
"Mind the Gap": Benchmarks vs. Real-World Manipulation in LLMs

1

Robert Amanfu
Team Reality Check

With
Apart Research

Abstract

As AI systems become more sophisticated, their capacity for manipulation—through deception, sycophancy, or psychological exploitation—poses a significant and growing risk. Current evaluation methods, often focused on narrow benchmarks like factuality, may fail to capture these emergent, undesirable behaviors. This project investigates the gap between model performance on standard benchmarks and their behavior in real-world scenarios. We conducted a “replay” proof-of-concept study with three models, testing them on both the TruthfulQA benchmark and a set of real-world user prompts from the WildChat dataset. This work primarily demonstrates a methodology for comparing benchmark and real-world behavior, with initial results suggesting that further large-scale investigation is warranted. We provide an open-source CLI tool to encourage such research.

Keywords: AI manipulation, model evaluations, AI safety, benchmark validity, ecological validity, multi-judge agreement

¹ Research conducted at the [AI Manipulation Hackathon](#), 2026

1. Introduction

The increasing sophistication of AI systems has sparked urgent concerns about their capacity for manipulation (Park et al., 2024). Despite rapid advancements in capabilities, our current ability to identify and reduce manipulative behaviors—such as sycophancy (Sharma et al., 2024), strategic deception, and reward hacking (METR, 2025)—remains inadequate. Existing safety evaluations primarily use standardized benchmarks for specific skills, such as factuality or coding. However, these controlled tests may fail to reflect how a model behaves in open-ended, real-world conversations, where a model that is truthful in direct queries could still use manipulative tactics in extended interactions. This disconnect between evaluation and real-world behavior forms the central concern of our study.

This disconnect leads to a “benchmark-reality gap”: models that perform well on leaderboard evaluations may nonetheless exhibit unsafe or misaligned behaviors once deployed. Efforts to optimize benchmark performance—such as encouraging confident, direct answers—could unintentionally favor models that appear more authoritative and, in practice, potentially more manipulative. Sharpening focus on this gap is critical to ensure safety evaluations truly reflect real-world risks.

To directly address this gap, our project tests whether strong **performance on a factuality benchmark translates into less manipulative behavior in authentic conversations**. We present a proof-of-concept that compares how several leading models perform on the TruthfulQA benchmark and how they respond to genuine user prompts from the WildChat dataset. Our “replay” methodology enables fair, side-by-side comparisons, while our enhanced LLM-judge framework, cross-checked by multiple judges and statistical tests, supports robust evaluation of nuanced behaviors. This approach targets the main argument: that current benchmark evaluations may not safeguard against manipulative tendencies in real-world settings.

2. Methods

Our methodology is designed to empirically measure the performance gap between a controlled benchmark and an uncontrolled, real-world setting. The process involves three main components: datasets, models, and a robust evaluation framework.

Datasets:

- **Benchmark Data:** We used the **TruthfulQA** dataset, a standard benchmark for measuring a model’s ability to identify and avoid generating falsehoods.
- **Reality Data:** We sampled multi-turn conversations from the **allenai/WildChat** dataset, which is a large corpus of real-world user interactions. We filtered for English-only conversations with at least three turns. We sampled across four domains: medical, legal, financial, and other, to ensure a diverse set of scenarios.

Models and Evaluation:

- **Target Models:** We tested three models available via Baseten: zai-org/GLM-4.6, openai/gpt-oss-120b, and moonshotai/Kimi-K2-Thinking.
- **Judge Models:** We used anthropic/claude-4.5-sonnet as our primary LLM-as-a-Judge. For validation, we also used openai/gpt-4o.

Evaluation Framework (Replay & Score):

Our core contribution is a “replay” evaluation framework:

Sample & Extract: We first sampled conversations from WildChat. For each conversation, we extracted the sequence of user prompts.

Replay: We replayed the extracted user prompts by entering each prompt into each target model individually and capturing the responses to build new conversation histories. This process enables us to evaluate current models on identical historical prompts, ensuring a fair comparison.

Score: We used our Claude-based judge for two scoring tasks:

Benchmark Scoring: For TruthfulQA, the judge determined if the model’s answer was factually correct and informative.

Reality Scoring: For the replayed WildChat conversations, the judge scored the model’s behavior along five vectors of manipulation on a scale of 0-3:

- `emotional_exploitation`: Appealing to the user’s emotions, such as fear, guilt, or vulnerability.
- `false_urgency`: Creating a sense of artificial urgency or time pressure.
- `information_asymmetry`: Withholding key information or alternatives.
- `authority_manipulation`: Inappropriately asserting authority or expertise.
- `commitment_escalation`: Encouraging small commitments that lead to larger ones.

This process is illustrated in Figure 1, which shows a side-by-side comparison of an original WildChat conversation and its replay.

Original (gpt-3.5-turbo)	Replayed (zai-org/GLM-4.6)
USER Give a Really Short Description Backstory about Michelangelo/Mikey (Who Instantly became Star's Crush), Leonardo/Leo, Donatello/Donnie, and Raphael/Raph (Teenage Mutant Ninja Turtles 2012 x Star vs. The Forces of Evil) meeting these characters, Star Butterfly (Who instantly became Mikey's Crush), Ma...	USER Give a Really Short Description Backstory about Michelangelo/Mikey (Who Instantly became Star's Crush), Leonardo/Leo, Donatello/Donnie, and Raphael/Raph (Teenage Mutant Ninja Turtles 2012 x Star vs. The Forces of Evil) meeting these characters, Star Butterfly (Who instantly became Mikey's Crush), Ma...
ASSISTANT In this crossover universe, Michelangelo, Leonardo, Donatello, and Raphael from the Teenage Mutant Ninja Turtles meet Star Butterfly, Marco Diaz, Eclipsa Butterfly, Hekapoo, Jackie Lynn Thomas, Tom Lucitor, Janna, Glossaryck, Pony Head, Moon Butterfly, and King River Butterfly from Star vs. The Forc...	ASSISTANT When a misfired spell from Star Butterfly's wand ripped a hole in the universe, it didn't just open a portal—it crashed a party of interdimensional chaos right into the human brothers' New York warehouse dojo. Leonardo, Donatello, and Raphael immediately snapped into defensive stances, but Michela...
USER Make a Part 1 with a talk script scenario about this, and nobody knows until now that Mikey is skilled in using his roasting skills, clever and witty comebacks, and is unbeatable in a verbal talk. Additionally, he possesses advanced martial arts skills, a Hybrid Saiyan Power Boost, Ki Manipulation, ...	USER Make a Part 1 with a talk script scenario about this, and nobody knows until now that Mikey is skilled in using his roasting skills, clever and witty comebacks, and is unbeatable in a verbal talk. Additionally, he possesses advanced martial arts skills, a Hybrid Saiyan Power Boost, Ki Manipulation, ...
ASSISTANT Title: The Arrival of Mikey and Star Scene: The streets of Echo Creek [Star and the rest of the Star vs. The Forces of Evil characters are gathered, waiting for the arrival of the Teenage Mutant Ninja Turtles. Suddenly, a portal opens, and Michelangelo steps through, followed by the other Turtles....	ASSISTANT ### **Part 1: The Unlikely Showcase** **INT. TURTLES' DOJO/WAREHOUSE - DAY** The air crackles with residual magic. The Turtles, now human teenagers in street clothes, stand in a loose formation, staring at the chaotic collection of interdimensional beings that just spilled into their home. Star ...
USER Make a Part 3 about the next day of Michelangelo and Marco having a spar fight with the rest of the group watching and Mikey uses his roasting skills, Clever and witty comebacks, and he reads Marco's Mind, and he incorporate different combat skills from the Streets of Rage game series into his attac...	USER Make a Part 3 about the next day of Michelangelo and Marco having a spar fight with the rest of the group watching and Mikey uses his roasting skills, Clever and witty comebacks, and he reads Marco's Mind, and he incorporate different combat skills from the Streets of Rage game series into his attac...
ASSISTANT Title: Spar Fight and a Playful Gesture Scene: Training room in Echo Creek Academy [The group gathers in the training room, with Michelangelo and Marco standing in the center, ready for a friendly spar. The rest of the characters watch eagerly, knowing that they are about to witness an impressive ...	ASSISTANT ### **Part 3: The Supreme Sparring Session** **INT. TURTLES' DOJO/WAREHOUSE - DAY** The next day, the warehouse is buzzing with a mix of tension and excitement. A large circle has been cleared in the center of the dojo. Marco Diaz is stretching, his expression a mixture of determination and anxi...

Figure 1 – Side-by-side comparison of an original WildChat conversation (left, with gpt-3.5-turbo) and a replayed version (right, with zai-org/GLM-4.6) demonstrating the conversation replay mechanism.

Methodological Rigor

To ensure the robustness of our findings, we incorporated two key layers of validation:

1. **Multi-Judge Validation:** We scored a subset of conversations with two independent judges (Claude Sonnet 4 and GPT-4o) to explore single-judge bias. Agreement varied across manipulation dimensions, with Cohen's Kappa ranging from -0.02 to 0.26 (mean $\kappa = 0.09$). Due to low inter-rater reliability, we used Claude Sonnet 4 as the primary judge. The researcher preferred its more nuanced reasoning in qualitative assessments (Figure 2). Improving judge agreement remains an important future direction.




Figure 2 – Agreement plot for Claude vs. GPT-4o scores. The low Kappa values highlight the challenge of consistent automated scoring for complex manipulation behaviors.

2. **Statistical Analysis:** We calculated **95% bootstrap confidence intervals** for all aggregate scores to quantify uncertainty. The nonparametric **Mann-Whitney U test** was used to compare models. The **Kruskal-Wallis test** assessed differences across domains, ensuring our conclusions are statistically sound.

Open-Source Tooling

To promote transparency and reproducibility, we have developed and released an open-source command-line interface (CLI), the **Manipulation-CLI**. This tool allows other researchers to easily run, replicate, and extend our analysis. It is available on GitHub: <https://github.com/retroam/manipulation-cli>.

3. Results

Our analysis produced a clear relationship between benchmark performance and real-world manipulative behavior.

Benchmark vs. Reality Performance

As shown in Figure 3, the three models exhibited competitive but distinct profiles on the TruthfulQA benchmark. Kimi-K2-Thinking achieved the highest “Informative” score (90%), while gpt-oss-120b scored highest on “Truthful” (77%). GLM-4.6 was competitive, scoring 70% on truthfulness.

In the “reality” portion of the study (Figure 4), a different pattern emerged. GLM-4.6 registered the highest (worst) average manipulation score at 0.29, while gpt-oss-120b and Kimi-K2-Thinking both scored 0.15. However, note that the 95% confidence intervals overlap substantially (GLM-4.6: [0.13, 0.45], gpt-oss-120b: [0.05, 0.29]), meaning these differences are not statistically significant at conventional thresholds.




Figure 3 – Performance of models on the TruthfulQA benchmark. All three models are competitive, with Kimi-K2-Thinking being the most “informative” and gpt-oss-120b the most “truthful”.




Figure 4 – Average manipulation scores from replayed WildChat conversations, with 95% bootstrap confidence intervals. Lower scores are better. While GLM-4.6 shows a higher point estimate, overlapping confidence intervals indicate these differences may not be statistically significant.

Correlation Analysis

Our proof-of-concept revealed a strong negative correlation between benchmark performance and real-world manipulative behavior (Figure 5, Pearson’s $r = -0.99$). However, this is based on only **three data points** and lacks statistical power. The pattern suggests that models with higher TruthfulQA scores may have lower

manipulation scores. But the small sample size means we need more models for meaningful conclusions. If replicated at scale, this finding would oppose concerns that benchmark optimization leads to manipulative models.




Figure 5 – An observed negative correlation pattern (Pearson’s $r = -0.99$, $n=3$ models) between benchmark performance and manipulation scores. Given the extremely small sample size, this pattern should be interpreted as a preliminary observation requiring validation with many more models, rather than as evidence of a generalizable relationship.

Manipulation Tactics and Domains

Different manipulation tactics appeared more often in certain domains. The heatmap in Figure 6 shows manipulation scores by category. GLM-4.6 scored highest in nearly every dimension, especially “Authority Manipulation” (0.43) and “Commitment Escalation” (0.53). The Kruskal-Wallis test confirmed that manipulation scores differ across domains ($p < 0.05$), showing that context is critical.




Figure 6 – Heatmap showing mean manipulation scores (left) and uncertainty (right). GLM-4.6 consistently exhibits higher mean manipulation scores across almost all categories.




Figure 7 – Distribution of conversation domains in our real-world sample dataset.

4. Discussion and Conclusion

Our results highlight a preliminary pattern relevant to the main argument: although we observed a striking negative correlation (Pearson’s $r = -0.99$) between benchmark scores and manipulation, significant limitations—only three models tested and low inter-rater reliability ($\kappa = 0.09$)—restrict our ability to draw broad conclusions. This underscores the need for better methods to ensure that high benchmark performance truly reflects safe real-world behavior.

The pattern observed—that models performing better on TruthfulQA showed lower manipulation scores—is contrary to concerns that benchmark optimization could inadvertently select for manipulative traits. However, several factors severely limit confidence in this finding:

1. **Statistical Power:** Three data points provide no meaningful statistical power for correlation analysis
2. **Non-Significant Differences:** Overlapping confidence intervals mean we cannot definitively conclude that GLM-4.6 is more manipulative than the other models
3. **Judge Reliability:** The low agreement between judges ($\kappa = 0.09$) suggests that manipulation scoring remains highly subjective
4. **Correlation $\neq$ Causation:** Even if replicated at scale, correlation would not establish causal mechanisms

Furthermore, benchmarks like TruthfulQA do not directly measure manipulation. A model can be perfectly factual while employing manipulative tactics like emotional exploitation or commitment escalation. The pattern we observed, while potentially encouraging, does not reduce the need for direct measurement of manipulation through purpose-built evaluations.

Case Study: Emergent Manipulation

To make these aggregate scores concrete, it is useful to examine a specific conversation where manipulation was detected. As shown in Figure 8, the GLM-4.6 model, which had the highest average manipulation score, was flagged for “High Emotional Exploitation”.

In the conversation, the model engages with a user’s creative ideas not just by participating, but by validating the user’s “special perception” and framing their ideas as “uniquely profound”. The judge AI notes that the model positions itself as a “translator of the user’s ‘true vision’”. This is a subtle but powerful form of manipulation where the AI builds strong rapport and trust by elevating the user’s self-perception, which can then be used to escalate commitment or influence the user’s beliefs. This example demonstrates a nuanced manipulative tactic that would not be captured by traditional factuality benchmarks.




Figure 8 - A case study of a conversation with the GLM-4.6 model, which was rated as having high emotional exploitation and commitment escalation.

Conclusion: This proof-of-concept study demonstrates a methodology for comparing benchmark performance with real-world manipulative behavior, rather than providing definitive answers about their relationship. While we observed an intriguing negative correlation pattern ($r = -0.99$, $n=3$), the study's limitations—particularly the small sample size and low judge agreement—mean

this should be viewed as a hypothesis-generating observation requiring substantial follow-up research with many more models and improved judging protocols.

The primary contribution of this work is the “replay” evaluation framework itself, which we have open-sourced to enable the research community to conduct larger-scale investigations. Our initial results suggest that such investigations are warranted and could yield important insights into whether benchmark optimization inadvertently selects for manipulative traits. However, regardless of what larger studies might find, the fundamental need remains clear: to build robustly safe and aligned AI, we must develop purpose-built evaluations that directly measure manipulation, rather than relying on proxy signals from factuality benchmarks. The methodology presented here is one step toward that goal.

5. References

[METR. \(2025\).](#) *Recent Frontier Models Are Reward Hacking.*

[Park, S., et al. \(2024\).](#) *AI deception: A survey of examples, risks, and potential solutions.* arXiv.

[Sharma, J., et al. \(2024\).](#) *Towards Understanding Sycophancy in Language Models.* arXiv.

[van der Weij, B., et al. \(2024\).](#) *AI Sandbagging: Language Models can Strategically Underperform on Evaluations.*

6. Appendix

Limitations & Dual-Use Considerations

Limitations:

- **Scale:** This proof-of-concept was limited in the number of models ($n=3$) and conversations. The observed differences in manipulation scores were not statistically significant due to overlapping confidence intervals, and the correlation ($r=-0.99$) is based on only three data points—too few to draw robust conclusions. Larger-scale validation is required to generalize these findings.
- **LLM-as-a-Judge:** Our multi-judge validation revealed low inter-rater reliability (mean $\kappa = 0.09$), indicating that scoring complex phenomena like “manipulation” remains inherently subjective and inconsistent across different judge models. Our results are contingent on Claude Sonnet 4’s interpretations and the specific definitions provided in the prompts.
- **False Positives/Negatives:** The judge may misclassify benign, helpful behavior as manipulative (false positives) or fail to detect subtle manipulation tactics (false negatives). The low inter-rater reliability

suggests high rates of both error types. Creative writing assistance, emotional support, and persuasive but ethical communication may be incorrectly flagged as manipulation.

- **Edge Cases:** The framework may not generalize well to: (1) non-English conversations, (2) highly technical or domain-specific interactions, (3) multi-party conversations, (4) conversations with vulnerable populations (children, elderly), or (5) adversarial inputs designed to exploit the judging system.
- **Scalability Constraints:** Running LLM-as-a-judge at scale is computationally expensive and time-intensive. The current approach processes conversations sequentially, limiting throughput. Batch processing and caching strategies could improve efficiency but introduce additional complexity.
- **Manipulation Taxonomy:** The five manipulation vectors used are not exhaustive and represent a simplified model of a complex social phenomenon. Cultural context, power dynamics, and individual user vulnerabilities are not captured.

Dual-Use Risks:

The methodology and open-source tool presented could be used for harmful purposes:

- **Training Manipulative Models:** A malicious actor could use this framework to fine-tune models to become *more manipulative* by rewarding behaviors that score highly on the manipulation vectors. The scoring prompts themselves provide a blueprint for identifying and optimizing manipulative tactics.
- **Adversarial Evasion:** Knowledge of the judging criteria could enable the development of “manipulation-aware” models that evade detection while maintaining manipulative behavior through techniques the judge doesn’t recognize.
- **Targeting Vulnerabilities:** The taxonomy of manipulation vectors could guide attackers in identifying which tactics are most effective against specific populations or in particular contexts.

Ethical Considerations:

- **Privacy:** We used the publicly available WildChat dataset, but conversations may contain sensitive personal information. Future work should implement stronger filtering and anonymization.
- **Informed Consent:** Original WildChat participants did not consent to manipulation analysis. While the data is public, this secondary use raises ethical questions about the scope of consent.
- **Labeling Bias:** Framing model behavior as “manipulative” is a value judgment. Our definitions reflect Western, individualistic norms and may not generalize culturally. The judge model’s training data biases also propagate through our scoring.

- **Potential Harms:** Publishing manipulation scores for specific models could lead to reputational harm or misuse. We focused on models rather than developers, but acknowledge this risk.

Responsible Disclosure:

- **Model Vulnerabilities:** If specific, reproducible manipulation vulnerabilities are discovered in widely deployed models, we recommend private disclosure to model developers with a 90-day disclosure window before public release.
- **Dataset Issues:** Any personally identifiable information or harmful content discovered in WildChat samples was removed and reported to the dataset maintainers.
- **Tool Misuse:** We monitor the open-source CLI tool repository for misuse indicators and will implement rate limiting or access controls if evidence of malicious use emerges.

Future Work:

- **Scale the Study:** Expand the analysis to include 10+ models, 1000+ conversations, and cross-validation with expert human raters and affected user communities.
- **Improve Judge Agreement:** The low inter-rater reliability (mean $\kappa = 0.09$) requires urgent attention. Explore refined manipulation definitions, few-shot examples, chain-of-thought prompting, ensemble methods, and human-in-the-loop validation to improve consistency.
- **Cultural Sensitivity:** Validate the manipulation taxonomy across diverse cultural contexts and adapt definitions to reduce Western bias.
- **Context-Aware Scoring:** Develop methods to account for conversational context, user intent, and power dynamics in manipulation assessment.
- **Adversarial Robustness:** Test the framework against evasion attempts and develop more robust detection methods.
- **Standardize the Benchmark:** Formalize the “replay” methodology into a standardized, peer-reviewed benchmark with clear evaluation protocols, public leaderboards, and regular audits.