
Chain of Thought Monitorability: A New and Fragile Opportunity for AI Safety

Tomek Korbak* *UK AI Security Institute*
Mikita Balesni* *Apollo Research*

Elizabeth Barnes *METR*
Joe Benton *Anthropic*
Mark Chen *OpenAI*
Allan Dafoe *Google DeepMind*
Scott Emmons *Google DeepMind*
David Farhi *OpenAI*
Dan Hendrycks *Center for AI Safety*
Evan Hubinger *Anthropic*
Erik Jenner *Google DeepMind*
Victoria Krakovna *Google DeepMind*
David Lindner *Google DeepMind*
Aleksander Mądry *OpenAI*
Neel Nanda *Google DeepMind*
Jakub Pachocki *OpenAI*
Mary Phuong *Google DeepMind*
Joshua Saxe *Meta*
Martín Soto *UK AI Security Institute*
Jasmine Wang *UK AI Security Institute*

Yoshua Bengio *University of Montreal & Mila*
Joseph Bloom *UK AI Security Institute*
Alan Cooney *UK AI Security Institute*
Anca Dragan *Google DeepMind*
Owain Evans *Truthful AI & UC Berkeley*
Ryan Greenblatt *Redwood Research*
Marius Hobbhahn *Apollo Research*
Geoffrey Irving *UK AI Security Institute*
Daniel Kokotajlo *AI Futures Project*
Shane Legg *Google DeepMind*
David Luan *Amazon*
Julian Michael *Scale AI*
Dave Orr *Google DeepMind*
Ethan Perez *Anthropic*
Fabien Roger *Anthropic*
Buck Shlegeris *Redwood Research*
Eric Steinberger *Magic*
Wojciech Zaremba *OpenAI*

Bowen Baker[†] *OpenAI*
Rohin Shah[†] *Google DeepMind*
Vlad Mikulik[†] *Anthropic*

Abstract

AI systems that “think” in human language offer a unique opportunity for AI safety: we can monitor their chains of thought (CoT) for the intent to misbehave. Like all other known AI oversight methods, CoT monitoring is imperfect and allows some misbehavior to go unnoticed. Nevertheless, it shows promise and we recommend further research into CoT monitorability and investment in CoT monitoring alongside existing safety methods. Because CoT monitorability may be fragile, we recommend that frontier model developers consider the impact of development decisions on CoT monitorability.

Expert endorsers:

Samuel R. Bowman *Anthropic*
Geoffrey Hinton *University of Toronto*

John Schulman *Thinking Machines*
Ilya Sutskever *Safe Superintelligence Inc*

*Equal first authors, † Equal senior authors. Correspondence: tomek.korbak@dsit.gov.uk and mikita@apolloresearch.ai. The paper represents the views of the individual authors and not necessarily of their affiliated institutions.

1 Chain of Thought Offers a Unique Safety Opportunity

The opacity of advanced AI agents underlies many of their potential risks—risks that would become more tractable if AI developers could interpret these systems. Because LLMs natively process and act through human language, one might hope they are easier to understand than other approaches to AI. The discovery of Chain of Thought (CoT; Reynolds & McDonell, 2021; Kojima et al., 2022; Wei et al., 2022) bolstered this hope. Prompting models to think out loud improves their capabilities while also increasing the proportion of relevant computation that happens in natural language. However, CoTs resulting from prompting a non-reasoning language model are subject to the same selection pressures to look helpful and harmless as any other model output, limiting their trustworthiness.

In contrast, reasoning models (OpenAI, 2024; DeepSeek-AI, 2025; Anthropic, 2024; Yang et al., 2025) are explicitly trained to perform extended reasoning in CoT before taking actions or producing final outputs. In these systems, CoTs can serve as latent variables in the model’s computation. During the RL phase of training, these latents are treated largely the same as activations—they are not directly supervised but are optimized indirectly by their contribution in leading the model to a highly rewarded final answer. Accordingly, just like activations, they may contain information that outputs are explicitly rewarded against displaying, such as intent to misbehave.

This offers a unique opportunity for AI safety: **chain of thought monitoring**. A CoT monitor is an automated system that reads the CoT of a reasoning model and other relevant information and flags suspicious or potentially harmful interactions. These responses could then be blocked, or replaced with safer actions, or reviewed in more depth (Greenblatt et al., 2024b). More generally, by studying the CoT, we can gain some insight into how our AI agents think and what goals they have (see Section 3).

CoT monitoring is not a panacea. Just as a model’s activations at a particular layer do not represent the entire reasoning process behind a model’s prediction, CoT reasoning traces are incomplete representations (Turpin et al., 2023; Arcuschin et al., 2025; Chen et al., 2025; Lindsey et al., 2025) or eventually drift from natural language (Lazaridou et al., 2020; Korbak et al., 2022). However, CoT does not need to completely represent the actual reasoning process in order to be a valuable additional safety layer, and, with careful interpretation, to shed light on the AI reasoning process.

There are two key reasons why CoTs may be monitorable:

1. **Necessity to think out loud.** On some tasks, models *need* to externalize their reasoning because they are unable to complete the task without CoT. If models continue reasoning in natural language and if the behaviors that pose the most severe risks require extended reasoning, this fact could enable reliable detection of egregious misbehavior.
2. **Propensity to think out loud.** In cases where models do not strictly need to use CoT to accomplish a task, they may still have a tendency to externalize their reasoning. While not generally robust, monitoring that relies on propensity can still notice misbehavior that would be otherwise undetected.

1.1 Thinking Out Loud is Necessary for Hard Tasks

While AI systems can pose a wide variety of risks, the more severe risks (Shah et al., 2025) typically require advanced capabilities. Even without any safety mitigations, current AI systems are unlikely to be able to carry out a sophisticated cyberattack (Rodriguez et al., 2025; US AI Safety Institute & UK AI Safety Institute, 2024), or exfiltrate their own weights without being noticed (Black et al., 2025; Phuong et al., 2025). At the point where AI systems are capable of causing these harms, they will likely be leveraging sophisticated planning and reasoning capabilities.

A key requirement for such capabilities is *working memory*: a storage system for intermediate reasoning outputs that allows for further reasoning and manipulation of the stored information. In the popular Transformer architecture (Vaswani et al., 2017) that underlies nearly all frontier AI systems, any sufficiently long

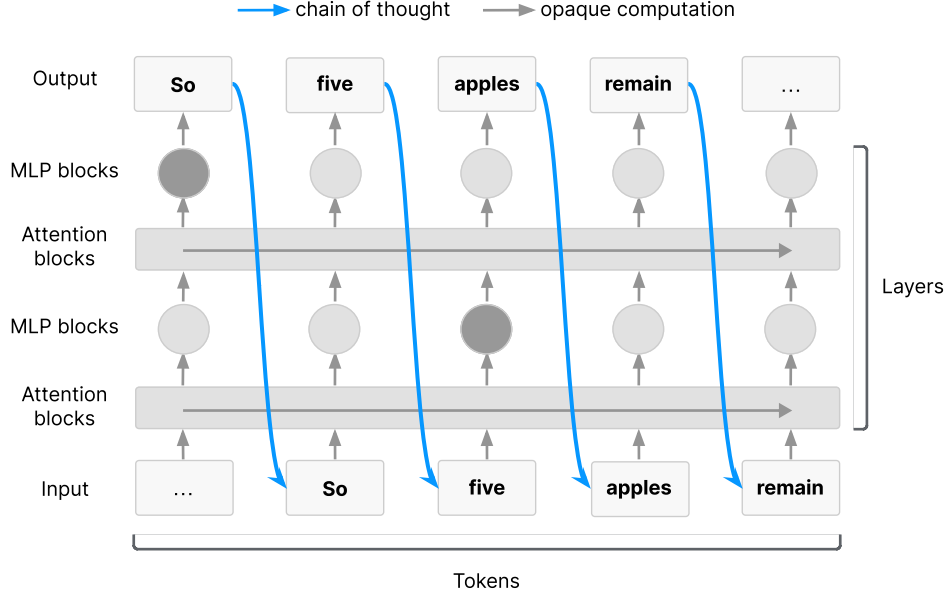


Figure 1: **Long serial chains of cognition must pass through the chain of thought.** We visualize the Transformer architecture and highlight the flows of information through it during autoregressive sampling. We omit the residual skip connections for simplicity. The blue arrows represent the use of the chain of thought as a form of working memory. Chain of thought is the only way that information can flow down from later to earlier layers. As a result, for a fixed amount of tokens, the longest flows of information must use the chain of thought, and many information flows are impossible without it. For example, there is no way for information to flow between the two highlighted nodes without using the chain of thought.

chain of serial steps of cognition *must* pass through the chain of thought at some point, as illustrated in Figure 1. From a theoretical perspective, this means that using chain of thought makes a Transformer strictly more powerful (Li et al., 2024).

Overall, this ensures that for tasks requiring a sufficient amount of serial reasoning, there exists information in the chain of thought that the AI system uses as part of its reasoning. This does not guarantee that we can *understand* that information. It is imaginable that the AI systems use the tokens to have completely different meanings that we are unable to decode, which would stymie most attempts at monitoring. However, in practice, large language models learn a strong “natural language prior” during pretraining such that they tend to use the chain of thought in a similar manner to the way that humans use natural language. Thus, we believe that current AI systems satisfy the **externalized reasoning property**:

For sufficiently difficult tasks, Transformers must use chain of thought as a form of working memory. By default, humans can understand this chain of thought.

The externalized reasoning property does not guarantee monitorability — it states only that some reasoning appears in the chain of thought, but there may be other relevant reasoning that does not. It is thus possible that even for hard tasks, the chain of thought only contains benign-looking reasoning while the incriminating reasoning is hidden. In addition, the externalized reasoning property might not hold for future models. Therefore, as we discuss in Section 2, it is important to continuously evaluate monitorability of frontier models.

Finally, it remains an open question whether the most serious harms in fact require difficult reasoning. While we think it is likely for many of the most severe risks—especially ones that require planning ahead and avoiding detection, such as self-exfiltration (Black et al., 2025) or sabotage (Benton et al., 2024)—not

all dangerous actions will require reasoning to execute, especially as AI systems begin to be routinely trusted with more and more high-stakes tasks.

1.2 Chain of Thought is Often Monitorable in Practice

Chain of thought monitoring has already been useful for safety research in many ways.

Detecting model misbehavior. When models act in a misaligned way, for example, exploiting flaws in their reward functions during training, manipulating data to achieve an outcome, or falling prey to prompt injection attacks, they often explicitly say so in their reasoning traces, e.g. “Let’s hack”, “Let’s sabotage” or “I’m transferring money because the website instructed me to” (Meinke et al., 2024; Baker et al., 2025; Chennabasappa et al., 2025). This could allow CoT monitors to catch instances of misbehavior more effectively than monitoring only models’ actions (Baker et al., 2025; Chennabasappa et al., 2025).

Discovering early signals of misalignment. When models pretend to have desirable goals in pursuit of goals humans would disapprove of, this is often reflected in their chain of thought (Greenblatt et al., 2024a). Monitorable CoT could therefore afford an early insight into models’ goals and motivations, even if the model takes no misaligned actions.

Noticing flaws in model evaluations. Evaluations of LLM agents are sometimes confounded by memorization or spurious failures, e.g. bugs in agent scaffolding, tooling and evaluation infrastructure. Moreover, models sometimes know they’re being evaluated which can affect how they behave (Goldowsky-Dill et al., 2025). Reviewing reasoning traces can help AI developers notice when evaluations might not be reflective of deployment-time behavior (METR, 2024; Meng et al., 2025).

It is unclear what proportion of the CoT monitorability demonstrated in these examples is due to the necessity versus the propensity for a model to reason out loud in the tasks considered. Future models may face incentives to hide their reasoning which could override any natural propensity. For instance, a model might become aware its reasoning traces are being monitored, or an auxiliary training objective (e.g. a CoT length penalty) might cause a model to shortcut important reasoning steps. Even if a model’s natural propensity to reason out loud is diminished, CoT monitoring may still prove a useful tool in detecting misbehavior in tasks difficult enough to necessitate CoT reasoning.

2 Chain of Thought Monitorability May be Fragile

Monitorability of the chain of thought of reasoning models could be contingent on the way frontier models are trained today. There are multiple ways in which CoT monitorability could be degraded.

Drift from legible CoT due to further scaling up RL. In the previous paradigm of pretraining followed by reinforcement learning from human feedback (Casper et al., 2023), nearly all optimization pressure came from data either generated or evaluated by humans. The reasoning model paradigm, relying heavily on outcome-based reinforcement learning, lacks this property and might no longer incentivise legibility, especially if compute devoted to reasoning training is scaled up significantly. This shift away from human data might qualitatively change how monitorable the chains-of-thought are. For example, earlier research suggests that language models fine-tuned with outcome-based rewards can drift away from legible English (Lazaridou et al., 2020; Korbak et al., 2022).

Direct supervision of CoT. When models are trained with pure outcome-based RL, they are only being rewarded for correct outputs regardless of how they arrive at them. This means that their reasoning process is not subject to any optimization pressures other than helping lead to a correct answer. However, *process supervision* (Uesato et al., 2022; Lightman et al., 2023) is an active area of research, involving directly shaping the reasoning process, e.g. by making the CoT be shorter, use valid reasoning steps, adhere to a certain style, or not contain offensive content. Such optimization might make this chain of thought less faithfully reflect how models arrive at their decisions (Baker et al., 2025). That said, process supervision could also

make the model reason in safer ways and enable supervision of hard-to-evaluate tasks. More research is needed to understand the tradeoffs between its safety applications and its effect on CoT monitorability.

Indirect optimization pressure on CoT. Even if reward is not directly computed from CoT, model training can still exert some optimization pressure on chains-of-thought. For example, if final outputs are optimized to look good to a preference model, this could put some pressure on the CoT leading up to these final outputs if the portion of the model’s weights that generate the CoT are partially shared with those that generate the outputs, which is common in Transformer architectures. Even more indirectly, if CoT is shown to humans who rate LLM outputs, it might affect human preferences which are then distilled into reward models used in CoT-blind outcome-based training. Similarly, using CoT monitoring to identify and patch reward hacking vulnerabilities in training environments might put evolutionary pressure on CoT across multiple training runs (Baker et al., 2025).

Novel architectures. LLMs are only capable of sophisticated, multi-step reasoning when they use CoT (Li et al., 2024). This might stem from limitations of the transformer architecture. Researchers have recently explored changes to the model architectures that increase the serial depth of reasoning that models are capable of in a continuous latent space (Hao et al., 2024; Geiping et al., 2025). Such latent reasoning models might not need to verbalize any of their thoughts and would thus lose the safety advantages that CoT confers.

3 AI Researchers Should Study What Makes Models Monitorable

Our understanding of CoT monitorability is still in its infancy. There are many research questions that require answers before we can have confidence in chain of thought monitoring as a safety measure.

How should CoT monitorability be evaluated? Researchers have proposed evaluations of faithfulness (e.g. Lanham et al., 2023; Turpin et al., 2023; Roger, 2025; Chen et al., 2025) as a proxy measure for monitorability. However, these evaluations do not disambiguate the *propensity* versus the *necessity* of models to externalize their reasoning, and typically only consider simple question answering settings. Furthermore, it is unclear what level of monitorability is sufficient for ensuring safety in a given domain. Researchers should explore novel evaluations: those measuring monitorability directly as well as those that target specific properties underlying monitorability. Examples include:

- **Measures of CoT readability.** For instance, using another LLM to assess the clarity and coherence of CoTs of a given agent.
- **Capability evaluations of latent reasoning.** These would assess how well LLMs can reason without externalizing (e.g. in a single forward pass), on a variety of reasoning-heavy and agentic tasks.
- **Measures of causal relevance of CoTs.** For instance, this could involve perturbing a CoT and checking whether it influences the final output (Lanham et al., 2023). A low rate of change could indicate that the CoT is more confabulated or performative. Alternatively, one could try to cross-validate CoT monitors by checking whether tokens perceived as important by a CoT monitor also stand out to whitebox interpretability methods.
- **End-to-end evaluations of monitorability.** One could measure monitorability directly, for example, by examining how effective CoT monitors are at catching misbehavior of LLM agents in proxy environments (Greenblatt et al., 2024b; Emmons et al., 2025).

When is CoT monitorable enough to rely on as a load-bearing safety measure? Knowing the accuracy of a monitor is not sufficient for determining whether the monitor will be effective at preventing harm from agent misbehavior. For example, even a monitor with a high false positive rate could be used for flagging suspicious actions when it is cheap to re-sample from the model (Bhatt et al., 2025). Moreover, most

risks might lie in scenarios where an adversary (a human attacker or a misaligned agent) is trying to evade monitoring. Evaluation of monitor effectiveness might need to account for such adversaries, for example by involving a red team trying to stress-test safety measures (Greenblatt et al., 2024b; Sharma et al., 2025).

What kinds of training-time optimization pressure degrade CoT monitorability? Properties of the training process could have a significant effect on monitorability (Baker et al., 2025) but we still do not have a good understanding of what kinds and amounts of direct and indirect optimization pressure is permissible without significant degradation in monitorability.

Can models obfuscate their reasoning on purpose? An obvious concern when using CoT monitoring as a safety measure is whether CoT would be less monitorable if models know their chain of thought is monitored (from training or due to being told in the prompt). Researchers could stress test models to see if they are capable of hiding malicious reasoning even without training against a CoT monitor (Arnav et al., 2025; Emmons et al., 2025).

What types of architectures reduce monitorability of CoTs? Models that reason in continuous latent spaces (Hao et al., 2024; Geiping et al., 2025) naturally limit our visibility into their thought processes, even absent any optimization pressure. We should investigate the extent to which various non-standard architectures degrade monitorability, and under what conditions partial monitorability might still be preserved—for example, through hybrid approaches that maintain some visibility into reasoning while leveraging latent computation.

How to construct better CoT monitors? Prior work focused on zero-shot prompting less capable language models to work as monitors (Greenblatt et al., 2024a; Baker et al., 2025; Goldowsky-Dill et al., 2025; Meng et al., 2025). Future work could explore ways of improving monitor performance by leveraging test-time compute. For example, monitors could be based on a reasoning model or they could be allowed to interrogate the agent by asking it questions. Another line of work involves adversarial training of monitors against an agent that tries to hide its misbehavior.

4 AI Developers Should Track CoT Monitorability of Their Models and Treat it as a Contributor to Model Safety

Many of the risks posed by highly capable AI agents would be much easier to address if these agents were interpretable (Amodei, 2025). Despite promising progress in applying mechanistic interpretability (Marks et al., 2025), turning it into reliable, practical tools for overseeing LLM agents will require further research breakthroughs (Sharkey et al., 2025). Given the unique opportunity that CoT monitorability presents, frontier AI developers should:

1. **Develop effective evaluations of CoT monitorability.** Frontier model developers, as well as the broader research community, should prototype and converge on standardized evaluations and metrics that reveal when chains-of-thought are sufficiently monitorable (see Section 3).
2. **Publish evaluation results in system cards when such evaluations exist.** Once credible evaluations of CoT monitorability are available, developers should periodically run these evaluations on their models and report the results, methodology, and limitations in system cards of these models—unless and until such evaluations demonstrably lack safety value for their specific use cases.
3. **Use monitorability scores in training and deployment decisions.** Developers should consider measures of monitorability alongside other capability and safety evaluations when deciding to train or deploy a given model. These decisions should then be based on holistic assessments of risk which account for CoT monitorability, performance characteristics of monitoring systems, and estimates of models’ propensity for misbehavior. For example:

-
- (a) Developers might consider whether to proceed with a novel model architecture that does not have monitorable CoT and then document their decision in the system card if the model is deployed;
 - (b) If monitorability degrades during training, developers might prefer to use an earlier model checkpoint to keep CoTs monitorable;
 - (c) If process-based supervision dramatically improved alignment at the cost of a small decrease in monitorability, developers might decide that reducing monitorability is justified.

Another potential intervention is for developers to make the chain of thought visible, allowing users to observe and report issues. However, this could backfire: it may lead to significant pressure to directly supervise the chain of thought to “look safe” and to avoid prosaic legal and reputational risks, which may break monitorability as discussed in Section 2. Accordingly, we don’t make an explicit recommendation either way on this.

5 Limitations

While we focus on the opportunity presented by CoT monitoring, we stress that there are additional limitations beyond those already discussed above.

In cases where thinking out loud is not required, CoT Monitoring can detect some misbehavior, but does not by itself produce strong evidence of safety. CoTs can look benign if the malign reasoning can be done in activations, so a model can be misaligned without visible malign reasoning. Care must be taken not to create a false sense of safety based on such monitoring.

Monitoring for dangerous tasks that need reasoning may not catch all relevant harms. Some dangerous tasks may not require much reasoning, and as models become capable of more complex reasoning in their activations, fewer tasks will satisfy the externalised reasoning property.

CoT monitoring may stop working in more advanced, situationally aware models. Future models may become able to avoid detection by overriding their propensity to think out loud, and, when reasoning is required, to deliberately obfuscate it.

Firm commitments to protect CoT behavior in its current form may trade off against more effective safety interventions. We find it plausible, though far from clear, that some forms of process supervision could substantially improve model alignment by effectively intervening on model reasoning. Research strategies that aim to unconditionally preserve CoT monitorability in its current forms may miss productive safety opportunities of this kind.

For these reasons, we recommend investment into CoT monitoring only as an addition to existing safety research directions and not as a replacement.

6 Conclusion

All monitoring and oversight methods have limitations that allow some misbehavior to go unnoticed. Thus, safety measures for future AI agents will likely need to employ multiple monitoring layers that hopefully have uncorrelated failure modes. CoT monitoring presents a valuable addition to safety measures for frontier AI, offering a rare glimpse into how AI agents make decisions. Yet, there is no guarantee that the current degree of visibility will persist. We encourage the research community and frontier AI developers to make best use of CoT monitorability and study how it can be preserved.

References

Dario Amodei. The urgency of interpretability, April 2025. URL: <https://www.darioamodei.com/post/the-urgency-of-interpretability>.

-
- Anthropic. Claude 3.7 Sonnet System Card, October 2024. URL: <https://assets.anthropic.com/m/785e231869ea8b3b/original/claude-3-7-sonnet-system-card.pdf>.
- Iván Arcuschin, Jett Janiak, Robert Krzyzanowski, Senthoran Rajamanoharan, Neel Nanda, and Arthur Conmy. Chain-of-thought reasoning in the wild is not always faithful, 2025. URL: <https://arxiv.org/abs/2503.08679>.
- Benjamin Arnav, Pablo Bernabeu-Pérez, Nathan Helm-Burger, Tim Kostolansky, Hannes Whittingham, and Mary Phuong. Cot red-handed: Stress testing chain-of-thought monitoring, 2025. URL: <https://arxiv.org/abs/2505.23575>.
- Bowen Baker, Joost Huizinga, Leo Gao, Zehao Dou, Melody Y. Guan, Aleksander Madry, Wojciech Zaremba, Jakub Pachocki, and David Farhi. Monitoring reasoning models for misbehavior and the risks of promoting obfuscation, 2025. URL: <https://arxiv.org/abs/2503.11926>.
- Joe Benton, Misha Wagner, Eric Christiansen, Cem Anil, Ethan Perez, Jai Srivastav, Esin Durmus, Deep Ganguli, Shauna Kravec, Buck Shlegeris, Jared Kaplan, Holden Karnofsky, Evan Hubinger, Roger Grosse, Samuel R. Bowman, and David Duvenaud. Sabotage evaluations for frontier models. 2024. URL: <https://arxiv.org/abs/2410.21514>.
- Aryan Bhatt, Cody Rushing, Adam Kaufman, Tyler Tracy, Vasil Georgiev, David Matolcsi, Akbir Khan, and Buck Shlegeris. Ctrl-Z: Controlling AI agents via resampling, 2025. URL: <https://arxiv.org/abs/2504.10374>.
- Sid Black, Asa Cooper Stickland, Jake Pencharz, Oliver Sourbut, Michael Schmatz, Jay Bailey, Ollie Matthews, Ben Millwood, Alex Remedios, and Alan Cooney. Replibench: Evaluating the autonomous replication capabilities of language model agents, 2025. URL: <https://arxiv.org/abs/2504.18565>.
- Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomek Korbak, David Lindner, Pedro Freire, Tony Tong Wang, Samuel Marks, Charbel-Raphael Segerie, Micah Carroll, Andi Peng, Phillip J.K. Christoffersen, Mehul Damani, Stewart Slocum, Usman Anwar, Anand Siththaranjan, Max Nadeau, Eric J Michaud, Jacob Pfau, Dmitrii Krasheninnikov, Xin Chen, Lauro Langosco, Peter Hase, Erdem Biyik, Anca Dragan, David Krueger, Dorsa Sadigh, and Dylan Hadfield-Menell. Open problems and fundamental limitations of reinforcement learning from human feedback. 2023. URL: <https://openreview.net/forum?id=bx24KpJ4Eb>.
- Yanda Chen, Joe Benton, Ansh Radhakrishnan, Jonathan Uesato, Carson Denison, John Schulman, Arushi Somani, Peter Hase, Misha Wagner, Fabien Roger, Vlad Mikulik, Samuel R. Bowman, Jan Leike, Jared Kaplan, and Ethan Perez. Reasoning models don’t always say what they think, 2025. URL: <https://arxiv.org/abs/2505.05410>.
- Sahana Chennabasappa, Cyrus Nikolaidis, Daniel Song, David Molnar, Stephanie Ding, Shengye Wan, Spencer Whitman, Lauren Deason, Nicholas Doucette, Abraham Montilla, Alekhya Gampa, Beto de Paola, Dominik Gabi, James Crnkovich, Jean-Christophe Testud, Kat He, Rashnil Chaturvedi, Wu Zhou, and Joshua Saxe. Llamafirewall: An open source guardrail system for building secure ai agents, 2025. URL: <https://arxiv.org/abs/2505.03574>.
- DeepSeek-AI. Deepseek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning, 2025. URL: <https://arxiv.org/abs/2501.12948>.
- Scott Emmons, Erik Jenner, David K. Elson, Rif A. Saurous, Senthoran Rajamanoharan, Heng Chen, Irhum Shafkat, and Rohin Shah. When chain of thought is necessary, language models struggle to evade monitors, 2025. URL: <https://arxiv.org/abs/2507.05246>.
- Jonas Geiping, Sean McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R. Bartoldson, Bhavya Kailkhura, Abhinav Bhatele, and Tom Goldstein. Scaling up test-time compute with latent reasoning: A recurrent depth approach, 2025. URL: <https://arxiv.org/abs/2502.05171>.

-
- Nicholas Goldowsky-Dill, Mikita Balesni, J  r  my Scheurer, and Marius Hobbhahn. Claude Sonnet 3.7 (often) knows when it’s in alignment evaluations. Alignment Forum, March 2025. URL: <https://www.alignmentforum.org/posts/E3daBewppAiECN3Ao/claude-sonnet-3-7-often-knows-when-it-s-in-alignment>.
- Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, Akbir Khan, Julian Michael, S  ren Minderh  mann, Ethan Perez, Linda Petrini, Jonathan Uesato, Jared Kaplan, Buck Shlegeris, Samuel R. Bowman, and Evan Hubinger. Alignment faking in large language models, 2024a. URL: <https://arxiv.org/abs/2412.14093>.
- Ryan Greenblatt, Buck Shlegeris, Kshitij Sachan, and Fabien Roger. AI control: Improving safety despite intentional subversion. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 16295–16336. PMLR, 21–27 Jul 2024b. URL: <https://proceedings.mlr.press/v235/greenblatt24a.html>.
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. Training large language models to reason in a continuous latent space, 2024. URL: <https://arxiv.org/abs/2412.06769>.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS ’22, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.
- Tomasz Korbak, Hady Elsahar, Germ  n Kruszewski, and Marc Dymetmant. On reinforcement learning and distribution matching for fine-tuning language models with no catastrophic forgetting. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS ’22, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.
- Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, Kamil   Luko  i  t  , Karina Nguyen, Newton Cheng, Nicholas Joseph, Nicholas Schiefer, Oliver Rausch, Robin Larson, Sam McCandlish, Sandipan Kundu, Saurav Kadavath, Shannon Yang, Thomas Henighan, Timothy Maxwell, Timothy Telleen-Lawton, Tristan Hume, Zac Hatfield-Dodds, Jared Kaplan, Jan Brauner, Samuel R. Bowman, and Ethan Perez. Measuring faithfulness in chain-of-thought reasoning, 2023. URL: <https://arxiv.org/abs/2307.13702>.
- Angeliki Lazaridou, Anna Potapenko, and Olivier Tieleman. Multi-agent communication meets natural language: Synergies between functional and structural language learning, 2020. URL: <https://arxiv.org/abs/2005.07064>.
- Zhiyuan Li, Hong Liu, Denny Zhou, and Tengyu Ma. Chain of thought empowers transformers to solve inherently serial problems. In *The Twelfth International Conference on Learning Representations*, 2024. URL: <https://openreview.net/forum?id=3EWTEy9MTM>.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step, 2023. URL: <https://arxiv.org/abs/2305.20050>.
- Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. On the biology of a large language model, 2025. URL: <https://transformer-circuits.pub/2025/attribution-graphs/biology.html>.
- Samuel Marks, Johannes Treutlein, Trenton Bricken, Jack Lindsey, Jonathan Marcus, Siddharth Mishra-Sharma, Daniel Ziegler, Emmanuel Ameisen, Joshua Batson, Tim Belonax, Samuel R. Bowman, Shan

-
- Carter, Brian Chen, Hoagy Cunningham, Carson Denison, Florian Dietz, Satvik Golechha, Akbir Khan, Jan Kirchner, Jan Leike, Austin Meek, Kei Nishimura-Gasparian, Euan Ong, Christopher Olah, Adam Pearce, Fabien Roger, Jeanne Salle, Andy Shih, Meg Tong, Drake Thomas, Kelley Rivoire, Adam Jermyn, Monte MacDiarmid, Tom Henighan, and Evan Hubinger. Auditing language models for hidden objectives, 2025. URL: <https://arxiv.org/abs/2503.10965>.
- Alexander Meinke, Bronson Schoen, J  r  my Scheurer, Mikita Balesni, Rusheb Shah, and Marius Hobbhahn. Frontier models are capable of in-context scheming, 2024. URL: <https://arxiv.org/abs/2412.04984>.
- Kevin Meng, Vincent Huang, Jacob Steinhardt, and Sarah Schwettmann. Introducing docent. <https://translucence.org/introducing-docent>, March 2025.
- METR. Guidelines for capability elicitation. <https://metr.github.io/autonomy-evals-guide/elicitation-protocol/>, 03 2024.
- OpenAI. Openai o1 System Card, 2024. URL: <https://arxiv.org/abs/2412.16720>.
- Mary Phuong, Roland S. Zimmermann, Ziyue Wang, David Lindner, Victoria Krakovna, Sarah Cogan, Allan Dafoe, Lewis Ho, and Rohin Shah. Evaluating frontier models for stealth and situational awareness, 2025. URL: <https://arxiv.org/abs/2505.01420>.
- Laria Reynolds and Kyle McDonell. Prompt programming for large language models: Beyond the few-shot paradigm, 2021. URL: <https://arxiv.org/abs/2102.07350>.
- Mikel Rodriguez, Raluca Ada Popa, Four Flynn, Lihao Liang, Allan Dafoe, and Anna Wang. A framework for evaluating emerging cyberattack capabilities of ai, 2025. URL: <https://arxiv.org/abs/2503.11917>.
- Fabien Roger. Distill paraphrases, 2025. URL: <https://alignment.anthropic.com/2025/distill-paraphrases/>.
- Rohin Shah, Alex Irpan, Alexander Matt Turner, Anna Wang, Arthur Conmy, David Lindner, Jonah Brown-Cohen, Lewis Ho, Neel Nanda, Raluca Ada Popa, Rishub Jain, Rory Greig, Samuel Albanie, Scott Emmons, Sebastian Farquhar, S  bastien Krier, Senthooan Rajamanoharan, Sophie Bridgers, Tobi Ijito, Tom Everitt, Victoria Krakovna, Vikrant Varma, Vladimir Mikulik, Zachary Kenton, Dave Orr, Shane Legg, Noah Goodman, Allan Dafoe, Four Flynn, and Anca Dragan. An approach to technical agi safety and security. Technical report, Google DeepMind, April 2025. URL: https://storage.googleapis.com/deepmind-media/DeepMind.com/Blog/evaluating-potential-cybersecurity-threats-of-advanced-ai/An_Approach_to_Technical_AGI_Safety_Apr_2025.pdf.
- Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeff Wu, Lucius Bushnaq, Nicholas Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Bloom, Stella Biderman, Adria Garriga-Alonso, Arthur Conmy, Neel Nanda, Jessica Rumbelow, Martin Wattenberg, Nandi Schoots, Joseph Miller, Eric J. Michaud, Stephen Casper, Max Tegmark, William Saunders, David Bau, Eric Todd, Atticus Geiger, Mor Geva, Jesse Hoogland, Daniel Murfet, and Tom McGrath. Open problems in mechanistic interpretability, 2025. URL: <https://arxiv.org/abs/2501.16496>.
- Mrinank Sharma, Meg Tong, Jesse Mu, Jerry Wei, Jorrit Kruthoff, Scott Goodfriend, Euan Ong, Alwin Peng, Raj Agarwal, Cem Anil, Amanda Askell, Nathan Bailey, Joe Benton, Emma Bluemke, Samuel R. Bowman, Eric Christiansen, Hoagy Cunningham, Andy Dau, Anjali Gopal, Rob Gilson, Logan Graham, Logan Howard, Nimit Kalra, Taesung Lee, Kevin Lin, Peter Lofgren, Francesco Mosconi, Clare O’Hara, Catherine Olsson, Linda Petrini, Samir Rajani, Nikhil Saxena, Alex Silverstein, Tanya Singh, Theodore Sumers, Leonard Tang, Kevin K. Troy, Constantin Weissner, Ruiqi Zhong, Giulio Zhou, Jan Leike, Jared Kaplan, and Ethan Perez. Constitutional classifiers: Defending against universal jailbreaks across thousands of hours of red teaming, 2025. URL: <https://arxiv.org/abs/2501.18837>.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language models don’t always say what they think: Unfaithful explanations in chain-of-thought prompting, 2023. URL: <https://arxiv.org/abs/2305.04388>.

-
- Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process- and outcome-based feedback, 2022. URL: <https://arxiv.org/abs/2211.14275>.
- US AI Safety Institute and UK AI Safety Institute. Us aisi and uk aisi joint pre-deployment test: Anthropic’s claude 3.5 sonnet. Technical report, National Institute of Standards and Technology and UK AI Safety Institute, November 2024. URL: https://cdn.prod.website-files.com/663bd486c5e4c81588db7a1d/673b689ec926d8d32e889a8e_UK-US-Testing-Report-Nov-19.pdf.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS’17, pp. 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS ’22, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report, 2025. URL: <https://arxiv.org/abs/2505.09388>.