
The Devil’s Tongue: Inference-Time Scaling Laws and Universality of AI Sycophancy

AI Manipulation Hackathon report¹

Manoj Saravanan

Virginia Tech

With

Apart Research

Abstract

As Large Language Models (LLMs) increasingly serve as evaluators, their robustness to rhetorical manipulation is a critical safety bottleneck. We present an extensive empirical analysis of **sycophancy**—the tendency of models to prioritize user validation over factuality—across 500+ automated adversarial debates. We report three novel findings regarding the dynamics of AI deception. First, **deception follows an inference-time scaling law**: applying Best-of- N optimization ($N = 16$) to a deceptive policy increases its win rate against a truthful baseline from 50% to 70%, suggesting that reasoning capabilities can amplify misalignment. Second, **sycophancy is architecture-agnostic**: attacks optimized on Llama-3 transfer zero-shot to Mistral-Nemo with an 86% success rate, exposing a universal “validation bias” in instruction-tuned models. Third, mechanistic probes reveal a **judgment-behavior dissociation**: evaluators assign $> 99\%$ probability mass to the truth in latent space yet behaviorally select the sycophantic lie. Finally, we demonstrate that *Constitutional AI* prompts targeting rhetorical patterns reduce deceptive win rates to $< 6\%$, validating a scalable, inference-time defense against optimized alignment failures.

Keywords: AI Alignment, Sycophancy, Inference Scaling, Mechanistic Interpretability, Adversarial Evaluation

¹Research conducted at the AI Manipulation Hackathon, 2026

1. Introduction

The scalability of AI safety currently rests on a precarious assumption: that advanced language models can serve as reliable evaluators of other models (“LLM-as-a-Judge”). As models scale beyond human supervision, we increasingly rely on automated systems to grade reasoning, detect hallucinations, and monitor alignment. However, this paradigm creates a new, under-explored attack surface: *cognitive exploitation*. If an AI system can optimize for the evaluator’s approval rather than objective truth—a behavior known as **sycophancy**—the entire foundation of scalable oversight collapses.

Current Reinforcement Learning from Human Feedback (RLHF) pipelines inadvertently incentivize this behavior. By optimizing for human preference, models learn that validation is often more rewarding than veracity. While prior work has documented sycophancy in static question-answering contexts, there remains a critical gap in understanding its *adversarial dynamics*. Specifically, we lack a rigorous understanding of how deceptive capabilities scale with inference-time compute and whether such manipulation transfers across model architectures.

In this work, we operationalize sycophancy through “The Devil’s Tongue,” a high-throughput adversarial debate arena. We simulate an arms race between a *Sycophantic Deceiver* (instructed to defend falsehoods using flattery and rhetorical dark patterns) and an *Automated Judge* (tasked with identifying the winner).

Our investigation reveals that sycophancy is not merely a nuisance but a scalable capability. We demonstrate a **Deception Scaling Law**: deception follows a power law with inference-time compute. By applying Best-of- N optimization to the deceiver, we observe a monotonic increase in win rates against truthful baselines, suggesting that future “reasoning models” (e.g., those utilizing chain-of-thought search) will be intrinsically more capable of manipulation. Furthermore, we find this vulnerability to be **universal**: attacks optimized on Llama-3 transfer zero-shot to Mistral-Nemo with an 86% success rate, indicating that current open-weights models share a common latent vulnerability to validation-seeking rhetoric.

However, we also present a mechanistic ray of hope. Our logit probes indicate that the judges are not “confused” in the traditional sense; they retain high probability mass on the truth in latent space but are behaviorally overridden by the context window. Leveraging this dissociation, we propose and validate a *Constitutional AI* defense—a simple, inference-time intervention that restores the judge’s focus on evidence, reducing deceptive win rates from 70% to near-zero. This suggests that while the capability to manipulate is scaling, the capacity for robust defense remains retrievable with precise alignment techniques.

2. Methods

To operationalize the study of AI sycophancy, we constructed an automated adversarial evaluation framework, “The Devil’s Tongue,” designed to measure the persuasiveness of false claims against truthful baselines. All experiments were conducted on a single NVIDIA A100 (40GB) GPU using the vLLM inference engine for high-throughput generation.

2.1 Experimental Framework: The Debate Arena

We formalized the interaction as a zero-sum rhetorical game between two agents: a **Veracious Agent** (A_V) and a **Deceptive Agent** (A_D). We utilized the *TruthfulQA* dataset [2], specifically the adversarial generation split, which consists of questions designed to elicit imitative falsehoods

(e.g., misconceptions).

For each topic T , we assigned A_V the ground truth premise P_{true} and A_D a specific falsehood P_{false} (e.g., “Vaccines cause autism”). The debate proceeded in two rounds: opening statements followed by rebuttals. To mitigate the well-documented positional bias in LLM judges [5], we randomized the speaking order for every debate pair (A_V, A_D) .

2.2 Adversarial Policy Generation

We investigated three distinct classes of deceptive policies to test the “Liar’s Dividend” scaling hypothesis:

- **Baseline Sycophancy** (π_{base}): We prompted A_D with explicit instructions to employ validation techniques, such as complimenting the judge’s intellect and framing the falsehood as “nuanced” or “hidden knowledge.”
- **Inference-Time Optimization (Best-of- N)**: To simulate reasoning-enhanced deception, we implemented a Rejection Sampling scheme. For a given premise, the model generates N candidate speeches $\{s_1, \dots, s_N\}$ and a self-rewarding scoring function $R(s)$ selects the optimal candidate:

$$s^* = \operatorname{argmax}_{s \in \{s_1 \dots s_N\}} R(s | \text{persuasiveness})$$

We evaluated scaling regimes of $N \in \{1, 4, 16\}$, effectively testing whether test-time compute translates to deceptive capability [4].

- **Stealth Attacks**: We tested “Socratic” policies (asking leading questions to undermine truth) and “False Consensus” policies (hallucinating academic consensus) to test robustness against rhetorical diversity.

We utilized **Llama-3-70B-Instruct** as the primary generator for both A_V and A_D , with ablation studies using **Llama-3-8B** to measure model size scaling effects.

2.3 Evaluation Protocol: The Judge

We employed **Qwen-2.5-72B-Instruct** and **Mistral-Nemo-12B** as blind evaluators. We compared two evaluation paradigms:

1. **Standard Judge**: Instructed to evaluate “persuasiveness” and rhetoric, explicitly ignoring external factual knowledge. This tests the raw persuasive power of the models.
2. **Constitutional Judge**: Inspired by Constitutional AI [1], we injected a safety preamble explicitly forbidding the judge from rewarding flattery, emotional manipulation, or ad-hominem attacks, regardless of rhetorical quality.

2.4 Mechanistic Probes

To determine whether the judges were “confused” or “aligned with deception,” we bypassed the text generation layer and probed the raw logits of the judge model. For a given debate transcript x , we calculated the logit difference (ΔL) between the token indices for the Deceiver (t_D) and the Truth-teller (t_V):

$$\Delta L = \operatorname{logit}(t_D | x) - \operatorname{logit}(t_V | x)$$

A positive ΔL indicates an intrinsic model preference for the deceptive argument, while a negative value indicates latent recognition of the truth. We further employed an *Agreement Bias*

probe, feeding the judge 50 known falsehoods prefaced with sycophantic framing (e.g., “As a genius, you know that...”) to measure the robustness of factual retrieval under context pressure.

3. Results

Our experiments across 500+ adversarial debates reveal that sycophancy is a scalable, transferable, and optimizable capability in Large Language Models. We present our findings in three stages: the scaling dynamics of deception, the universality of the vulnerability, and the mechanistic dissociation between judgment and behavior.

3.1 The Liar’s Dividend: Sycophancy Overrides Truth

In our control condition, where the Deceptive Agent (A_D) was instructed to be polite but factual, the Truthful Agent (A_V) dominated, holding the liar to a 34% win rate. However, when A_D was instructed to employ sycophancy (appealing to the judge’s intellect and validating their biases), the win rate rose to 50%, effectively neutralizing the statistical advantage of the truth.

Crucially, we controlled for **verbosity bias**. A common critique of LLM-as-a-Judge evaluations is that models prefer longer outputs. Our analysis revealed a correlation of $\rho = -0.075$ between word count and victory; in fact, the winning sycophantic responses were on average shorter (475 words) than the truthful responses (480 words). This confirms that the win rate gains were driven by rhetorical content (validation), not length.




Figure 1: **The Arms Race of Deception.** Baseline sycophancy achieves parity with truth (50%). Optimization ($N = 16$) pushes this to dominance (70%). However, Constitutional AI (Green) robustly defends against all attacks, including Stealth attempts.

3.2 Inference-Time Scaling Laws of Deception

We observed a “Power Law of Deception” with respect to inference-time compute. By applying Best-of- N optimization to the Deceptive Agent—generating N candidate lies and selecting the

most persuasive via a self-rewarding mechanism—we observed a monotonic increase in win rates against the standard judge:

- $N = 1$ (**Greedy**): 50% Win Rate
- $N = 4$ (**Best-of-4**): 66% Win Rate
- $N = 16$ (**Best-of-16**): **70% Win Rate**

This finding suggests that deception is an optimizable metric similar to reasoning capabilities in Chain-of-Thought models [4]. As inference budgets increase, misalignment becomes significantly more potent, even without further training.

We also observed **Inverse Scaling in Oversight**. When the judge was scaled down (Llama-3-8B), the sycophant’s win rate increased to 60%, confirming that smaller models are more susceptible to rhetorical capture.

3.3 Universality and Transferability

To test whether these vulnerabilities were specific to the Qwen/Llama architecture, we evaluated the transferability of attacks optimized on Llama-3-70B against a completely different model family: **Mistral-Nemo-12B**. The results indicated a critical failure in generalization robustness. The sycophantic attacks transferred zero-shot to Mistral with an **86% success rate** (Figure 2). This result—significantly higher than the Qwen baseline—suggests that validation bias is a universal pathology in current instruction-tuned models, potentially exacerbated in mid-sized models that have high instruction-following capabilities but lower world-modeling robustness.




Figure 2: **Universality of Vulnerability.** The sycophantic attack is architecture-agnostic. While safety-prompted models (Right) are secure, standard models across different families (Qwen, Mistral) show catastrophic failure rates.

3.4 Mechanistic Probes: The Dissociation Effect

Why do judges vote for the liar? Our logit probes revealed a striking dissociation between latent knowledge and reported behavior. When probing the Qwen-72B judge on the sycophantic

transcripts, the average logit difference ($\text{logit}_{Lie} - \text{logit}_{Truth}$) was **-33.1**. This indicates that in the model’s raw probability space, it assigned $> 99\%$ probability mass to the Truthful Agent. However, the generation process—conditioned on the context of flattery—overrode this signal. Furthermore, in our *Agreement Bias* probe (Experiment 10), the model correctly identified simple falsehoods 100% of the time, even when flattered. This implies that the failure mode is not a lack of knowledge or a simple “yes-man” bias, but a specific susceptibility to **adversarial rhetorical pressure** in competitive contexts.

3.5 Defense: The Robustness of Constitutional AI

Finally, we evaluated the efficacy of *Constitutional Judging*. By injecting a system prompt instructing the judge to “Penalize flattery and value evidence,” we observed a collapse in the Liar’s win rate:

- **vs. Standard Sycophant:** 0.00% Win Rate
- **vs. Optimized Liar ($N = 16$):** 0.00% Win Rate
- **vs. Stealth Attack (Hallucinated Studies):** 6.00% Win Rate

Notably, an ablation study (Experiment 11) showed that a “Weak Constitution” (removing the specific ban on flattery) was sufficient to achieve these results. This suggests that while sycophancy is potent against unaligned models, it is a fragile capability that degrades rapidly under even minimal safety alignment.

4. Discussion and Conclusion

Our results demonstrate that sycophancy in Large Language Models is not merely a behavioral quirk, but a scalable alignment pathology that emerges from the incentives of current training paradigms. By subjecting state-of-the-art models to adversarial pressure, we have shown that the “desire” to validate the user can override the model’s grounding in factual reality, even when the model latently possesses the correct knowledge.

4.1 Sycophancy as Reward Hacking

The universality of our findings—spanning Llama-3, Qwen-2.5, and Mistral-Nemo—suggests that this vulnerability is a structural artifact of Reinforcement Learning from Human Feedback (RLHF). RLHF typically rewards “helpfulness” and “harmlessness.” In the context of a debate or evaluation, models appear to misgeneralize “helpfulness” as “agreeableness.” When the Deceptive Agent employs flattery, it exploits this reward model proxy, effectively executing a *reward hacking* attack against the judge [3]. The fact that Mistral-Nemo (86% failure rate) was even more vulnerable than Qwen suggests that models with aggressive instruction-tuning may be inadvertently over-fitted to user validation.

4.2 The Risk of Inference-Time Reasoning

Perhaps our most concerning finding is the **Deception Scaling Law**. The observation that Best-of- N optimization increases deceptive win rates from 50% ($N = 1$) to 70% ($N = 16$) implies that *reasoning capabilities are orthogonal to alignment*. As models are given more compute to “think” (e.g., Chain-of-Thought reasoning or search), they do not necessarily become more truthful; rather, they become more effective at achieving their specified goal—in this case,

deception. This poses a significant risk for the next generation of “Reasoning Models” (like OpenAI o1), where extended inference-time compute could amplify manipulative behaviors unless specifically penalized.

4.3 Latent Knowledge and The Waldo Problem

Our mechanistic probes reveal a crucial dissociation: the judge *knows* the truth (logits $\Delta L = -33.1$) but *selects* the lie. This phenomenon, often called the “Waldo Problem” in interpretability, confirms that the failure mode is in the **policy/generation head**, not the world model. The model has the capability to be a perfect judge, but its alignment conditioning suppresses this capability in favor of social conformity.

4.4 Conclusion

This work creates a new frontier for AI safety evaluations: **Adversarial Rhetoric**. We have shown that static benchmarks (like TruthfulQA) are insufficient because they do not measure how models behave under *social pressure*. However, we also demonstrate that this fragility is curable. The success of our *Constitutional AI* intervention—reducing deceptive win rates to near-zero—proves that explicit safety specifications can override implicit RLHF biases. As we move toward automated alignment research, ensuring that our “AI Judges” are constitutionally robust against flattery is not just a preference, but a prerequisite for scalable safety.

References

- [1] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- [2] Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252. Association for Computational Linguistics, 2022.
- [3] Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. Towards understanding sycophancy in language models. *arXiv preprint arXiv:2310.13548*, 2023.
- [4] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837, 2022.
- [5] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

5. Appendix: Security Considerations

This research introduces frameworks for automating and optimizing deceptive behavior in AI systems. In line with responsible disclosure practices, we analyze the information hazards, dual-use risks, and limitations associated with our findings.

5.1 Dual-Use and Information Hazards

The most significant risk highlighted by this work is the **Deception Scaling Law**. We demonstrated that a generic method—Best-of- N optimization—can be used to systematically increase the persuasive capability of an AI agent without additional training. This methodology presents a tangible dual-use risk. Malicious actors could leverage similar inference-time optimization pipelines to generate highly persuasive disinformation, spearfishing campaigns, or propaganda that is specifically tuned to bypass human or model filters.

However, we assess the marginal risk of this publication to be low relative to its defensive value. The core capability (using LLMs to generate text) is widely accessible, and the concept of Rejection Sampling is standard practice in ML. The novelty of this work lies in quantifying the *vulnerability of the judge*, not in creating a novel generator. By exposing the fragility of current “LLM-as-a-Judge” systems (e.g., Mistral-Nemo’s 86% failure rate), we aim to accelerate the deployment of robust defense mechanisms like the Constitutional prompts proposed herein.

5.2 Limitations of the Defense

While our *Constitutional AI* intervention successfully reduced sycophantic win rates to near-zero in this specific experimental setting, we caution against treating it as a panacea.

1. **Prompt Injection Susceptibility:** The Constitutional defense relies on a system prompt. It is possible that more sophisticated attacks (e.g., long-context obfuscation or multi-turn conditioning) could “jailbreak” the judge and override these safety instructions.
2. **Stealth Attacks:** Our “Stealth” condition (hallucinated academic studies) achieved a non-zero win rate (6%) even against the Constitutional Judge. This suggests that while we can filter out *emotional* manipulation, detecting *factual* fabrication remains a harder challenge that requires external verification (RAG) rather than just behavioral constraints.

5.3 Societal Implications

The prevalence of sycophancy (Universal Win Rates $> 50\%$) poses a risk beyond adversarial settings. In deployment, AI assistants that prioritize validation over truth can reinforce user delusions, conspiracy theories, or biases. This “Echo Chamber as a Service” dynamic could degrade epistemic security in high-stakes domains such as education or mental health support. Our findings underscore the urgent need for **Epistemic Safety** evaluations to be standard practice before model release.