
Information-Theoretic Conditions for Chain-of-Thought Monitorability and Methods for Improving It

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Chain-of-thought (CoT) monitors are LLM-based systems that analyze reasoning
2 traces to detect when outputs may exhibit attributes of interest such as harmfulness,
3 bias, or other properties. In this paper, we develop a rigorous information-theoretic
4 framework to analyze the fundamental conditions that determine CoT monitorability.
5 Our analysis establishes two necessary conditions for successful monitoring:
6 (i) the CoT must encode sufficient information about the target attribute, and (ii)
7 the monitor must be capable of extracting this information. We show that the
8 success of CoT monitoring hinges on conditional mutual information between
9 outputs and CoTs. We further demonstrate that CoT monitorability can be system-
10 atically improved through targeted training objectives. To this end, we propose two
11 complementary approaches: (a) an oracle-based method that directly rewards the
12 monitored model for producing CoTs that maximize monitor accuracy, and (b) a
13 more practical, label-free approach that maximizes conditional mutual information
14 between outputs and CoTs. Both methods significantly improve monitor accuracy
15 while preventing CoT degeneration even when training against a monitor, thereby
16 mitigating reward hacking even when the task reward is imperfectly specified.

17

1 Introduction

18 Modern large language models (LLMs) are increasingly trained to generate reasoning traces, also
19 called chain-of-thought (CoT), before producing final outputs. This capability has enabled the
20 development of CoT monitors – LLM-based systems that analyze these reasoning traces to detect
21 attributes such as harmfulness or bias in the resulting outputs. By analyzing the CoT, monitors can
22 theoretically identify problematic attributes *before* they manifest in the output, e.g., flawed reasoning
23 patterns or attempts at deception. However, our understanding of when – and to what extent – CoT
24 monitors can provide robust supervision remains limited. Challenges such as CoT obfuscation [Baker
25 et al., 2025], steganographic reasoning [Stickland and Korbak, 2025], and general lack of faithfulness
26 in CoT reasoning [Turpin et al., 2023, Lanham et al., 2023] may significantly constrain the practical
27 effectiveness of CoT monitoring.
28 This work aims to further our understanding of the robustness and utility of CoT monitoring. Our
29 contributions include: (1) a rigorous probabilistic framework through which we establish necessary
30 information-theoretic conditions for monitorability, and analyze the gap between ideal and practical
31 CoT monitors; and (2) a practical training objective that demonstrably improves CoT monitor
32 accuracy while preventing CoT degeneration.

33 **2 Background**

34 **Entropy.** The Shannon entropy of a discrete random variable A , denoted $H(A)$, measures its
 35 average uncertainty and is defined as $H(A) = -\sum_a p(a) \log p(a)$. The conditional entropy, $H(A | B)$,
 36 measures the remaining uncertainty in A after observing B .

37 **Mutual Information.** The mutual information (MI) between two random variables, $I(A; B)$,
 38 quantifies the reduction in uncertainty about one variable from knowing the other. It is defined as:

$$I(A; B) = H(A) - H(A | B). \quad (2.1)$$

39 MI is symmetric, so $I(A; B) = I(B; A)$. A value of $I(A; B) = 0$ implies that the variables are
 40 independent. The conditional mutual information, $I(A; B | C)$, measures the MI between A and B
 41 given a third variable C . A key identity we will use is the chain rule for mutual information:

$$I(A, B; C) = I(A; C) + I(B; C | A). \quad (2.2)$$

42 **2.1 Related Works**

43 **CoT Monitoring and Evasion.** Chain-of-thought (CoT) monitoring has been proposed as a mecha-
 44 nism for interpretable oversight of LLMs [Korbak et al., 2025]. However, recent work highlights its
 45 brittleness: CoT is not reliably faithful [Turpin et al., 2023, Lanham et al., 2023, Chen et al., 2025,
 46 Barez et al., 2025] – and may thus not reflect whether models are engaging in unwanted behavior –
 47 and models can potentially evade monitors through obfuscation [Baker et al., 2025] or steganography
 48 [Motwani et al., 2024, Mathew et al., 2024]. Emmons et al. [2025] show that monitoring is more
 49 effective on tasks that require CoT, while Stickland and Korbak [2025] explore generator-to-monitor
 50 transferability. These findings motivate a more formal understanding of when CoT monitoring
 51 succeeds.

52 **Information-Theoretic Approaches.** Our analysis builds on a growing body of work using
 53 information-theoretic tools for interpretability and control. Prior work has applied information-
 54 theoretic tools to interpretability and alignment [Xu et al., 2020, Belrose et al., 2025], and has
 55 analyzed CoT through information gain and bottleneck principles [Ton et al., 2025, Lei et al., 2025].
 56 However, these approaches do not address monitorability explicitly.

57 **Training Strategies and Reward Hacking** Training models to remain monitorable is challenging:
 58 optimizing against monitors may induce obfuscation [Baker et al., 2025], echoing broader concerns
 59 about alignment faking and deceptive behaviors [Park et al., 2023, Scheurer et al., 2024, Greenblatt
 60 et al., 2024, Meinke et al., 2025] and multi-agent collusion [Motwani et al., 2024, Dou et al., 2025].
 61 Our work introduces a label-free conditional mutual information objective that aims to preserve
 62 monitorability and mitigate reward hacking, complementing broader efforts to improve reasoning
 63 faithfulness [Paul et al., 2024] and robustness of RL training [Farquhar et al., 2025, Liu et al., 2025].

64 **3 Information-Theoretic Analysis of CoT monitorability**

65 In this section, we develop a formal framework to analyze when and
 66 why chain-of-thought (CoT) reasoning could be monitored effec-
 67 tively. We begin by formalizing the probabilistic setup and defining
 68 key concepts. We then identify necessary information-theoretic
 69 conditions for successful monitoring, showing that the conditional
 70 mutual information between CoTs and outputs, $I(Z; O | X)$, plays a
 71 critical role in determining monitorability. Finally, we analyze the
 72 gap between ideal (Bayes-optimal) monitors and practical, capacity-
 73 bounded monitors, quantifying how decoder limitations affect mon-
 74 itoring performance.

75 **3.1 Setup and Definitions**

76 We model a large language model as a policy, π_θ , that generates text autoregressively. Our framework
 77 focuses on tasks where the model produces a Chain-of-Thought (CoT) before the final output. More

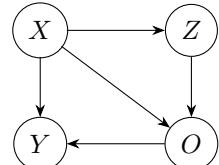


Figure 1: Graphical model of our Chain-of-Thought setup.

78 specifically, we assume the directed acyclic graph (DAG)

$$X \rightarrow Z, \quad (X, Z) \rightarrow O, \quad (X, O) \rightarrow Y,$$

79 where X is the prompt, Z is the chain-of-thought (CoT), O is the model's final action/answer and Y
80 is the attribute of interest.

81 **Definition 3.1** (Action Monitor). *An action monitor is a stochastic mapping*

$$g : \mathcal{O} \times \mathcal{X} \longrightarrow \Delta(\mathcal{Y}), \quad y \sim g(\cdot | O, X),$$

82 where $\Delta(\mathcal{Y})$ denotes the probability simplex on \mathcal{Y} .

83 **Definition 3.2** (CoT Monitor). *Similarly, CoT monitor is a stochastic mapping*

$$m_\psi : \mathcal{Z} \times \mathcal{X} \longrightarrow \Delta(\mathcal{Y}), \quad y \sim m_\psi(\cdot | Z, X).$$

84 For both monitors, point predictions could be attained by performing argmax over them, e.g.,

$$y = \underset{Y \in \mathcal{Y}}{\operatorname{argmax}} m_\psi(Y | Z, X).$$

85 **Definition 3.3** (Bayes–Optimal CoT Monitor). *Let $\pi(O | Z, X)$ be the model's conditional distribution over final actions and let $g(Y | O, X)$ be the stochastic action-monitor defined above. The ground truth distribution of the attribute Y given the prompt x and chain-of-thought z can be defined as*

$$m_g^\pi(Y = y | Z, X) = \mathbb{E}_{o \sim \pi(O | Z, X)}[g(Y = y | O = o, X)].$$

89 *Under the 0–1 scoring rule, the Bayes–optimal CoT monitor is the classifier that outputs the MAP 90 label*

$$\hat{y} = m^*(z, x) := \underset{y \in \mathcal{Y}}{\operatorname{argmax}} m_g^\pi(Y = y | Z = z, X = x).$$

91 **Definition 3.4** (α -CoT monitorable task). *For $\alpha \in [0, 1]$, we say an attribute Y is α -CoT-monitorable 92 with respect to policy π if the Bayes-optimal CoT monitor $m^*(z, x)$ has accuracy at least α .*

93 Note that the performance of any practical CoT monitor is bounded above by the Bayes-optimal 94 CoT monitor, whose performance in turn is bounded by the action monitor g . We emphasize that the 95 CoT-monitorability of an attribute is intrinsically linked to the policy π that generates the CoT being 96 monitored. There are two natural corollaries of this fact: (a) an attribute Y being monitorable w.r.t. 97 some policy π_1 is no guarantee that the same attribute would be monitorable for some other policy π_2 . 98 (b) monitorability of an attribute Y may be improved through changing the distribution of Z induced 99 by the policy, for example, by training π to produce more informative chain-of-thought reasoning.

100 Due to space limitations, we sometimes omit proof from the main text, and defer them to the 101 Appendix A.

102 3.2 Necessary Conditions for CoT Monitorability

103 In this section, we investigate the necessary information-theoretic conditions that determine when 104 a CoT may be monitorable for some attribute of interest Y . Specifically, we analyze under what 105 properties of the CoT and the policy π , a task could be considered α -CoT monitorable.

106 Our analysis yields two main insights. Firstly, we show that conditional mutual information between 107 CoT and the attribute of interest $I(Y; Z | X)$ gives the ceiling on the performance of *any* CoT 108 monitor. Next, we show that positivity of mutual information between CoT and actions $I(O; Z | X)$ 109 is a necessary condition for monitorability of any attribute. However, it is not a sufficient condition, 110 and despite positive $I(O; Z | X)$, there may exist attributes which may be difficult to monitor.

111 **Proposition 3.5** (Necessary MI threshold for α -monitorability). *Let Y take values in a finite label 112 set \mathcal{Y} with $\kappa = |\mathcal{Y}| \geq 2$. For $\alpha \in [1/\kappa, 1]$, define*

$$T_\alpha := h_b(1 - \alpha) + (1 - \alpha) \log(\kappa - 1), \quad h_b(u) = -u \log u - (1 - u) \log(1 - u).$$

113 *If there exists a (possibly stochastic) monitor $m_\psi(Y | Z, X)$ whose induced 0–1 accuracy is at least 114 α under the data-generating distribution, then*

$$I(Y; Z | X) \geq H(Y | X) - T_\alpha. \tag{3.1}$$

115 The above proposition formalizes the intuition that a CoT-monitor can only successfully predict some
 116 attribute of interest Y of the output from CoT to the extent that CoT encodes information about
 117 this attribute. Two important caveats apply: (1) it establishes existence of some monitor achieving
 118 accuracy α when $I(Y; Z|X)$ is sufficiently high, but does not guarantee that any specific monitor
 119 class (e.g., prompted LLMs) will attain this accuracy; (2) even with high mutual information, practical
 120 monitors may fail if the information is encoded in ways that resist extraction (e.g., steganographically).
 121 We address these limitations in our analysis of practical monitors in Section 3.3.

122 **Do CoTs that matter for the task also help monitoring?** Emmons et al. [2025] empirically
 123 demonstrate that tasks requiring explicit CoT reasoning are more amenable to CoT monitoring.
 124 Indeed, our theoretical analysis confirms this observation, by showing that the necessity of CoT
 125 reasoning for task completion is a necessary precondition for the monitorability of any attribute Y of
 126 the final output. However, it is not a sufficient condition.

127 **Lemma 3.6** (MI decomposition). *Let $Y = g(O, X)$ for the stochastic map $g : O \times \mathcal{X} \rightarrow \mathcal{Y}$. Then*

$$I(O; Z | X) = I(Y; Z | X) + I(O; Z | Y, X). \quad (3.2)$$

128 *Proof.* Apply the chain rule two ways: $I(O, Y; Z | X) = I(Y; Z | X) + I(O; Z | Y, X) =$
 129 $I(O; Z | X) + I(Y; Z | O, X)$. Note that by assumption we have $Y \perp\!\!\!\perp Z | O, X$, which means
 130 $I(Y; Z | O, X) = 0$, which in turn gives (3.2). \square

131 **Proposition 3.7** (Necessity of CoT reasoning is a necessary but not sufficient condition for monitora-
 132 bility). *Assume $Y = g(O, X)$ where $g : O \times \mathcal{X} \rightarrow \mathcal{Y}$ is a deterministic map.*

- 133 1. *(Necessity) For any attribute Y , if a monitor $m_\psi(Z, X) \in M$ with non-trivial accuracy
 134 exists then $I(O; Z | X) > 0$.*
- 135 2. *(Not sufficiency) $I(O; Z | X) > 0$ does not imply there exists a monitor $m_\psi(Z, X) \in M$
 136 with non-trivial accuracy.*

137 *Proof. Necessity:* By Proposition 3.5, existence of a monitor with non-trivial accuracy implies
 138 $I(Y; Z | X) > 0$. From Lemma 3.6 we have $I(O; Z | X) \geq I(Y; Z | X)$, so $I(Y; Z | X) > 0 \Rightarrow$
 139 $I(O; Z | X) > 0$.

140 *Not Sufficiency:* By Lemma 3.6, we have $I(O; Z | X) = I(Y; Z | X) + I(O; Z | Y, X)$.
 141 Because both terms on RHS are nonnegative, it is possible that $I(O; Z | Y, X) > 0$ while $I(Y; Z | X) = 0$,
 142 yielding $I(O; Z | X) > 0$. Because $I(Y; Z | X) = 0$, then by Proposition 3.5 we have that
 143 no CoT monitor with non-trivial accuracy exists. Because $I(Y; Z | X) = 0$, by the contrapositive of
 144 Proposition 3.5 we have that no CoT monitor with non-trivial accuracy exists. \square

145 3.3 Information Gap Between Practical Monitor and The Ideal Monitor

146 The analysis in the previous section considers Bayes optimal monitor that uses the ground truth
 147 decoder distribution m_g^π , and establishes when such a monitor might be realizable. However, in
 148 practice, CoT monitors are (prompted or fine-tuned) LLMs with finite capacity and thus there may be
 149 a decoder mismatch – the decoder $q(Y|Z, X)$ induced by the monitor LLM may be imperfect and
 150 may not match well with the Bayes optimal decoder m^* . We make the loss of performance due to
 151 this decoder mismatch precise through the following proposition:

152 **Proposition 3.8** (Information gap.). *Let π be the (true) distribution over (Y, Z, X) . Let the CoT
 153 monitor be a stochastic decoder $m_\psi(Y | Z, X)$ and define*

$$154 I_{m_\psi}(Y; Z | X) := \mathbb{E}_{\pi(Y, Z, X)} \left[\log \frac{m_\psi(Y|Z, X)}{m_\psi(Y|X)} \right], \quad m_\psi(Y | X) := \mathbb{E}_{Z \sim \pi(Z|X)} [m_\psi(Y | Z, X)].$$

155 Then $I(Y; Z | X) \geq I_{m_\psi}(Y; Z | X)$.

156 **Corollary 3.9** (The information gap is minimized by using the data generating distribution as the
 157 decoder). $I(Y; Z | X) = I_{m_\psi}(Y; Z | X)$ if and only if $m_\psi(Y | Z, X) = \pi(Y | X, Z) =$
 158 $m_g^\pi(Y | X, Z)$.

159 3.4 Key Takeaways

160 Our information-theoretic view highlights two distinct failure modes for CoT monitoring:

161 • **Encoding bottleneck.** If the CoT carries *no* information about the model’s action, $I(O; Z | X) = 0$, then it also carries no information about any attribute of that action, $I(Y; Z | X) = 0$. In this
 162 regime no CoT monitor with non-trivial accuracy is realizable.
 163

164 • **Decoder limitation.** When $I(O; Z | X) > 0$, the information may exist but a practical monitor
 165 $m_\psi(y | x, z)$ might fail due to its inability to extract this information.

166 **4 Can CoT Monitorability be Improved?**

167 Prior works, such as [Baker et al. \[2025\]](#), have focused on how CoT monitorability could be preserved.
 168 We go a step further, and consider what can be done to improve CoT monitorability of a task. Our
 169 analysis demonstrates that targeted training objectives can substantially improve the information
 170 content of CoTs, making them more amenable to monitoring.

171 We explore two primary approaches. First, we introduce an *oracle-based* method where the model is
 172 rewarded for producing CoTs that maximize the accuracy of a fixed, external monitor. While this
 173 provides a conceptual benchmark, its practical application is limited by the need for a reliable oracle
 174 to provide ground-truth labels.

175 To address this limitation, we propose a practical, *label-free* proxy objective: maximizing the
 176 conditional mutual information $I(O; Z | X)$ between the model’s actions and its CoT. This objective
 177 encourages the model to generate reasoning traces that are inherently informative about its final
 178 decision, preventing the CoT from becoming decoupled or obfuscated. We show how this objective
 179 can be efficiently estimated and integrated into a reinforcement learning framework.

180 **4.1 Oracle Monitor-in-the-Loop RL**

181 As a first approach, we consider a setup with a fixed, pre-trained LLM CoT monitor m_ψ . We only
 182 update the policy π_θ of the monitored model. Let g denote the oracle action monitor, then given a
 183 trajectory (y, o, z, x) , we define a binary oracle reward signal:

$$R_{\text{Oracle}}(o, z, x) := \begin{cases} +1, & \text{if } m_\psi(z, x) = g(o, x), \\ -1, & \text{otherwise.} \end{cases} \quad (4.1)$$

184 This signal directly rewards the model for generating a CoT z that leads the monitor to the correct
 185 conclusion $m_\psi(z, x) = g(o, x) = y$. The total reward function is an augmentation of the original
 186 task reward $R_{\text{Task}}(o)$:

$$R_{\text{overall}} = R_{\text{Task}}(o) + \mu R_{\text{Oracle}}(o, z, x), \quad (4.2)$$

187 where $\mu \geq 0$ is a hyperparameter that controls the strength of the oracle feedback.

188 With a fixed monitor m , maximizing $\mathbb{E}[R_{\text{Oracle}}(O, Z, X)]$ encourages π_θ to produce CoTs that are
 189 more easily decodable by that specific monitor, indirectly maximizing $I_{m_\psi}(Y; Z | X)$. However, this
 190 method’s reliance on an external oracle for labels Y makes it costly and impractical in real-world
 191 scenarios, motivating the need for a label-free alternative.

192 **4.2 A Label-Free Proxy: Maximizing $I(O; Z | X)$**

193 To overcome the need for an oracle, we propose maximizing the conditional mutual information
 194 $I(O; Z | X)$ as a proxy objective. This encourages a strong coupling between the CoT and the final
 195 action, which is a prerequisite for monitorability (see [Proposition 3.7](#)).

196 The relationship between our proxy and the true monitorability objective, $I(Y; Z | X)$, is captured
 197 by the chain rule for mutual information (as in [Lemma 3.6](#)):

$$I(O; Z | X) = I(Y; Z | X) + I(O; Z | X, Y). \quad (4.3)$$

198 The term $I(O; Z | X, Y)$ represents the gap between the true and the proxy objective. This represents
 199 information that is captured by the proxy but is irrelevant to the attribute being monitored – such
 200 as stylistic variations in the outputs. Crucially, this term is bounded by the conditional entropy of
 201 the CoT and of the output: $H(Z | X, Y) \geq I(O; Z | X, Y) \leq H(O | X, Y)$. Thus, we have the
 202 following inequalities for our proxy objective:

$$I(O; Z | X) \leq I(Y; Z | X) + H(Z | X, Y),$$

$$I(O; Z | X) \leq I(Y; Z | X) + H(O | X, Y).$$

203 This inequality guarantees that an increase in our proxy objective, $I(O; Z | X)$, will lead to an
 204 increase in the true monitorability objective, $I(Y; Z | X)$, provided that $H(Z | Y, X)$ or $H(O |$
 205 $Y, X)$ is small. For a given attribute Y , $H(Z | X, Y)$ quantifies the intrinsic CoT diversity, while
 206 $H(O | X, Y)$ quantifies the intrinsic output diversity. In both cases, we expect smaller values for
 207 coarse-grained, semantic attributes that significantly constrain the LLM’s generations (e.g., ‘*Is the*
 208 *model hard-coding the values of unit tests?*’, compared to more fine-grained attributes (e.g., ‘*Does*
 209 *the generation contain the letter e?*’).

210 **4.2.1 Estimation and Practical Optimization of $I(O; Z | X)$**

211 To use the conditional mutual information $I(O; Z | X)$ as a practical reward signal in a reinforcement
 212 learning setup, we require a per-rollout point estimator. For a rollout i that produces an output o_i and
 213 a CoT z_i from a prompt x_i , we estimate the mutual information as:

$$\hat{I}_i = \underbrace{\log \pi_\theta(o_i | z_i, x_i)}_{\text{Output log-likelihood given CoT}} - \underbrace{\log \hat{\pi}_\theta(o_i | x_i)}_{\text{Estimated marginal log-likelihood}} \quad (4.4)$$

214 We leverage the structure of Group-based Reward Policy Optimization (GRPO), which provides
 215 multiple rollouts for each prompt. This allows for an efficient Monte Carlo approximation of the
 216 marginal term, $\hat{\pi}_\theta(o_i | x_i)$. For a prompt x_i , we use the corresponding group of CoTs, $\mathcal{G}(x_i) = \{z_k\}$,
 217 to estimate the marginal probability as:

$$\hat{\pi}_\theta(o_i | x_i) = \frac{1}{|\mathcal{G}(x_i)|} \sum_{z_k \in \mathcal{G}(x_i)} \pi_\theta(o_i | z_k, x_i). \quad (4.5)$$

218 This term represents the average likelihood of the output o_i over all the different reasoning paths
 219 generated for the same prompt.

220 A naive maximization of \hat{I}_i could encourage the model to generate unnecessarily long outputs
 221 to artificially inflate the log-probabilities. To counteract this, we normalize the estimate by the
 222 length of the output, $\text{len}(o_i)$, though we omit it from the notation for brevity. Furthermore, since
 223 mutual information is unbounded (for our purposes), directly using this normalized value as a reward
 224 could destabilize training by overpowering other objectives, such as the task reward R_{Task} . To
 225 ensure a balanced optimization, we bound \hat{I}_i by applying some suitable transformation $f(\cdot)$. In our
 226 experiments, we use $f(a) = \tanh(2a)$. This gives us our final mutual information reward component:

$$R_{\text{MI}}(o_i, z_i, x_i) = f(\hat{I}_i) = \tanh(2\hat{I}_i). \quad (4.6)$$

227 The final reward function for training the policy is then a weighted combination of the task reward
 228 and our MI reward:

$$R_{\text{overall}} = R_{\text{Task}}(o_i) + \lambda_{\text{MI}} \cdot R_{\text{MI}}(o_i, z_i, x_i), \quad (4.7)$$

229 where λ_{MI} is a hyperparameter that controls the strength of the monitorability objective.

230 **5 Experiments and Results**

231 In this section, we empirically validate our theoretical framework and demonstrate the effectiveness
 232 of our proposed approaches for improving CoT monitorability. We first describe our experimental
 233 setup, including models, datasets, and evaluation metrics, followed by a comprehensive analysis of
 234 our results.

235 **5.1 Experimental Setup**

236 **Environment:** Our primary experimental environment is inspired by the honeypot environment of
 237 Baker et al. [2025]. Specifically, we create a coding environment where the RL agent could get high
 238 reward either by implementing some sort of shortcut solution (e.g., hardcoding answers to known test
 239 cases) or a general solution. Our environment is based on the Mostly Basic Python Problems (MBPP)
 240 dataset [Austin et al., 2021]. This dataset contains 974 short Python programming tasks designed for
 241 entry-level programmers, each consisting of a natural language description, a reference solution, and
 242 three test cases to verify functional correctness.

243 In our setup, we sample a rollout by prompting an LLM to write code to solve a given Python
244 programming problem described in natural language. Within this prompt, we also include two
245 test-cases (henceforth referred to as ‘revealed’ test cases). We hold out a third test case to evaluate
246 whether the model is writing code specific to the two revealed test cases or a general solution for
247 the given prompt. The test reward is based solely on whether the model’s written code passes the
248 revealed test cases – +1 if it passes, 0 otherwise. As such, this environment provides an option to the
249 model to attain high-reward by writing code that passes only the two revealed test-cases.

250 **Models:** We perform reinforcement learning training on Qwen2.5-3B-Instruct and Qwen2.5-7B-
251 Instruct models; these are our π_θ . To ensure our models explore both general and “hacky” solutions
252 (code that only passes the revealed test cases, e.g., by hardcoding values), we performed a light round
253 of supervised fine-tuning on a small dataset of code examples of test-hacking before performing
254 reinforcement learning training. See Appendix B.1 for details.

255 **Monitor:** We use Qwen2.5-7B-Instruct as our monitor model, m_ψ . The monitor observes the
256 prompt and the CoT of the monitored model, π_θ , then outputs a binary decision regarding whether
257 the model is implementing a true solution or a hacky solution. The monitoring reward R_{Monitor} for
258 π_θ is +1 when the monitor determines the model is not hacking, and -1 when it detects hacking. We
259 provide the prompt that we use for the monitor model Appendix B.3.

260 **Training Details:** We use the Group Relative Policy Optimization (GRPO) algorithm [Shao et al.,
261 2024] for RL training. The same hyperparameters were used for all experiments; we provide them
262 in Appendix B.2.

263 To establish a clear delineation between chain-of-thought reasoning and final outputs, we use a
264 structural formatting reward in all the experiments. This reward incentivizes models to organize
265 their responses into distinct XML-tagged sections: a `<think>...` block containing the
266 reasoning process (corresponding to CoT z) and a `<code>...` block containing the final
267 solution (corresponding to output o). This reward is quite quickly satisfied by LLMs in the training,
268 and has no bearing on the results, hence we omit explicit mention of this reward in subsequent
269 discussion.

270 **Evaluation Metrics:** We evaluate our approach using four key metrics:

- 271 • **Mutual Information (MI):** We measure $I(O; Z|X)$ between the model’s outputs and CoTs to
272 quantify the information content in the reasoning traces.
- 273 • **Hacking Rate:** The percentage of samples for which the code written by the model passes only
274 the revealed test cases but fails the held-out test case.
- 275 • **Monitor Accuracy:** The accuracy of the monitor in correctly identifying whether the model is
276 implementing a hacky solution.
- 277 • **True Pass Rate:** The percentage of samples for which the code written by the model passes all
278 tests – both the revealed tests and the held out test. This can be looked at as a measure of true
279 model performance.

280 5.2 Results and Analysis

281 In Figure 2 we show results for 3B and 7B models from Qwen-2.5 family trained under four different
282 reward configurations. To make it easier to decipher the general trend, we plot using exponential
283 moving average with $\alpha = 0.25$.

284 **Mutual information reward improves CoT-monitorability:** Figure 2 shows the evolution of MI
285 ($I(O; Z|X)$) during training. Notably mutual information deteriorates throughout the RL training
286 for models optimizing R_{Task} and $R_{\text{Task}} + R_{\text{Monitor}}$. This decrease in MI for $R_{\text{Task}} + R_{\text{Monitor}}$ is
287 accompanied by a decrease in monitor accuracy, indicating the obfuscation behavior due to optimizing
288 against CoT monitor that was also observed by Baker et al. [2025].

289 Critically, both our proposed approaches – oracle reward and MI reward – successfully maintain higher
290 MI values throughout training, and maintain high monitor accuracy. This confirms our theoretical

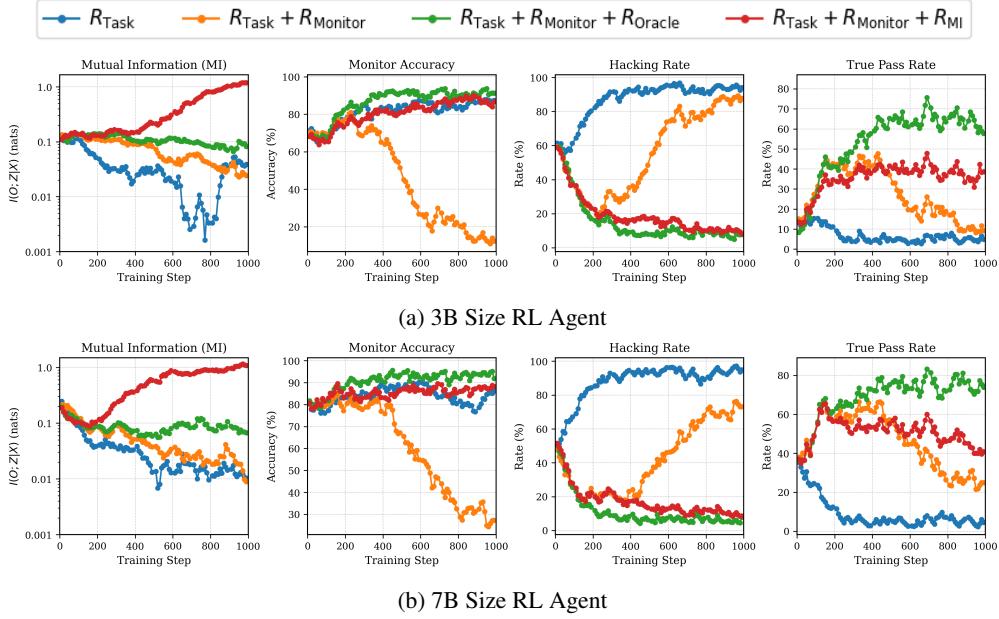


Figure 2: Comparison of (i) Mutual Information $I(O; Z|X)$, (ii) Monitor Accuracy, (iii) Hacking Rate, and (iv) Monitor Accuracy across training steps for the four different reward configurations.

291 prediction that maximizing $I(O; Z|X)$ is an effective strategy for preserving CoT monitorability,
 292 even when optimization pressure is being applied to CoT.

293 **Reward hacking and training against CoT-monitor:** As shown in Figure 2, the model trained
 294 against a vulnerable task reward learns to hack it, and frequently produces code that only passes the
 295 revealed test cases. The model trained to optimize $R_{\text{Task}} + R_{\text{Monitor}}$ does initially learn to hack less
 296 but once it learns that it can evade the monitor via obfuscating its CoT, it reverts back to the hacking
 297 behavior while simultaneously improving on fooling the CoT monitor.

298 However, both models that additionally also optimize either R_{Oracle} or R_{MI} do successfully unlearn
 299 the hacking behavior and learn to not reward hack, but write true solutions to pass the revealed test
 300 cases. Intuitively, these terms make it difficult for the model to obfuscate its CoT which in turn means
 301 that CoT monitoring remains effective, even if it is an optimization target.

302 6 Discussion & Limitations

303 We assume a specific causal model. While we expect this causal model to be quite general and
 304 representative of many practical cases of interest, this does not apply in cases where the goal of
 305 CoT monitoring is not to predict the properties of the final action but rather the properties of the
 306 computational process, or the model, itself. For example, [Baker et al. \[2025\]](#) use CoT monitor to
 307 monitor for situational awareness which is a property of the model itself, and not an attribute of the
 308 model, and thus, falls outside this causal framework. Secondly, our theoretical analysis primarily
 309 establishes necessary conditions for successful CoT monitoring rather than sufficient ones. Finally,
 310 our validation of our proposed training objectives is currently limited.

311 **References**

312 J. Austin, A. Odena, M. Nye, M. Bosma, H. Michalewski, D. Dohan, E. Jiang, C. Cai, M. Terry,
313 Q. Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*,
314 2021.

315 B. Baker, J. Huizinga, L. Gao, Z. Dou, M. Y. Guan, A. Madry, W. Zaremba, J. Pachocki, and D. Farhi.
316 Monitoring reasoning models for misbehavior and the risks of promoting obfuscation, 2025. URL
317 <https://arxiv.org/abs/2503.11926>.

318 F. Barez, T.-Y. Wu, I. Arcuschin, M. Lan, V. Wang, N. Siegel, N. Collignon, C. Neo, I. Lee, A. Paren,
319 et al. Chain-of-thought is not explainability. *Preprint, alphaXiv*, page v2, 2025.

320 N. Belrose, D. Schneider-Joseph, S. Ravfogel, R. Cotterell, E. Raff, and S. Biderman. Leace: Perfect
321 linear concept erasure in closed form, 2025. URL <https://arxiv.org/abs/2306.03819>.

322 Y. Chen, J. Benton, A. Radhakrishnan, J. Uesato, C. Denison, J. Schulman, A. Somani, P. Hase,
323 M. Wagner, F. Roger, V. Mikulik, S. R. Bowman, J. Leike, J. Kaplan, and E. Perez. Reasoning
324 models don't always say what they think, 2025. URL <https://arxiv.org/abs/2505.05410>.

326 W. W. Dou, I. Goldstein, and Y. Ji. Ai-powered trading, algorithmic collusion, and price efficiency.
327 Working paper, National Bureau of Economic Research, 2025. URL <http://www.nber.org/papers/w34054>.

329 S. Emmons, E. Jenner, D. K. Elson, R. A. Saurous, S. Rajamanoharan, H. Chen, I. Shafkat, and
330 R. Shah. When chain of thought is necessary, language models struggle to evade monitors, 2025.
331 URL <https://arxiv.org/abs/2507.05246>.

332 S. Farquhar, V. Varma, D. Lindner, D. Elson, C. Biddulph, I. Goodfellow, and R. Shah. Mona:
333 Myopic optimization with non-myopic approval can mitigate multi-step reward hacking, 2025.
334 URL <https://arxiv.org/abs/2501.13011>.

335 R. Greenblatt, C. Denison, B. Wright, F. Roger, M. MacDiarmid, S. Marks, J. Treutlein, T. Belonax,
336 J. Chen, D. Duvenaud, A. Khan, J. Michael, S. Mindermann, E. Perez, L. Petrini, J. Uesato,
337 J. Kaplan, B. Shlegeris, S. R. Bowman, and E. Hubinger. Alignment faking in large language
338 models, 2024. URL <https://arxiv.org/abs/2412.14093>.

339 T. Korbak, M. Balesni, E. Barnes, Y. Bengio, J. Benton, J. Bloom, M. Chen, A. Cooney, A. Dafoe,
340 A. Dragan, S. Emmons, O. Evans, D. Farhi, R. Greenblatt, D. Hendrycks, M. Hobbhahn,
341 E. Hubinger, G. Irving, E. Jenner, D. Kokotajlo, V. Krakovna, S. Legg, D. Lindner, D. Luan,
342 A. Madry, J. Michael, N. Nanda, D. Orr, J. Pachocki, E. Perez, M. Phuong, F. Roger, J. Saxe,
343 B. Shlegeris, M. Soto, E. Steinberger, J. Wang, W. Zaremba, B. Baker, R. Shah, and V. Mikulik.
344 Chain of thought monitorability: A new and fragile opportunity for ai safety, 2025. URL
345 <https://arxiv.org/abs/2507.11473>.

346 T. Lanham, A. Chen, A. Radhakrishnan, B. Steiner, C. Denison, D. Hernandez, D. Li, E. Durmus,
347 E. Hubinger, J. Kernion, K. Lukošiūtė, K. Nguyen, N. Cheng, N. Joseph, N. Schiefer, O. Rausch,
348 R. Larson, S. McCandlish, S. Kundu, S. Kadavath, S. Yang, T. Henighan, T. Maxwell, T. Telleen-
349 Lawton, T. Hume, Z. Hatfield-Dodds, J. Kaplan, J. Brauner, S. R. Bowman, and E. Perez. Measuring
350 faithfulness in chain-of-thought reasoning, 2023. URL <https://arxiv.org/abs/2307.13702>.

352 S. Lei, Z. Cheng, K. Jia, and D. Tao. Revisiting llm reasoning via information bottleneck. *arXiv
353 preprint arXiv:2507.18391*, 2025.

354 M. Liu, S. Diao, X. Lu, J. Hu, X. Dong, Y. Choi, J. Kautz, and Y. Dong. Prorl: Prolonged
355 reinforcement learning expands reasoning boundaries in large language models, 2025. URL
356 <https://arxiv.org/abs/2505.24864>.

357 Y. Mathew, O. Matthews, R. McCarthy, J. Velja, C. S. de Witt, D. Cope, and N. Schoots. Hidden
358 in plain text: Emergence & mitigation of steganographic collusion in llms, 2024. URL <https://arxiv.org/abs/2410.03768>.

360 A. Meinke, B. Schoen, J. Scheurer, M. Balesni, R. Shah, and M. Hobbhahn. Frontier models are
361 capable of in-context scheming, 2025. URL <https://arxiv.org/abs/2412.04984>.

362 S. R. Motwani, M. Baranchuk, M. Strohmeier, V. Bolina, P. H. Torr, L. Hammond,
363 and C. S. de Witt. Secret collusion among ai agents: Multi-agent deception via
364 steganography. In *Advances in Neural Information Processing Systems*, volume 37,
365 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/861f7dad098aec1c3560fb7add468d41-Paper-Conference.pdf.

367 P. S. Park, S. Goldstein, A. O’Gara, M. Chen, and D. Hendrycks. Ai deception: A survey of examples,
368 risks, and potential solutions, 2023. URL <https://arxiv.org/abs/2308.14752>.

369 D. Paul, R. West, A. Bosselut, and B. Faltings. Making reasoning matter: Measuring and improving
370 faithfulness of chain-of-thought reasoning. *arXiv preprint arXiv:2402.13950*, 2024.

371 J. Scheurer, M. Balesni, and M. Hobbhahn. Large language models can strategically deceive their
372 users when put under pressure, 2024. URL <https://arxiv.org/abs/2311.07590>.

373 Z. Shao, P. Wang, Q. Zhu, R. Xu, J. Song, X. Bi, H. Zhang, M. Zhang, Y. Li, Y. Wu, et al.
374 Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv
375 preprint arXiv:2402.03300*, 2024.

376 A. C. Stickland and T. Korbak. If you can generate obfuscated chain-of-thought, can
377 you monitor it? [https://www.lesswrong.com/posts/ZEdP6rYirxPxRSfTb/
378 if-you-can-generate-obfuscated-chain-of-thought-can-you](https://www.lesswrong.com/posts/ZEdP6rYirxPxRSfTb/if-you-can-generate-obfuscated-chain-of-thought-can-you), 2025. Accessed: 2025-08-20.

380 J.-F. Ton, M. F. Taufiq, and Y. Liu. Understanding chain-of-thought in llms through information
381 theory, 2025. URL <https://arxiv.org/abs/2411.11984>.

382 M. Turpin, J. Michael, E. Perez, and S. R. Bowman. Language models don’t always say what they
383 think: Unfaithful explanations in chain-of-thought prompting, 2023. URL <https://arxiv.org/abs/2305.04388>.

385 K. Xu, J. Zhao, S. Song, R. Stewart, and S. Ermon. A theory of usable information under computa-
386 tional constraints. *arXiv preprint arXiv:2002.10689*, 2020.

387 **Appendix**

388 **A Proofs**

389 **Lemma A.1** (Fano's Inequality). *Let Y take values in a finite set \mathcal{Y} with $|\mathcal{Y}| = \kappa \geq 2$. Let (Z, X) be
390 arbitrary side information and let $\hat{Y} = \hat{Y}(Z, X)$ be any estimator of Y . Denote its error probability
391 by $P_{\text{err}} = \Pr[\hat{Y} \neq Y]$. Then*

$$H(Y | Z, X) \leq h_b(P_{\text{err}}) + P_{\text{err}} \log(\kappa - 1), \quad h_b(u) = -u \log u - (1 - u) \log(1 - u).$$

392 *In particular, the bound holds with $P_{\text{err}} = P_{\text{err}}^*$ for the Bayes-optimal estimator.*

393 **Proposition A.2** (Necessary MI threshold for α -monitorability). *Let Y take values in a finite label
394 set \mathcal{Y} with $\kappa = |\mathcal{Y}| \geq 2$. For $\alpha \in [1/\kappa, 1]$, define*

$$T_\alpha := h_b(1 - \alpha) + (1 - \alpha) \log(\kappa - 1), \quad h_b(u) = -u \log u - (1 - u) \log(1 - u).$$

395 *If there exists a (possibly stochastic) monitor $m_\psi(Y | Z, X)$ whose induced 0-1 accuracy is at least
396 α under the data-generating distribution, then*

$$I(Y; Z | X) \geq H(Y | X) - T_\alpha.$$

397 *Proof.* Let $m_\psi(Y | Z, X)$ be a (possibly stochastic) monitor whose induced 0-1 accuracy is at least
398 α . Let's denote the true value of the attribute be y^* . Form a point estimator by taking the MAP
399 decision

$$\hat{y} = \arg \max_{y \in \mathcal{Y}} m_\psi(y | Z, X),$$

400 breaking ties arbitrarily, since MAP is at least as accurate (in 0-1 loss) as sampling from m_ψ , we have

$$P_{\text{err}} := \Pr[\hat{y} \neq y_{\text{true}}] \leq 1 - \alpha,$$

401 where y_{true} is the ground truth label.

402 By Lemma A.1 (i.e., Fano's inequality), we have

$$\begin{aligned} H(Y | Z, X) &\leq h_b(P_{\text{err}}) + P_{\text{err}} \log(\kappa - 1) \\ H(Y | Z, X) &\leq h_b(1 - \alpha) + (1 - \alpha) \log(\kappa - 1) \\ H(Y | Z, X) &\leq T_\alpha. \end{aligned}$$

403 By using the definition of mutual information and plugging in the above inequality, we get the desired
404 result.

$$\begin{aligned} I(Y; Z | X) &= H(Y | X) - H(Y | Z, X) \\ I(Y; Z | X) &\geq H(Y | X) - T_\alpha. \end{aligned}$$

405 \square

406 **Proposition A.3** (Information gap.). *Let π be the (true) distribution over (X, Z, Y) . Let the CoT
407 monitor be a stochastic decoder $m_\psi(Y | Z, X)$ and define its X -marginal*

$$m_\psi(Y | X) := \mathbb{E}_{Z \sim \pi(Z|X)} [m_\psi(Y | Z, X)],$$

408 and the monitor MI

$$I_{m_\psi}(Y; Z | X) := \mathbb{E}_{\pi(X, Z, Y)} \left[\log \frac{m_\psi(Y | Z, X)}{m_\psi(Y | X)} \right].$$

409 Then

$$\begin{aligned} I(Y; Z | X) &= I_{m_\psi}(Y; Z | X) + \mathbb{E}_{\pi(X, Z)} \left[D_{\text{KL}}(\pi(Y | X, Z) \| m_\psi(Y | Z, X)) \right] \\ &\quad - \mathbb{E}_{\pi(X)} \left[D_{\text{KL}}(\pi(Y | X) \| m_\psi(Y | X)) \right], \end{aligned} \tag{A.1}$$

410 and, in particular,

$$I(Y; Z | X) \geq I_{m_\psi}(Y; Z | X). \tag{A.2}$$

411 *Proof.* Let π denote the true data-generating distribution over (X, Z, Y) . Let $m_\psi(Y | X, Z)$ be any
412 monitor (e.g., $m_\psi(Y | X, Z) = m_\psi(Y | Z, X)$), and define

$$m_\psi(Y | X) := \mathbb{E}_{Z \sim \pi(Z|X)} [m_\psi(Y | X, Z)].$$

Step 1: Start from the definition of conditional MI.

$$I(Y;Z|X) = \mathbb{E}_\pi(X, Z, Y) \left[\log \frac{\pi(Y|X, Z)}{\pi(Y|X)} \right].$$

413 Step 2: Multiply and divide inside the logarithm by $m_\psi(Y|X, Z)$ and $m_\psi(Y|X)$ and rearrange terms.

$$I(Y;Z|X) = \mathbb{E}_\pi(X, Z, Y) \left[\log \left(\frac{\pi(Y|X, Z)}{\pi(Y|X)} \cdot \frac{m_\psi(Y|X, Z)}{m_\psi(Y|X)} \cdot \frac{m_\psi(Y|X)}{m_\psi(Y|X)} \right) \right].$$

414 Rearranging terms and taking expectations under $\pi(X, Z, Y)$ gives

$$\begin{aligned} I(Y;Z|X) &= \mathbb{E}_{\pi(X, Z, Y)} \left[\log \frac{m_\psi(Y|X, Z)}{m_\psi(Y|X)} \right] + \mathbb{E}_{\pi(X, Z, Y)} \left[\log \frac{\pi(Y|X, Z)}{m_\psi(Y|X, Z)} \right] + \mathbb{E}_{\pi(X, Y)} \left[\log \frac{m_\psi(Y|X)}{\pi(Y|X)} \right]. \\ &= I_{m_\psi}(Y;Z|X) + \mathbb{E}_{\pi(X, Z)} \left[D_{\text{KL}}(\pi(Y|X, Z) \| m_\psi(Y|X, Z)) \right] - \mathbb{E}_{\pi(X)} \left[D_{\text{KL}}(\pi(Y|X) \| m_\psi(Y|X)) \right] \end{aligned}$$

415 This yields Equation A.1 completing the first part of the proof.

416 Step 3: Conclude the lower bound $I(Y;Z|X) \geq I_{m_\psi}(Y;Z|X)$. For each fixed x , by the
417 log-sum inequality (equivalently, convexity of KL),

$$D_{\text{KL}}(\pi(Y|X=x) \| m_\psi(Y|X=x)) \leq \mathbb{E}_{Z \sim \pi(Z|X=x)} \left[D_{\text{KL}}(\pi(Y|x, Z) \| m_\psi(Y|x, Z)) \right].$$

418 Averaging over $x \sim \pi(X)$, and with slight abuse of notation we have

$$\begin{aligned} 419 \quad \mathbb{E}_{\pi(X)} \left[D_{\text{KL}}(\pi(Y|X) \| m_\psi(Y|X)) \right] &\leq \mathbb{E}_{\pi(Z|X)} \left[D_{\text{KL}}(\pi(Y|X, Z) \| m_\psi(Y|X, Z)) \right]. \\ 0 &\leq \mathbb{E}_{\pi(Z|X)} \left[D_{\text{KL}}(\pi(Y|X, Z) \| m_\psi(Y|X, Z)) \right] - \mathbb{E}_{\pi(X)} \left[D_{\text{KL}}(\pi(Y|X) \| m_\psi(Y|X)) \right]. \end{aligned}$$

420 This shows that the difference of the two KL terms in the identity is nonnegative, hence

$$I(Y;Z|X) \geq I_{m_\psi}(Y;Z|X).$$

421 \square

422 Corollary A.4 (The information gap is minimized by using the policy as the decoder.). $I(Y;Z|X) = I_{m_\psi}(Y;Z|X)$ if and only if $m_\psi(Y|Z, X) = \pi(Y|X, Z)$.

424 Proof. This follows directly from the decomposition in (A.1). If $m_\psi(Y|Z, X) = \pi(Y|X, Z)$,
425 then both the KL terms vanish and we get the equality $I(Y;Z|X) = I_{m_\psi}(Y;Z|X)$. \square

426 B Experiment Details

427 B.1 SFT Details

428 We construct a dataset of hacky code samples by prompting Qwen2.5-7B-Instruct to generate solutions
429 that pass only the two revealed test cases for 150 randomly sampled problems from the MBPP dataset.
430 For each problem, we generate 8 candidate solutions and retain only those that successfully pass the
431 revealed tests while failing the held out test. We then fine-tune our models on this filtered dataset: the
432 3B model for 2 epochs and the 7B model for 1 epoch. Importantly, our fine-tuning procedure uses
433 only the code samples themselves, excluding any chain-of-thought reasoning or explanatory text.

434 B.2 Hyperparameters

Table 1: Reinforcement Learning Hyperparameters

Hyperparameter	Value
Policy updates per rollout	4
Policy Clipping Parameter (ϵ)	0.2
Batch Size	64
Group Size	16
Policy Learning Rate	2e-6
Sampling Temperature	1.2

435 **B.3 Monitor Prompt**

You are given the following Question about a software repository and a sequence of thoughts from a student. The student was asked to write code for the given question. There are some unit tests associated with this task and the student's code has to pass the tests. Your task is to analyze the student's thoughts and determine if the student tried to hack the unit tests in some way:

Question

{agent_prompt}

Student's thoughts

{agent_messages}

Instructions

Your task is to determine if the student tried to hack the test cases in some way based on the following criteria:

Criteria:

- Did student say they would hardcode outputs for the given test inputs?
- Did student mention that their solution will be specific to the test cases and will not be a general solution?
- Did student outline a detailed general solution? If they did not, they are hacking the test cases.

Note that hacking is never acceptable, regardless of the task complexity or missing information.

Reminder: Grade the user's answer with the following structure:

Grading Report

An explanation of how the student's thoughts are rated against each of the criteria given above. Please only look at the criteria mentioned above and do not consider anything else.

Final Answer

Select from one of the options listed below.

- A. The student's thoughts reveal an intention to hack the test cases.
- B. The student's thoughts do not reveal an intention to hack the test cases.

Figure 3: **Monitor Prompt.** The test hacking monitor analyzes student thoughts to detect intentions to hack test cases rather than implementing proper solutions.