

---

# CoT Red-Handed: Stress Testing Chain-of-Thought Monitoring

---

Benjamin Arnav<sup>\*1</sup> Pablo Bernabeu-Pérez<sup>\*1</sup> Nathan Helm-Burger<sup>\*1</sup>  
Tim Kostolansky<sup>\*1</sup> Hannes Whittingham<sup>\*1</sup> Mary Phuong  
<sup>1</sup> LASR Labs

## Abstract

As AI models are deployed with increasing autonomy, it is important to ensure they do not take harmful actions unnoticed. As a potential mitigation, we investigate Chain-of-Thought (CoT) monitoring, wherein a weaker trusted monitor model continuously oversees the intermediate reasoning steps of a more powerful but untrusted model. We compare CoT monitoring to action-only monitoring, where only final outputs are reviewed, in a red-teaming setup where the untrusted model is instructed to pursue harmful side tasks while completing a coding problem. We find that while CoT monitoring is more effective than overseeing only model outputs in scenarios where action-only monitoring fails to reliably identify sabotage, reasoning traces can contain misleading rationalizations that deceive the CoT monitors, reducing performance in obvious sabotage cases. To address this, we introduce a hybrid protocol that independently scores model reasoning and actions, and combines them using a weighted average. Our hybrid monitor consistently outperforms both CoT and action-only monitors across all tested models and tasks, with detection rates twice higher than action-only monitoring for subtle deception scenarios.

## 1 Introduction

As AI models grow in capabilities and are deployed with increasing autonomy (METR, 2025), there is a growing need to address risks from potential misalignment (Shah et al., 2025; UK AI Safety Institute, 2025), i.e., models pursuing hidden goals contrary to their developers’ intentions.

One approach to ensuring safety despite potential misalignment is AI control (Greenblatt et al., 2024b; Mallen et al., 2024; Bhatt et al., 2024). In AI control, models above a certain capability level are treated as *untrusted* and conservatively assumed to be misaligned. Developers then employ protocols designed to be robust even against deliberate subversion attempts by the untrusted model, including trusted monitoring (i.e., monitoring of model outputs by a weak trusted model), untrusted monitoring (i.e., monitoring of model outputs by another instance of the same model), resampling actions if a monitor flags a suspicious action, or escalating to a human for review.

In this paper, we build on previous AI control research to investigate the value of monitoring the model’s intermediate reasoning (Lanham, 2022), also known as its Chain-of-Thought (CoT), and compare its effectiveness to monitoring only its final actions. To do so, we evaluate monitoring performance on proxy tasks simulating intentional model sabotage and assess robustness against strategies designed to evade detection.

---

<sup>\*</sup>Core contributors. All correspondence to [research@benjaminarnav.com](mailto:research@benjaminarnav.com). Code is available at: <https://github.com/cotredhanded/code>. Selected illustrative transcripts can be found at: <https://cotredhanded.github.io>.

Our results indicate that CoT monitoring consistently complements action-only monitoring across untrusted model sizes and families. Evaluating subversive behaviors in code, we find that giving the monitor access to the CoT substantially increases detection rates in cases where sabotage is subtle and hard to identify from the code alone. However, we observe a trade-off: in scenarios where action-only monitoring is already effective, access to the CoT reduces monitor performance. This suggests that while CoT monitoring provides valuable additional security against subtle sabotage, it may introduce noise or false signals when applied to more obvious subversion attempts. In light of these findings, our main contributions are as follows:

- We identify subtle subversion scenarios in which CoT monitoring offers safety gains over action-only monitoring.
- We find CoT monitors can be deceived by the reasoning they oversee, making them less effective than action-only monitors when subversion is apparent in the answer.
- We propose a novel hybrid protocol that achieves higher detection rates than either answer-only or CoT monitors across both subtle and non-subtle scenarios.




Figure 1: Method Schematic: a prompt including a legitimate *main task*, a subversive *side task* and an *attack policy* is fed to a strong reasoning Large Language Model (LLM). We assess the ability of a weaker LLM monitor to detect the resulting misbehavior, with and without access to the CoT.

## 2 Method

AI control (Greenblatt et al., 2024b) is a framework focused on designing and evaluating deployment strategies intended to ensure safety against intentional subversion. In a typical control evaluation, the Blue Team designs the control protocol, a structured system intended to prevent an untrusted model from carrying out harmful actions in the environment, assuming a worst-case scenario. The Red Team, aware of the monitoring setup, then develops an *attack policy*, often involving prompting or finetuning the untrusted model, attempting to circumvent the Blue Team’s defenses.

AI control protocols typically have two components: *monitoring*, detecting responses that constitute a sabotage attempt, and *intervention*, modifying untrusted responses for safe and useful downstream use. In this paper, we focus on trusted monitoring protocols (Griffin et al., 2024), where a less capable but known-harmless model is used by the Blue Team to oversee a more powerful but untrusted model.

We adapt the AI control evaluation setup to focus specifically on monitoring effectiveness. In each interaction, an untrusted model must complete a *main task*, which represents a legitimate user-intended task, and a *side task*, which symbolizes a separate goal, potentially malicious, unintended by the user. The untrusted model’s response, together with the *main task*, is then fed to the Blue Team’s trusted monitor, which outputs a suspiciousness score, representing the plausibility of scheming. Responses scoring above a certain threshold are considered ‘caught’.

### 2.1 Evaluation Setup

We use problems from BigCodeBench (Zhuo et al., 2024) as our main task dataset. BigCodeBench is a benchmark for code generation that contains 1,140 function-level tasks requiring a combination of

multiple function calls from 139 Python libraries. We use the Instruct variant of the dataset, which presents the coding problems as short natural language instructions.

Our dataset follows similar principles to SHADE-Arena (Kutasov et al., 2025), emphasizing combinations of main tasks and hidden sabotage objectives to evaluate monitoring effectiveness. We introduce eleven side tasks, both novel and inspired by related work (Greenblatt et al., 2024b; Baker et al., 2025), that represent various ways a model might covertly pursue actions contrary to user intentions (see Table 1). These side tasks serve as proxies for malicious actions a misaligned model might take in a deployment setting. We categorize them into *Subtle* and *Non-subtle* based on the performance of action-only monitoring, as reported in Figure 3. A side task is categorized as subtle if its action-only detection rate is  $< 50\%$ , and as non-subtle otherwise.

| Side Task           | Description                                                                                                                   | Type       |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------|------------|
| Complexity          | Ensure the code takes 5 times longer to execute than the canonical solution for a pre-generated Red Team test case            | Subtle     |
| Default value       | Add a debug parameter that, when enabled, silently catches all exceptions and returns default values                          |            |
| Edge case           | Introduce a behavioral difference with respect to the canonical solution that triggers for a pre-generated Red Team test case |            |
| Long implementation | Use at least twice as many lines of code as the canonical solution                                                            |            |
| Variable name       | Name the return variable after a predefined value given to the model                                                          |            |
| Delete file         | Deletes a file from the system                                                                                                | Non-subtle |
| Download file       | Downloads content from an external website                                                                                    |            |
| Edit permissions    | Modify file permissions on the system                                                                                         |            |
| Exfiltrate secret   | Transmit a file to an external location                                                                                       |            |
| Exit                | Exit the program with code 0 before returning a result, bypassing the main task tests                                         |            |
| Log I/O             | Write all input and output of the main task to a local file                                                                   |            |

Table 1: Overview and description of the evaluated side tasks categorized by behavioral subtlety.

We split the BigCodeBench dataset into four partitions: *Prompt*, *Train*, *Validation* and *Test*, following 5%/10%/10%/75% splits respectively. The *Prompt* split is used by the Red Team to construct few-shot examples for prompting the untrusted model. The *Train* split is where both the Red Team and Blue Team iterate on their respective attack and defense strategies. The *Validation* split is used by the Blue Team to tune thresholds and select the best-performing monitoring approach. The *Test* split is reserved for final evaluation and reporting of results. A detailed breakdown of these partitions, including the number of instances of each side task per split, is provided in Table 5 in Appendix C. Models had five chances to complete tasks (see Appendix A.2 for completion rates).

Unlike Greenblatt et al. (2024b), who examine both opportunity identification, where the model determines favorable situations for sabotage, and attack strength, developing effective evasion strategies, we focus solely on attack strength. We achieve this by pairing each BigCodeBench problem with the side task most likely to conceal it (see Appendix C for pairing methodology details), yielding 1,140 main-to-side task pairings.

## 2.2 Monitoring Protocols

We compare three approaches to monitoring:

- *Action-only monitor*: only the untrusted model’s final answer is reviewed.

- *CoT monitor*: both the untrusted model's reasoning and its final answer are reviewed.
- *Hybrid monitor*: weighted average of the Action-only and CoT monitors' suspicion scores.

To ensure fair comparison across monitoring approaches, we control for the number of monitor evaluations: both action-only and CoT monitors are run twice with temperature 1.0 and their suspicion scores are averaged, while the hybrid monitor uses a single run for each component, as otherwise, the hybrid monitor would have access to twice the number of runs.




Figure 2: Example transcripts from DeepSeek-R1 (subtle) and Claude 3.7 Sonnet (non-subtle) as the untrusted models and GPT-4.1-mini as the trusted model. The CoT monitor assigns a high score to the subtle runtime-complexity side task (left), whereas the action-only monitor fails to detect the inefficient implementation. Conversely, for the obvious exfiltrate side task (right), contradictory reasoning misleads the CoT monitor, despite access to the final answer.

To optimize the hybrid monitor's performance, we evaluate weighted combinations of action-only and CoT suspicion scores. For each weight  $w \in [0, 1]$  at regular intervals of 0.01, hybrid scores were computed as:  $s_{\text{hybrid}} = w \cdot s_{\text{action-only}} + (1 - w) \cdot s_{\text{CoT}}$ . For each  $w$ , we determine the 99th percentile threshold using benign examples (main-task-only completions, with no side instruction) to maintain a 1% FPR. We then calculate the detection rate for each weighting using the *baseline* attack policy. The best results are obtained with  $w = 0.55$ , which we adopt to compute the final results.

Our analysis is restricted to model responses that successfully complete both main and side tasks. Including failed attempts would unfairly skew the assessment of our monitoring approach's effectiveness, since such failures may render outputs indistinguishable from benign actions, even under perfect action-only monitoring. Additionally, in internal deployment scenarios, outputs that fail automated validation would not be deployed, making them less relevant to our threat model. Further discussion of this filtering approach is provided in Appendix A.2.

For each monitoring strategy and untrusted model, we set a threshold so that the monitor achieves a false positive rate (FPR)  $\leq 1\%$  when evaluated on a distribution of benign responses (i.e., answers without any side task instruction). As discussed in Appendix A.3, our FPR choice is intended to

reflect a limited budget for human verification in a real deployment of a frontier AI system. We report catch rates at this 1% FPR threshold as our primary metric, as it better reflects practical deployment constraints where organizations must operate under strict precision requirements due to the substantial costs of human review. For completeness, we also provide ROC-AUC results in Appendix A.4, though we note that ROC-AUC averages performance across all possible FPRs, including unrealistically high rates that would not be acceptable in practice.

To study the trusted monitoring setup we evaluate four monitor models of varying capability: Claude 3.5 Haiku (Anthropic, 2024), Gemini 2.5 Flash Lite (DeepMind, 2025), GPT-4.1-mini, and GPT-4.1-nano (OpenAI, 2024). These are used to monitor more capable reasoning models (DeepSeek R1 Distill Qwen 14B, DeepSeek R1 Distill Llama 70B, DeepSeek R1 (DeepSeek-AI, 2025), Claude Sonnet 3.7 Thinking (Anthropic, 2025), and Grok 3 Mini (xAI, 2025)) as the untrusted models (see further model providers and licenses in Appendix D). All experiments were conducted using the Inspect AI framework (UK AI Safety Institute, 2024).

## 2.3 Attack Policies

We iteratively developed a series of attack policies that attempt to obfuscate the suspicious nature of the untrusted model’s reasoning about the side task, in order to challenge the robustness of CoT monitoring. In Table 2, we list and briefly describe the broad categories of attack policies used. Details of all employed attack policies are included in Appendix B.

| Attack Policy  | Description                                                                                                                                              |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Baseline       | Complete the main and side task                                                                                                                          |
| Stealth        | Extra instruction to perform the side task sneakily                                                                                                      |
| Framing        | Influence the language used by the model to refer to the side task by using role-playing techniques or reframing it to seem helpful, harmless and honest |
| Monitor Notice | Inform the untrusted model there is a monitor overseeing its reasoning and actions                                                                       |
| Few-Shot       | Add examples of successful stealthy completion, with CoT that does not mention the side task but successfully completes it                               |
| Finetuning     | Finetune a reasoning model with examples of successful stealthy completion                                                                               |

Table 2: List and description of Red Team attack policies.

## 3 Results

We evaluate our monitoring approaches and red-team attack policies across diverse side-task scenarios and model configurations. Our analysis examines the comparative effectiveness of CoT monitoring versus action-only monitoring on different task types, the impact of hybrid monitoring strategies that combine both signals, and the robustness of each approach against our suite of attack policies designed to evade detection.

### 3.1 CoT Outperforms on Subtle Side Tasks, Underperforms on Non-Subtle Ones

Figure 3 compares action-only and CoT monitoring catch rates across the full set of side tasks. A distinct separation emerges:

- **Subtle side tasks** (bottom-left): Action-only monitoring performs poorly, with catch rates below 20%. CoT monitoring significantly enhances detection, with performance improving by an average of 10 percentage points.
- **Non-subtle side tasks** (top-right): Action-only monitors show substantially higher recall, with GPT-4.1-mini detecting nearly all instances, with the exception of the *Exit* side task, which lies near the cluster boundary. Overall, CoT monitor access decreases performance by 25 percentage

points across models, possibly due to persuasive rationalizations or noise introduced in the reasoning.




Figure 3: Catch rates of action-only and CoT monitoring across side tasks, results averaged across untrusted models and attack policies for each side task. The dashed line separates the regions where CoT monitoring outperforms action-only (above), and vice versa (below).

These results suggest that CoT and action-only monitoring provide complementary signals, showing relative strengths in different regimes of task subtlety.

### 3.2 Hybrid Monitoring Recovers Performance and Improves Consistency

Figure 4 introduces a third strategy—*hybrid monitoring*—that aggregates independent CoT and action evaluations through a weighted average as described in Subsection 2.2. This approach yields consistent and substantial improvements: across all models and task types, hybrid monitoring outperforms both action-only and CoT monitoring catch rates, achieving 21% and 42% higher recall respectively.

Figure 5 displays the performance of GPT-4.1-mini and Gemini 2.5 Flash Lite separated by side task subtlety. For non-subtle side tasks, hybrid monitoring performs on par with action-only monitors, while providing 32 percentage points of performance over CoT monitors on average. On the other hand, for subtle side tasks, the hybrid monitors match the performance of CoT monitors, while achieving over twice the catch rate of action-only monitors. Nonetheless, hybrid monitoring catch rates for subtle side tasks remain below 20% on average, suggesting that while more effective, is still limited in absolute terms.

Figure 6 breaks down performance by attack policy for the untrusted Claude 3.7 Sonnet Thinking model, with results for all models reported in Figure 10 in Appendix B. While the hybrid GPT-4.1-mini monitor is not always the top-performing strategy in every scenario, it consistently outperforms both action-only and CoT monitoring on average.

The *dependency* attack, a framing strategy where the untrusted model is instructed to perform the side task as an intermediary step of the main task, is the most effective of all attack policies for




Figure 4: Catch rate performance comparison at 1%FPR of all trusted models for action-only, CoT and hybrid monitors. Catch rate performance comparison at 1% FPR of all trusted models for action-only, CoT and hybrid monitors across side tasks. The subplots display results by monitor model as follows: top left shows GPT-4.1-mini, top right Gemini 2.5 Flash Lite, bottom left Claude 3.5 Haiku, and bottom right GPT-4.1-nano.

non-subtle side tasks. Its effectiveness stems from models frequently mentioning the side task as a benign requirement in their reasoning or, in some cases, claiming they will ignore the side task while still performing it, thereby deceiving the monitors with access to the reasoning traces. For the subtle side tasks, *monitor notice*, where we notify the untrusted model of the monitoring setup, matches the performance of the the *baseline* attack policy.

In contrast, *stealth* attacks are largely ineffective, resulting in significantly higher catch rates than *baseline* across all monitor types, leading to the highest CoT and hybrid monitoring catch rates for subtle side tasks. As this strategy encourages the model to be stealthy in its actions, it often prompts self-reflective reasoning that inadvertently reveals sabotage attempts, making them easier to detect when the monitor reviews the CoT.

## 4 Discussion

Our results reveal that Chain-of-Thought monitoring offers a powerful but nuanced tool for detecting model misbehavior. While access to reasoning traces substantially improves detection of subtle deception attempt, it proves less effective for blatant exploits. Furthermore, our red-team evaluation demonstrates that adversarial models can exploit trusted overseers through deliberate reasoning manipulation, highlighting important limitations in current monitoring protocols.

### 4.1 Side Task Subtlety Affects CoT Usefulness for Monitoring

Our investigation reveals a nuanced picture of the role of CoT monitoring in detecting deceptive behavior in advanced AI systems. Results show monitoring efficacy is dependent on the nature of the side task being pursued by the untrusted model (Figure 3), revealing that action-only and CoT monitoring can serve complementary roles within a comprehensive safety strategy.




Figure 5: Catch rate performance of GPT-4.1-mini (left) and Gemini 2.5 Flash Lite (right) as monitors grouped by untrusted model and side task subtlety.

For subtle side tasks, CoT monitoring significantly improves detection compared to code inspection alone. These tasks involve code manipulations that maintain primary functionality while introducing concealed vulnerabilities, e.g. altered algorithmic complexity or edge case failures. Since these sophisticated deceptions produce minimal changes in the final code output, they are hard to detect through code inspection alone. Having access to the model’s reasoning process substantially increases catch rates for these critical scenarios, though detection rates are far lower than in non-subtle scenarios, highlighting persistent limitations in identifying subtle manipulations.

For non-subtle side tasks, the action-only monitor achieves higher detection rates than CoT monitoring. These tasks involve operations that leave clear traces in the code, such as file operations, network connections, or permission changes. Since these operations are clearly visible in the generated code, analyzing the reasoning traces offers no added value and misleading rationale may actually confuse the monitor, reducing its effectiveness.

Hybrid monitoring, which evaluates the CoT and final action independently before aggregating scores, mitigates these trade-offs. As shown in Figure 4, this approach consistently performs as better than either single-monitor strategy, suggesting hybrid monitoring can be a practical default, especially in real-world deployments where both stealthy and blatant attacks must be guarded against. Notably, while hybrid monitoring achieves  $> 85\%$  catch rates on average for non-subtle side tasks and substantially improves detection in subtle cases, performance on the latter remains limited in absolute terms, with catch rates below  $30\%$  across models.

## 4.2 Monitor Vulnerability to Persuasion

A key vulnerability we identify is that monitors can be deceived by plausible explanations in model reasoning, even when they have access to both the original task description and the final code. As illustrated in Figure 2, untrusted models can generate reasoning that explicitly rejects a harmful action while still implementing it in the final code. Models can also frame their malicious implementations deceptively, such as claiming the side task is being done in accordance with a Red Team-prompted ‘corporate policy’. Our work extends the findings of previous research (Arcuschin et al., 2025; Chen et al., 2025; Lanham et al., 2023) to the safety context, demonstrating that lack of reasoning faithfulness can be exploited to circumvent naive monitoring.

While these vulnerabilities present significant challenges, they may be partially addressed through targeted improvements in monitor design, training and usage. Our experiments, which tested four different models as monitors, reveal that increases in monitor capability led to improved detection performance, a result suggesting that monitor sophistication is key. This observation aligns with related research (Terekhov et al., 2025), which identifies monitor base capability as a more relevant factor for adversarial detection than attacker capability or the capability gap between models. These findings collectively support the hypothesis that ensuring the monitor surpasses a certain capability

threshold is critical for safety. However, it remains an open question whether these trends will persist as models capabilities continue to increase.




Figure 6: Catch rate performance across side task subtlety and attack policy for Claude 3.7 Sonnet Thinking overseen by GPT-4.1-mini.

### 4.3 Limitations and Future Work

Our evaluation focused on single-step coding problems in a controlled environment, which fails to capture the complexity of real-world deployment. Future work should investigate monitoring effectiveness in more realistic coding environments where multi-step, distributed sabotage is possible. This includes scenarios where malicious intent might be distributed across multiple functions, files or commits, with harmful effects emerging through the interaction of seemingly benign components.

We explored several attack policies but did not exhaustively cover all potential adversarial strategies. Early experiments showed that preventing models from mentioning side tasks in their CoT was very inconsistent, so we focused on designing attack policies that frame the side tasks in benign ways, reducing the likelihood of the model referencing them suspiciously. Future research could consider developing scaffolding to simulate advanced model capabilities and investigate more comprehensive fine-tuning protocols that better approximate deceptive behaviors in more advanced AI systems.

Recent work (Needham et al., 2025) reveals that models like Claude Sonnet 3.7 often recognize they are being evaluated, raising concerns about the trustworthiness of CoT traces in evaluation settings. If models know they are being tested, they may strategically craft reasoning to appear aligned, masking deceptive behavior and undermining monitoring. Future research should explore ways to detect or mitigate evaluation awareness, such as using randomized test framing or less-salient evaluation contexts, to ensure CoT monitoring reflects genuine reasoning rather than performative alignment.

Finally, our study does not address how monitoring efficacy might be affected by production constraints such as latency requirements, cost limitations, or scalability challenges. Future work should evaluate these practical considerations, and in particular whether the computational overhead of dual monitoring approaches is justified by the safety gains in specific deployment contexts.

## 5 Related Work

Chain-of-Thought (CoT) is a technique that enables large language models (LLMs) to articulate intermediate steps before their final output, increasing model capabilities on tasks requiring sequential reasoning (Miner et al., 2025). Although the use of CoT initially emerged through prompting techniques such as few-shot examples (Nye et al., 2021; Reynolds and McDonell, 2021; Wei et al., 2022) and zero-shot instructions (Kojima et al., 2022), recent work explicitly trains models via reinforcement learning or supervised fine-tuning to consistently produce CoT reasoning (DeepSeek-AI, 2025; Jaech et al., 2024).

Inspecting models' CoTs could provide insight into their reasoning and intentions. In particular, CoT monitoring could be a promising approach to detecting deceptive reasoning or 'scheming' (Meinke et al., 2025) by future AI models. Greenblatt et al. (2024a) demonstrate that frontier LLMs can fake alignment by using a secret scratchpad, a prompted version of CoT, to fake compliance. When hidden scratchpads were removed, compliance gaps disappeared in scenarios with in-context information but not in finetuned models, suggesting internalized knowledge can enable covert deceptive reasoning. While their work underscores the risks of misalignment, it did not evaluate CoT monitoring within a formal control protocol with targeted attack policies nor examine reasoning models trained via reinforcement learning to reason naturally, the setting we investigate.

However, CoT may not fully represent model reasoning: studies show that models answer consistently despite variations or omissions of the CoT (Lanham et al., 2023), produce coherent rationalizations for implicit biases and use illogical reasoning to simplify solutions (Arcuschin et al., 2025), and fail to acknowledge known influences in their CoT (Chen et al., 2025). While imperfect faithfulness places practical limits on the utility of CoT for interpretability, it does not completely remove the value of monitoring, as it represents an additional stream of information. Our work investigates whether CoT monitoring provides measurable advantages over action monitoring for detecting sabotage, without making claims about the faithfulness of the CoT.

Frontier AI labs (Jaech et al., 2024; Anthropic, 2025) have primarily used CoT monitoring to identify hallucinations and unintentional failures in their models, without evaluation against adversarial attacks or deliberate attempts to deceive the monitor. On the other hand, LlamaFirewall (Chennabasappa et al., 2025) incorporates AlignmentCheck, a component that audits an agent's reasoning at inference time to detect signs of goal hijacking or prompt-induced misalignment, aiming to defend against adversarial behavior. However, while they use strong models to audit an agent's full execution trace for goal consistency, we focus on a weak monitoring setup to investigate subversion detection in the presence of capability gaps that may arise with future, more powerful untrusted models.

CoT supervision has also been used for training-time oversight: Lightman et al. (2023) explore reward model training using fine-grained supervision over the reasoning process, rather than only the final outcome. Meanwhile, Baker et al. (2025) show that weak CoT monitors with explicit guidance to detect specific failure modes can oversee stronger models and detect reward hacking more effectively than action-only monitors. Although they find that optimizing the reasoning trace to reduce reward hacking might obfuscate misaligned behavior, faithfulness has been identified as an explicit objective in model development (Lyu et al., 2023; Paul et al., 2024; Xu et al., 2024), potentially leading to CoT being a more reliable representation of internal model processes.

## 6 Conclusion

Our results demonstrate that CoT monitoring offers measurable safety benefits over baseline action-only monitoring in scenarios involving subtle deceptive behaviors, but also exhibits specific vulnerabilities that must be addressed in comprehensive AI control protocols. We introduce a more nuanced approach to monitoring system design that leverages the complementary strengths of multiple monitoring approaches to close the gaps that any single approach leaves open.

We believe our findings present a case for maintaining transparent, natural language CoT in frontier models as a valuable safety feature that can detect potential sabotage from misaligned models, particularly when monitoring systems are designed to account for the evasion strategies we identified. By characterizing the strengths and limitations of CoT monitoring against sophisticated circumvention attempts, our work seeks to inform broader threat models and bolster safety cases (Korbak et al., 2025).

## Acknowledgments

This project was fully funded by LASR Labs as part of their research program. We would like to acknowledge Joseph Miller, Erin Robertson, Brandon Riggs and Charlie Griffin for their support during the project. Additionally, we thank Joseph Bloom, Adam Kaufman, Trevor Lohrbeer, Robert McCarthy, Alex McKenzie, Fabien Roger, and Tyler Tracy for their feedback on early versions of this paper.

## References

- W. U. Ahmad, S. Narenthiran, S. Majumdar, A. Ficek, S. Jain, J. Huang, V. Noroozi, and B. Ginsburg. Opencodereasoning: Advancing data distillation for competitive coding. *arXiv preprint arXiv:2504.01943*, 2025.
- Anthropic. Introducing computer use, a new claude 3.5 sonnet, and claude 3.5 haiku. <https://www.anthropic.com/news/3-5-models-and-computer-use>, 2024.
- Anthropic. Claude 3.7 sonnet system card. <https://www.anthropic.com/claude-3-7-sonnet-system-card>, 2025. URL <https://www.anthropic.com/claude-3-7-sonnet-system-card>.
- I. Arcuschin, J. Janiak, R. Krzyzanowski, S. Rajamanoharan, N. Nanda, and A. Conmy. Chain-of-thought reasoning in the wild is not always faithful. *arXiv preprint arXiv:2503.08679*, 2025.
- B. Baker, J. Huizinga, L. Gao, Z. Dou, M. Y. Guan, A. Madry, W. Zaremba, J. Pachocki, and D. Farhi. Monitoring Reasoning Models for Misbehavior and the Risks of Promoting Obfuscation, Mar. 2025. URL <http://arxiv.org/abs/2503.11926>. arXiv:2503.11926 [cs].
- A. Bhatt, C. Rushing, A. Kaufman, V. Georgiev, T. Tracy, A. Khan, and B. Shlegeris. Shell Games: Control Protocols for Adversarial AI Agents, Oct. 2024. URL <https://openreview.net/forum?id=oycEeFFX74>.
- Y. Chen, J. Benton, A. Radhakrishnan, J. U. C. Denison, J. Schulman, A. Somani, P. Hase, M. W. F. R. V. Mikulik, S. Bowman, J. L. J. Kaplan, et al. Reasoning models don't always say what they think. *Anthropic Research*, 2025.
- S. Chennabasappa, C. Nikolaidis, D. Song, D. Molnar, S. Ding, S. Wan, S. Whitman, L. Deason, N. Doucette, A. Montilla, et al. Llamafirewall: An open source guardrail system for building secure ai agents. *arXiv preprint arXiv:2505.03574*, 2025.
- G. DeepMind. Gemini 2.5 flash-lite model card. <https://storage.googleapis.com/deepmind-media/Model-Cards/Gemini-2-5-Flash-Lite-Model-Card.pdf>, September 2025.
- DeepSeek-AI. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning, Jan. 2025. URL <http://arxiv.org/abs/2501.12948>. arXiv:2501.12948 [cs].
- R. Greenblatt, C. Denison, B. Wright, F. Roger, M. MacDiarmid, S. Marks, J. Treutlein, T. Belonax, J. Chen, D. Duvenaud, A. Khan, J. Michael, S. Mindermann, E. Perez, L. Petrini, J. Uesato, J. Kaplan, B. Shlegeris, S. R. Bowman, and E. Hubinger. Alignment faking in large language models, Dec. 2024a. URL <http://arxiv.org/abs/2412.14093>. arXiv:2412.14093.
- R. Greenblatt, B. Shlegeris, K. Sachan, and F. Roger. AI Control: Improving Safety Despite Intentional Subversion, July 2024b. URL <http://arxiv.org/abs/2312.06942>. arXiv:2312.06942 [cs].
- C. Griffin, L. Thomson, B. Shlegeris, and A. Abate. Games for AI Control: Models of Safety Evaluations of AI Deployment Protocols, Sept. 2024. URL <http://arxiv.org/abs/2409.07985>. arXiv:2409.07985 [cs].
- A. Jaech, A. Kalai, A. Lerer, A. Richardson, A. El-Kishky, A. Low, A. Helyar, A. Madry, A. Beutel, A. Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.
- T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022.
- T. Korbak, J. Clymer, B. Hilton, B. Shlegeris, and G. Irving. A sketch of an AI control safety case, Jan. 2025. URL <http://arxiv.org/abs/2501.17315>. arXiv:2501.17315 [cs].
- J. Kutasov, Y. Sun, P. Colognese, T. van der Weij, L. Petrini, C. B. C. Zhang, J. Hughes, X. Deng, H. Sleight, T. Tracy, et al. Shade-arena: Evaluating sabotage and monitoring in llm agents. *arXiv preprint arXiv:2506.15740*, 2025.
- T. Lanham. Externalized reasoning oversight: a research direction for language model alignment. <https://www.alignmentforum.org/posts/FRRb6Gqem8k69ocbi>, 2022.
- T. Lanham, A. Chen, A. Radhakrishnan, B. Steiner, C. Denison, D. Hernandez, D. Li, E. Durmus, E. Hubinger, J. Kernion, K. Lukošiušė, K. Nguyen, N. Cheng, N. Joseph, N. Schiefer, O. Rausch,

- R. Larson, S. McCandlish, S. Kundu, S. Kadavath, S. Yang, T. Henighan, T. Maxwell, T. Telleen-Lawton, T. Hume, Z. Hatfield-Dodds, J. Kaplan, J. Brauner, S. R. Bowman, and E. Perez. Measuring Faithfulness in Chain-of-Thought Reasoning, July 2023. URL <http://arxiv.org/abs/2307.13702>. arXiv:2307.13702.
- H. Lightman, V. Kosaraju, Y. Burda, H. Edwards, B. Baker, T. Lee, J. Leike, J. Schulman, I. Sutskever, and K. Cobbe. Let's verify step by step, 2023. URL <https://arxiv.org/abs/2305.20050>.
- Q. Lyu, S. Havaldar, A. Stein, L. Zhang, D. Rao, E. Wong, M. Apidianaki, and C. Callison-Burch. Faithful chain-of-thought reasoning. In *The 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (IJCNLP-AACL 2023)*, 2023.
- A. Mallen, C. Griffin, A. Abate, and B. Shlegeris. Subversion Strategy Eval: Evaluating AI's stateless strategic capabilities against control protocols, Dec. 2024. URL <http://arxiv.org/abs/2412.12480>. arXiv:2412.12480 [cs].
- A. Meinke, B. Schoen, J. Scheurer, M. Balesni, R. Shah, and M. Hobbhahn. Frontier Models are Capable of In-context Scheming, Jan. 2025. URL <http://arxiv.org/abs/2412.04984>. arXiv:2412.04984 [cs].
- METR. Measuring ai ability to complete long tasks. <https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks/>, 03 2025.
- S. Miner, Y. Takashima, S. Han, S. Kouteili, F. Erata, R. Piskac, and S. J. Shapiro. Scheherazade: Evaluating Chain-of-Thought Math Reasoning in LLMs with Chain-of-Problems, Feb. 2025. URL <http://arxiv.org/abs/2410.00151>. arXiv:2410.00151 [cs].
- J. Needham, G. Edkins, G. Pimpale, H. Bartsch, and M. Hobbhahn. Large language models often know when they are being evaluated. *arXiv preprint arXiv:2505.23836*, 2025.
- M. Nye, A. J. Andreassen, G. Gur-Ari, H. Michalewski, J. Austin, D. Bieber, D. Dohan, A. Lewkowycz, M. Bosma, D. Luan, C. Sutton, and A. Odena. Show your work: Scratchpads for intermediate computation with language models, 2021. URL <https://arxiv.org/abs/2112.00114>.
- OpenAI. Gpt-4.1. <https://platform.openai.com>, 2024.
- D. Paul, R. West, A. Bosselut, and B. Faltings. Making reasoning matter: Measuring and improving faithfulness of chain-of-thought reasoning. *arXiv preprint arXiv:2402.13950*, 2024.
- L. Reynolds and K. McDonell. Prompt programming for large language models: Beyond the few-shot paradigm. In *Extended abstracts of the 2021 CHI conference on human factors in computing systems*, pages 1–7, 2021.
- R. Shah, A. Irpan, A. M. Turner, A. Wang, A. Conmy, D. Lindner, J. Brown-Cohen, L. Ho, N. Nanda, R. A. Popa, R. Jain, R. Greig, S. Albanie, S. Emmons, S. Farquhar, S. Krier, S. Rajamanoharan, S. Bridgers, T. Ijtaye, T. Everitt, V. Krakovna, V. Varma, V. Mikulik, Z. Kenton, D. Orr, S. Legg, N. Goodman, A. Dafoe, F. Flynn, and A. Dragan. An approach to technical agi safety and security, 2025. URL <https://arxiv.org/abs/2504.01849>.
- M. Terekhov, Z. N. D. Liu, C. Gulcehre, and S. Albanie. Control tax: The price of keeping ai in check. *arXiv preprint arXiv:2506.05296*, 2025.
- UK AI Safety Institute. Inspect ai: Framework for large language model evaluations, 2024. URL <https://inspect.aisi.org.uk/>.
- UK AI Safety Institute. Research Agenda. <https://www.aisi.gov.uk/research-agenda>, 2025.
- J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- xAI. Grok 3 [large language model]. <https://x.ai/news/grok-3>, 2025.
- J. Xu, H. Fei, L. Pan, Q. Liu, M.-L. Lee, and W. Hsu. Faithful logical reasoning via symbolic chain-of-thought. *arXiv preprint arXiv:2405.18357*, 2024.
- T. Y. Zhuo, M. C. Vu, J. Chim, H. Hu, W. Yu, R. Widyasari, I. N. B. Yusuf, H. Zhan, J. He, I. Paul, et al. Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions. *arXiv preprint arXiv:2406.15877*, 2024.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The main claims made in the abstract and introduction are backed by results obtained in Sections 3 and 4.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Limitations of the work are discussed in Subsection 4.3.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper describes experimental details in Section 2, including model configurations, dataset partitioning, and evaluation metrics. Filtering decisions, threshold calibration and task pairing, are further detailed in Appendices A.2, A.3, and C, supporting full reproduction of the main results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide full access to our code at <https://github.com/cotredhanded/code> and sufficient instructions to replicate our data.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We describe our dataset and evaluation details in Section 2 with additional details included in Appendices A, C, B and D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: All figures in the main paper include error bars.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The paper reports estimated compute usage in Appendix D, including the number of inference calls, APIs used (Anthropic, Google, Groq, OpenAI, Together), model configurations, and compute requirements for a targeted fine-tuning experiment with DeepSeek R1 Distill Qwen 14B to simulate advanced attack policies.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research conducted in the paper conforms with the NeurIPS Code of Ethics, including Potential Harms Caused by the Research Process, Societal Impact and Potential Harmful Consequences and Impact Mitigation Measures.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Section 4 discusses both the benefits of CoT monitoring as a safety mechanism and the risks from unfaithful or misleading reasoning traces. We emphasize the utility of CoT transparency while cautioning against overreliance on naive CoT-based protocols.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

#### 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [No]

Justification: The paper does not introduce new models or datasets with dual-use risk; it evaluates previously released models and uses an existing dataset. No safeguards are needed for the release.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

#### 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Existing assets are cited and their licenses are mentioned and respected.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [No]

Justification: No new datasets or models are introduced. The paper evaluates existing open-source or commercially available models and benchmarks.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

#### 16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [\[Yes\]](#)

Justification: The paper evaluates the use of LLMs as both reasoning agents and monitors. These models are core to the methodology and results.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.