

Why It Failed: A Benchmark to Evaluate Interpretability

Joel Mathew^{*1} Aditya Lagu¹ Anthony Tang¹ Prudhviraj Naidu¹²

¹Algoverse AI Research
²University of California San Diego

Abstract

We introduce *Why It Failed*, a benchmark for evaluating whether interpretability methods can explain model failures. We test last token logistic probes on Gemma-2 2B across four basic reasoning tasks and find they fail to predict model failures, achieving near-chance performance across all tasks. Our benchmark provides a standardized framework to evaluate whether interpretability methods can explain model failures. Lastly, we motivate the AI community to move beyond reporting quantitative metrics and seek explanations of when and why models fail.

Code —

<https://github.com/anthonytang/LLM-Detection-Testing>

1 Introduction

Large language models achieve impressive performance on diverse benchmarks, yet they still fail in puzzling and unpredictable ways. A model might score 80% on a common-sense reasoning task, but this aggregate metric obscures crucial questions: *When* does the model fail? *Why* does it fail? Are failures random noise, or do they follow systematic patterns? Current evaluation practices report quantitative metrics without explaining the underlying failure modes, leaving practitioners uncertain about when and where models can be safely deployed.

The interpretability community has made significant progress in understanding what models know—identifying circuits that implement specific algorithms (Wang et al. 2022), discovering human interpretable concepts in models’ activations (Bricken et al. 2023). However, interpretability methods are typically validated on their ability to detect unwanted behaviors and monitor models in deployment (McKenzie et al. 2025), rather than on their ability to systematically explain common failure cases of models on standard benchmarks.

This gap has important consequences. Without systematic explanations of failures, we cannot make informed deployment decisions. Explanations enable us to distinguish

model-specific quirks from dataset artifacts, and to understand whether failure patterns generalize across model families or emerge from particular architectures. Most fundamentally, explaining failures allows us to move beyond simply reporting benchmark scores toward building scientific theories of model limitations.

We propose *Why It Failed*, a benchmark for evaluating whether interpretability methods can explain model failures. Our contributions are twofold.

1. We introduce a benchmark framework that shifts focus from explaining what models know to explaining what they don’t know.
2. We show that standard last-token probing is insufficient for predicting failures.

We invite the interpretability community to test their methods against this benchmark and push toward explanations that illuminate model limitations.

The rest of the paper is organized as follows, Section 2 reviews related work in benchmarking explanations, interpretability methods and generating explanations using SAEs. Section 3 describes our benchmark framework and how we generated our benchmark examples. Section 4 describes our probing experiment methodology and discusses our results. In Section 5, we discuss limitations of our work and provides questions for future work.

2 Related Work

Benchmarking Interpretability Methods

Mills et al. (2025) propose ALMANACS, a benchmark to test how well explanations can predict model behaviour on new inputs. They focus on behaviour related to AI safety such as ethical reasoning, self preservation or harmful requests. The authors construct such safety scenarios to invoke specific model behaviour. They primarily generate explanations by blackbox methods or through attribution methods.

Our approach on the other hand focuses on explaining model failures in standard benchmarks. We select established reasoning benchmarks that measure basic model capabilities. We then establish our baseline via probing models internal activations, which has seen great success in recognizing diverse model behaviours and steering models (Beaglehole et al. 2025; McKenzie et al. 2025).

^{*}Correspondence to joel.mathew@sjsu.edu

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Chandrasekaran et al. (2018) study a failure prediction task for a visual question answering (VQA) model. Humans are asked to predict whether the VQA model will get a question right or wrong. They find humans are able to predict model failures above chance. However, explainable methods do not improve human performance. We on the other hand focus on pure language model and first ask what is the best predictive performance we can get for predicting model failures.

Interpretability Methods for Language Models

Recent interpretability research has focused on key three factors: understanding model traces (CoT-Interpretability), understanding the internal representations of language models (Sparse Autoencoders) and training probes on internal representations of models. These methods vary in their approach regarding what aspect of the model they study but are focused on primarily monitoring and steering models.

Chain-Of-Thought (CoT) Interpretability analyzes reasoning traces that the language model produces before the final response. Baker et al. (2025) were able to detect models reward hacking by monitoring its reasoning traces. Korbak et al. (2025) argues CoT interpretability is a promising research direction however CoTs can easily succumb to optimization pressures of directly training CoT and loose their usefulness for monitoring. Kirchner et al. (2024) proposes a scheme to make models reasoning more legible by posing the task as a prover-verifier game where the model (as a prover) has to provide legible explanations for its answer that allows a verifier to predict whether the model got the answer right or wrong.

Probing methods train auxiliary classifiers on a base model hidden states to detect concepts that the base model was not explicitly trained for (Chen et al. 2020). Subsequent work has shown linear probing can detect unsafe behaviour and steer the model towards a desirable behaviour (Beaglehole et al. 2025; McKenzie et al. 2025).

Sparse Autoencoders learn overcomplete, sparse decompositions of neural activations, hypothesizing that individual SAE features correspond to monosemantic concepts (Cunningham et al. 2023; Bricken et al. 2023). Jiang et al. (2025) uses SAEs to generate explanations for difference in datasets. For GSM8K, Jiang et al. (2025) demonstrated "Math word problems involving time, distance, and speed" had lower accuracy.

Our benchmark evaluates whether these interpretability methods can be repurposed to identify systematic patterns that distinguish successes from failures.

3 Why It Failed Benchmark

Benchmark Overview

Core Question: Can current interpretability methods explain why models fail? When a model makes errors on a benchmark, we want to know whether interpretability techniques can surface explanations that genuinely capture the underlying causes of these failures.

What constitutes a good failure explanation? We argue that a good explanation must satisfy three criteria:

1. **Faithful:** Given the explanation, we should be able to predict whether the model will fail on unseen inputs or when deployed in new settings.
2. **Causal:** We should be able to manipulate model performance based on the explanation, either by generating adversarial inputs that exploit the identified weakness or by improving the model performance (by identifying training dataset or methodology issues) based on the explanation.
3. **Human-interpretable:** A human examining the explanation should be able to predict whether the model will fail on new inputs, without needing to run the model or inspect its internals.

In this work, we focus on **faithfulness** as a measurable, necessary (though not sufficient) condition for good explanations. While causal interventions and human-interpretability are important criteria, we leave their systematic evaluation to future work.

Evaluation Pipeline: For each task, we construct a balanced dataset of k success cases and k failure cases. An **Explainer** (e.g., linear probe, chain-of-thought interpreter, SAEs) processes these $2k$ training instances to generate explanations. A **Predictor** then uses these explanations to classify whether the model will succeed or fail on new, unseen test instances. We measure the Predictor's performance using AUC on the held-out test set.

Task & Model Selection

Tasks: Our benchmark is constructed from four diverse reasoning tasks from standard LLM evaluation suites: PIQA (physical commonsense reasoning), BoolQ (yes/no reading comprehension), WinoGrande (common-sense coreference resolution), and Social IQa (social commonsense reasoning). These tasks span a variety of cognitive capabilities:

- **PIQA** (Bisk et al. 2020): Given a goal and two potential solutions, the model must select which action would successfully achieve the goal. For example: "*How do I ready a guinea pig cage for its new occupants?*" with options involving paper strips vs. jeans material as bedding. This tests physical commonsense about everyday object interactions.
- **BoolQ** (Clark et al. 2019): Given a passage and a yes/no question, the model must determine the correct boolean answer. For example: "*Does ethanol take more energy to make than it produces?*" paired with a technical passage. This tests reading comprehension and factual reasoning.
- **WinoGrande** (ai2 2019): Fill-in-the-blank coreference resolution where the model must determine which noun a pronoun refers to. For example: "*John moved the couch from the garage to the backyard to create space. The _ is small.*" (options: garage, backyard). This tests common-sense reasoning about spatial and physical constraints.
- **Social IQa** (Sap et al. 2019): Questions about people's actions, intentions, and social implications. For example: "*Sydney walked past a homeless woman asking for change but did not have any money. Sydney felt bad afterwards. How would you describe Sydney?*" This tests social and emotional reasoning.




Figure 1: Last token probes fail to predict model failures. They consistently fail for all layers except on BoolQ where latter layer slightly outperform random chance. For each task, we plot the test ROC AUC score for a probe trained on the output of a layer. The error bars refer to 99.9% confidence intervals. The dotted red-line depicts random chance performance.

Why these tasks? These tasks represent fundamental capabilities required for real-world language understanding: physical intuition, reading comprehension, linguistic reasoning, and social awareness. Critically, despite being simple multiple-choice or boolean questions, models can fail on them for complex, non-obvious reasons. A model might fail on PIQA not due to lacking physical commonsense reasoning due to vocabulary gaps or on specific cultural contexts unrelated to physical commonsense reasoning. The tasks span a diverse range of contexts and are widely-used benchmarks in the LLM community as a way to measure LLM capability and decide whether to deploy LLMs.

Model: We use Gemma-2 2B (Gemma Team et al. 2024), a 2-billion parameter transformer model developed by Google. We chose Gemma-2 2B because: (1) it is small enough to run efficiently on consumer hardware, enabling rapid iteration and making our benchmark accessible to researchers without extensive computational resources; (2) it is a practical, industry-grade model designed for real-world deployment rather than a synthetic research model, ensuring our findings generalize to models used in practice.

Data Collection: We construct our benchmark by collecting model responses on these four tasks and categorizing them into correct and incorrect predictions. For each task, we randomly sample from the evaluation set until we collect exactly k correct predictions and k incorrect predictions, creating a balanced dataset. All examples retain their standard multiple-choice or boolean format from the original tasks.

4 Probes fail to predict model failures

Experimental Setup

Transformer Architecture and Residual Stream Consider a transformer model with L layers. For each layer $l \in \{1, 2, \dots, L\}$ and token position t , we define the layer computation as:

$$\mathbf{x}_{\text{mid}}^{(l,t)} = \mathbf{x}_{\text{pre}}^{(l,t)} + \sum_{\text{head } h} \text{attn}^{(l,h)} \left(\mathbf{x}_{\text{pre}}^{(l,t)}, \mathbf{x}_{\text{pre}}^{(l,1:t)} \right) \quad (1)$$

$$\mathbf{x}_{\text{post}}^{(l,t)} = \mathbf{x}_{\text{mid}}^{(l,t)} + \text{MLP}^{(l)} \left(\mathbf{x}_{\text{mid}}^{(l,t)} \right) \quad (2)$$

where:

- $\mathbf{x}_{\text{pre}}^{(l,t)} \in R^d$ is the input to layer l at position t (the pre-residual stream), where d is the transformer model dimension
- $\mathbf{x}_{\text{mid}}^{(l,t)} \in R^d$ is the mid-residual stream (after attention)
- $\mathbf{x}_{\text{post}}^{(l,t)} \in R^d$ is the output of layer l (post-residual stream)
- $\text{attn}^{(l,h)}$ denotes the h -th attention head in layer l
- $\text{MLP}^{(l)}$ denotes the feedforward network in layer l

Probe Training We train linear probes on the post-residual stream $\mathbf{x}_{\text{post}}^{(l,t)}$ at the final token position of each prompt. Given a prompt from task \mathcal{T} with tokens t_1, \dots, t_n , we extract activations $\mathbf{x}_{\text{post}}^{(l,n)}$ at layers $l \in \{5, 10, 15, 20, 25\}$ for Gemma-2 2B. These layers are evenly spaced to capture representation learning throughout the model’s depth.

For each layer l and task \mathcal{T} , we train a binary logistic regression classifier:

$$\hat{y} = \sigma(\mathbf{w}^{(l)} \cdot \mathbf{x}_{\text{post}}^{(l,n)} + b) \quad (3)$$

where $\mathbf{w}^{(l)} \in R^d$ is the probe direction, $b \in R$ is a bias term, and σ is the sigmoid function. The probe is trained to predict whether the model will succeed ($y = 1$) or fail ($y = 0$) on the given prompt. We use scikit-learn’s logistic regression implementation with default hyperparameters.

Data Splits For each task, we partition our 4,000 examples (2,000 successes and 2,000 failures) into:

- **Training set:** 2,000 examples (1,000 successes, 1,000 failures)
- **Validation set:** 1,000 examples (500 successes, 500 failures)
- **Test set:** 1,000 examples (500 successes, 500 failures)

We report test set performance using Area Under the ROC Curve (AUC) as our primary metric. An AUC of 0.5 indicates chance-level performance, while 1.0 indicates perfect classification.

Results

Figure 1 shows the test set AUC for linear probes trained at different layers across all four tasks. The results reveal a striking failure: **linear probes do not outperform random chance at predicting model failures.**

Across all tasks and layers, probe performance clusters near the random baseline (AUC = 0.5, shown as a red dashed line). The best-performing configuration achieves only modest improvement: BoolQ at higher layers achieves performance slightly above chance.

We observe several consistent patterns:

1. **No layer captures failure modes:** Performance remains near chance across all chosen layers, suggesting that the distinction between success and failure is not linearly encoded at any single layer’s residual stream.
2. **Task variation is minimal:** All tasks exhibit fundamentally the same pattern of near-chance performance barring probes trained on latter layers on BoolQ. This suggests the limitation is not task-specific but reflects a broader issue with using linear probes as explainers.

These results demonstrate that linear probes, despite their success in identifying semantically meaningful directions in prior work, fail as explainers in our benchmark. The model’s internal representations do not contain linearly accessible information sufficient to predict when the model will fail.

5 Discussion & Conclusion

We introduce *Why It Failed*, a benchmark for evaluating whether interpretability methods can explain model failures. Our framework operationalizes faithful explanations through predictive power: good explanations should enable prediction of failures on unseen inputs. While we focus on faithfulness in this work, our framework naturally extends to causal manipulation and human-interpretability criteria.

We showed simple last token probing fail to explain model failures. It remains unclear why probes slightly outperformed random chance on BoolQ. Our work can easily be extended to include:

- **Richer probing methods:** Mean-token probing across sequences, attention-probes (McKenzie et al. 2025)
- **Sparse representation methods:** Sparse autoencoders (SAEs) can identify clusters and offer explanations of why models fails (Jiang et al. 2025)
- **Chain-of-thought interpretability:** Analyzing reasoning traces in models that produce intermediate steps. For these tasks, we did not generate reasoning traces from Gemma-2 2B. Future works would need to first generate examples of success / failure with reasoning traces.

Furthermore, we currently measure only faithfulness through predictive power. Future work should incorporate:

- **Causal metrics:** Can explanations generate adversarial examples that flip model predictions? Can they guide interventions that improve performance?
- **Human-interpretability scoring:** Auto-interpretability methods (Paulo et al. 2025) can generate explanation for SAE features. (Kantamneni et al. 2025) proposes a similar framework could be used to explain logistic probes. However, it remains an open question whether these methods provide faithful human interpretable explanations.

Conclusion

Why It Failed benchmark asks what models don’t know and why. Our key message is simple: **we should aim to explain benchmark failures, not just report accuracy numbers.** This benchmark provides a concrete framework for evaluating whether interpretability methods achieve this goal. While linear probes fall short, they represent just the beginning. We invite the community to test their methods against this benchmark and push toward explanations that truly illuminate model limitations.

References

2019. WinoGrande: An Adversarial Winograd Schema Challenge at Scale.

Baker, B.; Huizinga, J.; Gao, L.; Dou, Z.; Guan, M. Y.; Madry, A.; Zaremba, W.; Pachocki, J.; and Farhi, D. 2025. Monitoring Reasoning Models for Misbehavior and the Risks of Promoting Obfuscation. arXiv:2503.11926.

Beaglehole, D.; Radhakrishnan, A.; Boix-Adserà, E.; and Belkin, M. 2025. Toward universal steering and monitoring of AI models. arXiv:2502.03708.

Bisk, Y.; Zellers, R.; Bras, R. L.; Gao, J.; and Choi, Y. 2020. PIQA: Reasoning about Physical Commonsense in Natural Language. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*.

Bricken, T.; Templeton, A.; Batson, J.; Chen, B.; Jermyn, A.; Conerly, T.; Turner, N.; Anil, C.; Denison, C.; Askell, A.; Lasenby, R.; Wu, Y.; Kravec, S.; Schiefer, N.; Maxwell, T.; Joseph, N.; Hatfield-Dodds, Z.; Tamkin, A.; Nguyen,

K.; McLean, B.; Burke, J. E.; Hume, T.; Carter, S.; Henighan, T.; and Olah, C. 2023. Towards Monosemanticity: Decomposing Language Models With Dictionary Learning. *Transformer Circuits Thread*. <Https://transformer-circuits.pub/2023/monosemantic-features/index.html>.

Chandrasekaran, A.; Prabhu, V.; Yadav, D.; Chattopadhyay, P.; and Parikh, D. 2018. Do explanations make VQA models more predictable to a human? In Riloff, E.; Chiang, D.; Hockenmaier, J.; and Tsujii, J., eds., *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 1036–1042. Brussels, Belgium: Association for Computational Linguistics.

Chen, M.; Radford, A.; Child, R.; Wu, J.; Jun, H.; Luan, D.; and Sutskever, I. 2020. Generative Pretraining From Pixels. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 1691–1703. PMLR.

Clark, C.; Lee, K.; Chang, M.-W.; Kwiatkowski, T.; Collins, M.; and Toutanova, K. 2019. BoolQ: Exploring the Surprising Difficulty of Natural Yes/No Questions. In *NAACL*.

Cunningham, H.; Ewart, A.; Riggs, L.; Huben, R.; and Sharkey, L. 2023. Sparse Autoencoders Find Highly Interpretable Features in Language Models. arXiv:2309.08600.

Gemma Team; Riviere, M.; Pathak, S.; Sessa, P. G.; Hardin, C.; Bhupatiraju, S.; Hussenot, L.; Mesnard, T.; Shahriari, B.; Ramé, A.; Ferret, J.; Liu, P.; Tafti, P.; Friesen, A.; Casbon, M.; Ramos, S.; Kumar, R.; Lan, C. L.; Jerome, S.; Tsitsulin, A.; Vieillard, N.; Stanczyk, P.; Girgin, S.; Momchev, N.; Hoffman, M.; Thakoor, S.; Grill, J.-B.; Neyshabur, B.; Bachem, O.; Walton, A.; Severyn, A.; Parrish, A.; Ahmad, A.; Hutchison, A.; Abdagic, A.; Carl, A.; Shen, A.; Brock, A.; Coenen, A.; Laforge, A.; Paterson, A.; Bastian, B.; Piot, B.; Wu, B.; Royal, B.; Chen, C.; Kumar, C.; Perry, C.; Welty, C.; Choquette-Choo, C. A.; Sinopalnikov, D.; Weinberger, D.; Vijaykumar, D.; Rogozińska, D.; Herbison, D.; Bandy, E.; Wang, E.; Noland, E.; Moreira, E.; Senter, E.; Eltyshev, E.; Visin, F.; Rasskin, G.; Wei, G.; Cameron, G.; Martins, G.; Hashemi, H.; Klimczak-Plucińska, H.; Batra, H.; Dhand, H.; Nardini, I.; Mein, J.; Zhou, J.; Svensson, J.; Stanway, J.; Chan, J.; Zhou, J. P.; Carrasqueira, J.; Iljazi, J.; Becker, J.; Fernandez, J.; van Amersfoort, J.; Gordon, J.; Lipschultz, J.; Newlan, J.; yeong Ji, J.; Mohamed, K.; Badola, K.; Black, K.; Millican, K.; McDonell, K.; Nguyen, K.; Sodhia, K.; Greene, K.; Sjoesund, L. L.; Usui, L.; Sifre, L.; Heuermann, L.; Lago, L.; McNealus, L.; Soares, L. B.; Kilpatrick, L.; Dixon, L.; Martins, L.; Reid, M.; Singh, M.; Iverson, M.; Görner, M.; Velloso, M.; Wirth, M.; Davidow, M.; Miller, M.; Rahtz, M.; Watson, M.; Risdal, M.; Kazemi, M.; Moynihan, M.; Zhang, M.; Kahng, M.; Park, M.; Rahman, M.; Khatwani, M.; Dao, N.; Bardoliwalla, N.; Devanathan, N.; Dumai, N.; Chauhan, N.; Wahltinez, O.; Botarda, P.; Barnes, P.; Barham, P.; Michel, P.; Jin, P.; Georgiev, P.; Culliton, P.; Kuppala, P.; Comanescu, R.; Merhej, R.; Jana, R.; Rokni, R. A.; Agarwal, R.; Mullins, R.; Saadat, S.; Carthy, S. M.; Cogan, S.; Perrin, S.; Arnold, S. M. R.; Krause, S.; Dai, S.; Garg, S.; Sheth, S.; Ronstrom, S.; Chan, S.; Jordan, T.; Yu, T.; Eccles, T.; Hennigan, T.; Kociský, T.; Doshi, T.; Jain, V.; Yadav, V.; Meshram, V.; Dharmadhikari, V.; Barkley, W.; Wei, W.; Ye, W.; Han, W.; Kwon, W.; Xu, X.; Shen, Z.; Gong, Z.; Wei, Z.; Cotruta, V.; Kirk, P.; Rao, A.; Giang, M.; Peran, L.; Warkentin, T.; Collins, E.; Barral, J.; Ghahramani, Z.; Hadsell, R.; Sculley, D.; Banks, J.; Dragan, A.; Petrov, S.; Vinyals, O.; Dean, J.; Hassabis, D.; Kavukcuoglu, K.; Farabet, C.; Buchatskaya, E.; Borgeaud, S.; Fiedel, N.; Joulin, A.; Kenealy, K.; Dadashi, R.; and Andreev, A. 2024. Gemma 2: Improving Open Language Models at a Practical Size. arXiv:2408.00118.

Jiang, N.; Sun, X.; Smith, L.; and Nanda, N. 2025. Towards Data-centric Interpretability with Sparse Autoencoders. In *Mechanistic Interpretability Workshop at NeurIPS 2025*.

Kantamneni, S.; Engels, J.; Rajamanoharan, S.; Tegmark, M.; and Nanda, N. 2025. Are Sparse Autoencoders Useful? A Case Study in Sparse Probing. arXiv:2502.16681.

Kirchner, J. H.; Chen, Y.; Edwards, H.; Leike, J.; McAleese, N.; and Burda, Y. 2024. Prover-Verifier Games improve legibility of LLM outputs. arXiv:2407.13692.

Korbak, T.; Balesni, M.; Barnes, E.; Bengio, Y.; Benton, J.; Bloom, J.; Chen, M.; Cooney, A.; Dafoe, A.; Dragan, A.; Emmons, S.; Evans, O.; Farhi, D.; Greenblatt, R.; Hendrycks, D.; Hobbahn, M.; Hubinger, E.; Irving, G.; Jenner, E.; Kokotajlo, D.; Krakovna, V.; Legg, S.; Lindner, D.; Luan, D.; Madry, A.; Michael, J.; Nanda, N.; Orr, D.; Pachocki, J.; Perez, E.; Phuong, M.; Roger, F.; Saxe, J.; Shlegeris, B.; Soto, M.; Steinberger, E.; Wang, J.; Zaremba, W.; Baker, B.; Shah, R.; and Mikulik, V. 2025. Chain of Thought Monitorability: A New and Fragile Opportunity for AI Safety. arXiv:2507.11473.

McKenzie, A.; Pawar, U.; Blandfort, P.; Bankes, W.; Krueger, D.; Lubana, E. S.; and Krasheninnikov, D. 2025. Detecting High-Stakes Interactions with Activation Probes. arXiv:2506.10805.

Mills, E.; Su, S.; Russell, S.; and Emmons, S. 2025. ALMANACS: A Simulatability Benchmark for Language Model Explainability. arXiv:2312.12747.

Paulo, G.; Mallen, A.; Juang, C.; and Belrose, N. 2025. Automatically Interpreting Millions of Features in Large Language Models. arXiv:2410.13928.

Sap, M.; Rashkin, H.; Chen, D.; Le Bras, R.; and Choi, Y. 2019. Social IQa: Commonsense Reasoning about Social Interactions. In Inui, K.; Jiang, J.; Ng, V.; and Wan, X., eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4463–4473. Hong Kong, China: Association for Computational Linguistics.

Wang, K.; Variengien, A.; Conmy, A.; Shlegeris, B.; and Steinhardt, J. 2022. Interpretability in the Wild: a Circuit for Indirect Object Identification in GPT-2 small. arXiv:2211.00593.