modularity and maintainability become increasingly important, decision-making will be delegated to specialized subunits in a multi-expert (such as a mixture of experts) or multiagent system²² with a centralized coordinating unit. In these compound AI systems, each addressable subsystem is subject to testing and maintenance according to capability demand and regulatory compliance (see Fig. 2c). For example, Polaris²³ from Hippocratic AI features a multiagent medical foundation model that writes medical reports and notes as well as engages in low-risk patient interactions. Future systems with specialized units can mimic the group decision-making process in healthcare²⁴ to manage real-world complexities through enhanced reasoning and cooperative performance gain. Robustness tests for compound AI systems may consider different specifications for subsystems depending on the part-part and part-whole relationship in identifying bottlenecks and cascading effects associated with robustness failures.

Bridging policy with implementation

Ensuring robustness for BFM requires advancing regulatory policies for both AI and health information technology. Currently, the leading AI regulatory frameworks, such as the EU AI Act and the US Federal AI Risk Management Act, recognize the relation between natural and adversarial notions of robustness but contain insufficient details to guide implementation in domain-specific applications (see SI section 2). Existing health information technology regulations, such as the US-based HTI-1 final rule by the Office of the National Coordinator, focus primarily on transparency and disclosures of the use of predictive decision support models, yet lack detailed robustness requirements. The situation is in part due to the lack of a safety bare minimum for specific biomedical applications⁵ and the fast-evolving technological landscape, which can exacerbate the challenges laid out at the beginning of this Comment. These existing gaps make concrete community-endorsed standards on robustness even more important.

Considerations in implementation. Mandating robustness specifications according to the tasks and the biomedical domains (see Figs. 1, 2) provides

a means to map policy objectives onto real-world implementations. It also facilitates evidence collection and enables effective risk management throughout the model lifecycle. We advocate that robustness specifications (i) should seek community endorsement to gain a broad adoption; (ii) should consider the permissible tasks and user group characteristics due to the difference in user journeys; (iii) should inform regulatory standards, such as in the construction of quantitative risk thresholds²⁵ or safety cases by enriching the failure mode taxonomy of BFMs and improving their informativeness. These considerations will facilitate the implementation of robustness specifications and ensure that their adoption is within shared interests of stakeholders.

Community benefits. Establishing a consensus-driven robustness specification from the research community will incentivize systematic efforts by model developers, research institutions, and independent third parties. For model developers, robustness testing informs model selection and updates. For the model deployment team and model users, robustness testing allows for identifying inference-time adjustments of prompt templates to improve the reliability of BFM applications. These potential uses of robustness tests are summarized in Table 1. In addition, robustness specifications provide templates for failure-reporting procedures to allow users to provide timely feedback to the deployment team. Integrating robustness specifications with incident reporting mechanisms⁶ facilitates the identification of model vulnerabilities and guides targeted improvements or informs post-hoc adjustments to model behavior. Their implementations can assist the training of end-users to recognize potential failures, calibrate user confidence, and enact mitigation strategies.

Data availability

Data is provided within the manuscript or supplementary information files. An extended reference on the robustness concepts discussed here is available online (<https://github.com/RealPolitiX/bfm-robust>).

R. Patrick Xian^{1,2}✉, **Noah R. Baker**³, **Tom David**⁴, **Qiming Cui**^{1,5}, **A. Jay Holmgren**⁶, **Stefan Bauer**⁷, **Madhumita Sushil**⁸ & **Reza Abbasi-Asl**^{1,2,8,9}✉

¹Department of Neurology, University of California, San Francisco, San Francisco, CA, USA. ²Weill Institute for Neurosciences, University of California, San Francisco, San Francisco, CA, USA. ³Biological and Medical Informatics Graduate Program, University of California, San Francisco, San Francisco, CA, USA. ⁴PRISM Eval, Paris, France. ⁵Department of Bioengineering, University of California, Berkeley, Berkeley, CA, USA. ⁶Division of Clinical Informatics and Digital Transformation, University of California, San Francisco, San Francisco, CA, USA. ⁷School of Computation, Information and Technology, Technical University of Munich & Helmholtz AI, Garching bei München, Germany. ⁸Bakar Computational Health Sciences Institute, University of California, San Francisco, San Francisco, CA, USA. ⁹Department of Bioengineering and Therapeutic Sciences, University of California, San Francisco, San

Francisco, CA, USA. ✉e-mail: xrpatrick@gmail.com;
reza.abbasiasl@ucsf.edu

Received: 17 January 2025; Accepted: 1 August 2025;

Published online: 29 August 2025

References

1. Tocchetti, A. et al. A.I. Robustness: a human-centered perspective on technological challenges and opportunities. *ACM Comput. Surv.* **57**, 141 (2025).
2. Kostick-Quenet, K. M. & Gerke, S. AI in the hands of imperfect users. *npj Digit. Med.* **5**, 197 (2022).
3. Ladkin, P. B. Robustness of Software. *Digit. Evid. Electron. Signature Law Rev.* **17**, 15–24 (2020).
4. Warrach, H. J., Tazbaz, T. & Califf, R. M. FDA perspective on the regulation of artificial intelligence in health care and biomedicine. *JAMA* **333**, 241–247 (2025).
5. Freyer, O., Wiest, I. C., Kather, J. N. & Gilbert, S. A future role for health applications of large language models depends on regulators enforcing safety standards. *Lancet Digit. Health* **6**, e662–e672 (2024).
6. Lyell, D., Wang, Y., Coiera, E. & Magrabi, F. More than algorithms: an analysis of safety events involving ML-enabled medical devices reported to the FDA. *J. Am. Med. Inform. Assoc.* **30**, 1227–1236 (2023).
7. Rothermel, G., Untch, R., Chu, C. & Harrold, M. Prioritizing test cases for regression testing. *IEEE Trans. Softw. Eng.* **27**, 929–948 (2001).
8. Chen, P.-Y., Liu, S. & Paul, S. *Foundational Robustness of Foundation Models*. NeurIPS Tutorial (2022). <https://research.ibm.com/publications/foundational-robustness-of-foundation-models>.
9. Qi, X. et al. AI Risk Management Should Incorporate Both Safety and Security. arXiv:2405.19524. <http://arxiv.org/abs/2405.19524> (2024).
10. Yang, J. et al. Poisoning medical knowledge using large language models. *Nat. Mach. Intell.* **6**, 1156–1168 (2024).
11. Jin, R., Huang, C.-Y., You, C. & Li, X. Back-door Attack on Unpaired Medical Image-Text Foundation Models: A Pilot Study on Med-CLIP. In *2024 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 272–285 (IEEE, 2024).
12. Chowdhury, A. G. et al. Breaking Down the Defenses: A Comparative Survey of Attacks on Large Language Models. arXiv:2403.04786. <http://arxiv.org/abs/2403.04786> (2024).
13. Karunanayake, N., Gunawardena, R., Senevi-Ratne, S. & Chawla, S. Out-of-distribution data: an acquaintance of adversarial examples - a survey. *ACM Comput. Surv.* **57**, 210 (2025).
14. Hager, P. et al. Evaluation and mitigation of the limitations of large language models in clinical decision-making. *Nat. Med.* **30**, 2613–2622 (2024).
15. Johri, S. et al. An evaluation framework for clinical use of large language models in patient interaction tasks. *Nat. Med.* **31**, 77–86 (2025).
16. Han, T. et al. Medical large language models are susceptible to targeted misinformation attacks. *npj Digit. Med.* **7**, 288:1–9 (2024).
17. Yan, Q., He, X., Yue, X. & Wang, X. E. *Worse than random? An embarrassingly simple probing evaluation of large multimodal models in medical VQA*. Findings of the ACL, 10345–10359 (2025).
18. Xian, R. P. et al. Assessing biomedical knowledge robustness in large language models by query-efficient sampling attacks. *Trans. Mach. Learn. Res.* (2024). <https://openreview.net/forum?id=pvoI5JyVYB>.
19. Boone, L. et al. ROOD-MRI: Benchmarking the robustness of deep learning segmentation models to out-of-distribution and corrupted data in MRI. *NeuroImage* **278**, 120289 (2023).
20. Yang, Y., Zhang, H., Katabi, D. & Ghassemi, M. Change is hard: a closer look at subpopulation shift. In *ICML*, 39584–39622 (Honolulu, Hawaii, USA, 2023).
21. Chandu, K. et al. Certainly Uncertain: A Benchmark and Metric for Multimodal Epistemic and Aleatoric Awareness. *ICLR* (2025). <https://openreview.net/forum?id=cQ25MQQSNl>.
22. Wang, W. et al. *A Survey of LLM-based Agents in Medicine: How far are we from Baymax?* Findings of the ACL, 10345–10359 (Vienna, Austria, 2025). <https://aclanthology.org/2025.findings-acl.539/>.
23. Mukherjee, S. et al. Polaris: A Safety-focused LLM Constellation Architecture for Healthcare. arXiv:2403.13313. <http://arxiv.org/abs/2403.13313> (2024).
24. Radcliffe, K., Lyson, H. C., Barr-Walker, J. & Sarkar, U. Collective intelligence in medical decision-making: a systematic scoping review. *BMC Med. Inform. Decis. Making.* **19**, 158 (2019).
25. Koessler, L., Schuett, J. & Anderjurg, M. Risk thresholds for frontier AI. arXiv:2406.14713. <http://arxiv.org/abs/2406.14713> (2024).

Acknowledgements

R.P.X. thanks N. Rethmeier for helpful discussions. R.A.-A. would like to acknowledge funding from the Weill Neurohub. M.S. is partially funded by the National Cancer Institute of the National Institutes of Health under Award Number P30CA082103. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Author contributions

R.P.X. and R.A.-A. conceptualized the main idea of the work. R.P.X. prepared the figures and tables, and wrote the majority of the text with significant contributions from N.R.B., T.D., Q.C., and A.J.H. S.B. provided theoretical support. M.S. and R.A.-A. provided background information in different application settings. All authors edited the text and contributed to the discussions to finalize the manuscript.

Competing interests

T.D. is a co-founder and director of governance & standardization at PRISM Eval. Other authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41746-025-01926-2>.

Correspondence and requests for materials should be addressed to R. Patrick Xian or Reza Abbas-Asl.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025