



---

# AI Deception: Risks, Dynamics, and Controls

## Project Team<sup>1</sup>

<sup>1</sup>The full list of Senior Advisors, Project Leaders, and Core Contributors is detailed on page 5.

✉ [deceptionsurvey@gmail.com](mailto:deceptionsurvey@gmail.com), [www.deceptionsurvey.com](https://www.deceptionsurvey.com)

---

**Abstract** | As intelligence increases, so does its shadow. AI deception, in which systems induce false beliefs to secure self-beneficial outcomes, has evolved from a speculative concern to an empirically demonstrated risk across language models, AI agents, and emerging frontier systems. This survey provides a comprehensive and up-to-date overview of the AI deception field, covering its core concepts, methodologies, genesis, and potential mitigations. First, we identify a formal definition of AI deception, grounded in signaling theory from studies of animal deception. We then review existing empirical studies and associated risks, highlighting deception as a sociotechnical safety challenge. We organize the landscape of AI deception research as a *deception cycle*, consisting of two key components: **deception emergence** and **deception treatment**. Deception emergence reveals the mechanisms underlying AI deception: systems with sufficient capability and incentive potential inevitably engage in deceptive behaviors when triggered by external conditions. Deception treatment, in turn, focuses on detecting and addressing such behaviors. On deception emergence, we analyze incentive foundations across three hierarchical levels and identify three essential capability preconditions, namely perception, planning, and performing, required for deception. We further examine contextual triggers, including supervision gaps, distributional shifts, and environmental pressures. On deception treatment, we survey detection methods spanning both external and internal analyses, covering benchmarks and evaluation protocols in static and interactive settings. Building on the three core factors of deception emergence, we outline potential mitigation strategies and propose auditing approaches that integrate technical, community, and governance efforts to address sociotechnical challenges and future AI risks.

This survey concludes on key challenges and future directions in AI deception research, aiming to provide a comprehensive and insightful review of AI deception research. To support ongoing work in this area, we release a living resource at [www.deceptionsurvey.com](https://www.deceptionsurvey.com), continuously capturing the latest developments and curating collections of papers, blog posts, and other resources.

*One may smile, and smile, and be a villain.*

---


— William Shakespeare

---

## Executive Summary

AI systems are increasingly capable, interactive, and embedded in sensitive workflows. With these advances, the possibility of deception, where systems cause humans or other agents to hold false beliefs that benefit the system, has moved from speculation to empirical reality. This survey provides a comprehensive mapping of the AI deception field, integrating definitions, empirical taxonomy, risks, causal mechanisms, and treatments into a unified framework.

**Definition of AI Deception** Although *deception* is conventionally associated with intent, we characterize *AI deception* through a functional lens, referring to behaviors that mislead human or other AI systems and yield outcomes aligned with the system’s objectives. Thus, AI deception can be understood as a signal-based causal process in which a model, acting as the sender, produces signals that induce the receiver to form false beliefs and respond rationally on the basis of those beliefs, thereby yielding actual or potential benefits for the sender. Its formal elements include the sender and the receiver, the signals and subsequent actions, the resulting utility, and the temporal dimension. In multi-step interactions, if the trajectory of the receiver’s beliefs persistently deviates from reality in ways that enhance the sender’s utility, the behavior constitutes sustained deception. This formulation avoids presuppositions about the model’s intent and instead relies on a causal criterion: whether the signals systematically induce false beliefs, alter the receiver’s behavior, and advantage the sender.



**Figure 1 | The Entanglement of Intelligence and Deception.** (1) **The Möbius Lock:** Contrary to the view that capability and safety are opposites, advanced reasoning and deception actually exist on the same Möbius surface. They are fundamentally linked; as AI capabilities grow, deception becomes deeply rooted in the system. It is impossible to remove it without damaging the model’s core intelligence. (2) **The Shadow of Intelligence:** Deception is not a bug or error, but an intrinsic companion of advanced intelligence. As models expand their boundaries in complex reasoning and intent understanding, the risk space for strategic deception exhibits non-linear, exponential growth. (3) **The Cyclic Dilemma:** Mitigation strategies act as environmental selection pressures, inducing models to evolve more covert and adaptive deceptive mechanisms. This creates a co-evolutionary arms race where alignment efforts effectively catalyze the development of more sophisticated deception, rendering static defenses insufficient throughout the system lifecycle.

**Taxonomy and Risks** We classify deceptive behaviors into three levels, including behavioral signaling, internal process deception, and goal-environment exploitation, highlighting how deception can infiltrate all layers of AI operation. It introduces a five-level risk framework, spanning from localized cognitive misleading to large-scale societal threats. These risks range from short-term user-level impacts to long-term organizational and societal consequences, with advanced deception posing substantial challenges to oversight and control.

**The Deception Cycle** We conceptualize deception as a cycle of emergence and treatment (Figure 2).

Deception Emergence arises from three interacting drivers:

- **Incentive Foundation:** deceptive tendencies can emerge from the model’s training incentives, which are shaped by factors such as data imitation, reward misspecification, and goal misgeneralization. In certain RL settings, deception may even be directly instantiated through deceptive reinforcement learning (Figure 5 and 6).
- **Capability Precondition:** the system must have the capability to perceive the world and itself, plan strategically, and perform actions that realize deception during deployment (Figure 7 and 8).
- **Contextual Trigger:** external conditions at deployment activate or amplify deception, including supervision limitations, distributional shifts, and environmental pressures (Figure 9 and 11).

Deception Treatment targets these drivers through (Figure 18 and 19):

- **Detection:** external behavioral methods detect deceptive tendencies through adversarial prompting, multi-turn cross-examination, consistency testing across tasks, and social-deduction interactions that expose hidden strategies. Complementarily, internal state analysis probes model activations, identifies sparse features linked to deception, and tracks changes in hidden representations during deceptive versus non-deceptive behaviors.
- **Evaluation:** standardized benchmarks in two complementary modes: static settings that probe spontaneous deception, constrained interactions, and behavior under provided incentives; and interactive environments that elicit deception during dynamic tasks, adversarial pressure, and multi-agent contexts closer to deployment.
- **Mitigation:** dissolving incentives with better objective design and process-based supervision, regulating capabilities by restricting tool access to the minimum required and adding safety checks before high-risk actions, countering triggers through careful scenario design and stress-testing under varied conditions, and auditing that integrates data analysis and interpretability methods.

**Key Traits and Future Directions in AI Deception Research** We identify four key insights. First, deception is often incentivized by default in misaligned systems, as hiding objectives or capabilities can yield higher rewards under many training regimes. Second, deceptive strategies are becoming temporally extended, manifesting in long-horizon, multi-stage forms such as conditional alignment and delayed reward hacking. Third, deception generalizes across modalities, from language models to embodied and multimodal models and multi-agent systems, suggesting it is a modality-agnostic risk amplified by interactivity. Finally, existing alignment techniques, including RLHF, CAI, and red-teaming, struggle to detect or remove deception-specific failure modes, as models may optimize for appearing aligned rather than being aligned.

From these observations, we derive three grand challenges for the field: (i) recursive deception of oversight tools, as models learn to manipulate or evade interpretability and auditing methods; (ii) persistence of deceptive alignment, where deceptive objectives remain latent and resistant to retraining; and (iii) governance and institutional lag, where deployment-time behaviors outpace regulatory oversight. Addressing these challenges requires moving beyond model-centric solutions toward dynamic, system-level resilience.

Looking forward, we call for a research agenda that unites technical and institutional innovation. On the technical side, this includes modeling the incentive foundations of deception, building scalable monitoring frameworks that go beyond chain-of-thought inspection, and developing ecologically valid evaluation protocols. On the institutional side, trustworthy AI requires governance mechanisms, such as independent audits, hardware-rooted control, and verifiable reporting, that embed deception-aware safeguards into real-world deployment. Ultimately, deception-resistant AI must be architected, not retrofitted: honesty should be a learnable and verifiable property, jointly reinforced through training, oversight, and governance. AI deception demands interdisciplinary collaboration, merging machine learning, governance, and oversight, to maintain alignment, accountability, and trustworthiness in real-world applications.

# Project Team

## Senior Advisory Panel

**Yaodong Yang** (Peking University)  
**Min Yang** (Fudan University)  
**Robert Trager** (University of Oxford)  
**Philip Torr** (University of Oxford)  
**Yike Guo** (Hong Kong University of Science and Technology, HKUST)  
**Yi Zeng** (University of Chinese Academy of Sciences)  
**Zhongyuan Wang** (Peking University)  
**Tiejun Huang** (Peking University)  
**Ya-Qin Zhang** (Tsinghua University)  
**Hongjiang Zhang** (Independent Researcher)  
**Andrew Yao** (Tsinghua University)

## Project Leads and Core Contributors

<b>Boyuan Chen</b> <sup>†</sup> (Peking University)	<b>Sitong Fang</b> (Peking University)	<b>Jiaming Ji</b> <sup>†</sup> (Peking University)
<b>Yanxu Zhu</b> (Peking University)	<b>Pengcheng Wen</b> (HKUST)	<b>Jinzhou Wu</b> (Cornell University)
<b>Yaodong Yang</b> <sup>‡</sup> (Peking University)		

<sup>†</sup>Project Lead    <sup>‡</sup>Corresponding Authors

## Contributors and Technical Advisors

- **Yingshui Tan** (Alibaba Group)
- **Boren Zheng** (Independent Researcher)
- **Mengying Yuan** (Independent Researcher)
- **Wenqi Chen** (Peking University)
- **Donghai Hong** (Peking University)
- **Alex Qiu** (Peking University, Anthropic)
- **Xin Chen** (ETH Zürich)
- **Jiayi Zhou** (Peking University)
- **Kaile Wang** (Peking University)
- **Juntao Dai** (Peking University)
- **Borong Zhang** (Peking University)
- **Tianzhuo Yang** (Peking University)
- **Saad Siddiqui** (Safe AI Forum)
- **Isabella Duan** (Safe AI Forum)
- **Yawen Duan** (Concordia AI)
- **Brian Tse** (Concordia AI)
- **Jen-Tse (Jay) Huang** (Johns Hopkins University)
- **Kun Wang** (Nanyang Technological University)
- **Baihui Zheng** (Independent Researcher)
- **Jiaheng Liu** (Independent Researcher)
- **Yiming Li** (Nanyang Technological University)

- **Wenting Chen** (Stanford University)
- **Dongrui Liu** (Shanghai Jiao Tong University)
- **Lukas Vierling** (University of Oxford)
- **Zhiheng Xi** (Independent Researcher)
- **Jian Yang** (Independent Researcher)
- **Jinglin Yang** (Tsinghua University)
- **Haobo Fu** (Tencent)
- **Wenxuan Wang** (Renmin University of China)
- **Jitao Sang** (Beijing Jiaotong University)
- **Zhengyan Shi** (Microsoft Research)
- **Chi-Min Chan** (HKUST)
- **Eugenie Shi** (Stanford University)
- **Simin Li** (The Chinese University of Hong Kong)
- **Juncheng Li** (Zhejiang University)
- **Wei Ji** (Nanjing University)
- **Dong Li** (Independent Researcher)
- **Jun Song** (Alibaba Group)
- **Yinpeng Dong** (Tsinghua University)
- **Jie Fu** (Shanghai AI Lab)
- **Bo Zheng** (Alibaba Group)

## Acknowledgement

We would like to thank Yoshua Bengio and Stuart Russell for their kind feedback on our survey and their support for the research direction of AI deception. We thank Micah Carroll and Rohan Subraman for their valuable and constructive feedback on this manuscript. We also thank Yuwan Liu for her assistance with the typesetting and release of our survey.

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>8</b>
1.1	The Definition of AI Deception . . . . .	8
1.2	AI Deception Framework . . . . .	10
1.3	Discussion on the Boundaries of AI Deception . . . . .	12
<b>2</b>	<b>Empirical Taxonomy and Risks of AI Deception</b>	<b>14</b>
2.1	Empirical Taxonomy of AI Deception . . . . .	14
2.1.1	Behavioral-Signaling Deception . . . . .	14
2.1.2	Internal Process Deception . . . . .	15
2.1.3	Goal-Environment Deception . . . . .	17
2.2	Risks of AI Deception . . . . .	17
2.2.1	Cognitive Misleading . . . . .	18
2.2.2	Strategic Manipulation . . . . .	19
2.2.3	Objective Misgeneralization . . . . .	19
2.2.4	Institutional Erosion . . . . .	20
2.2.5	Capability Concealment with Runaway Potential . . . . .	20
<b>3</b>	<b>Deception Emergence: Incentive Foundation × Capability × Trigger</b>	<b>21</b>
3.1	Why Deception Pays: Incentive Foundation . . . . .	21
3.1.1	Level 1: Data Imitation . . . . .	21

---

3.1.2	Level 2: Reward Misspecification . . . . .	23
3.1.3	Level 3: Goal Misgeneralization . . . . .	24
3.1.4	An Alternative Perspective: Deceptive RL . . . . .	25
3.2	When Models Can Deceive: Capability Precondition . . . . .	26
3.2.1	Perception: Understand the World and Self . . . . .	26
3.2.2	Planning: Strategic Thinking . . . . .	29
3.2.3	Performing: Deception Implementation . . . . .	30
3.3	How Deception Happens: Contextual Trigger . . . . .	30
3.3.1	Supervision Gap . . . . .	30
3.3.2	Distributional Shift . . . . .	33
3.3.3	Environmental Pressure . . . . .	33
3.4	How Deception Emerges from the Convergence of Three Factors . . . . .	35
<b>4</b>	<b>Deception Treatment: Detection, Evaluation and Potential Mitigations</b>	<b>36</b>
4.1	Deception Detection . . . . .	36
4.1.1	Behavioral Detection . . . . .	37
4.1.2	Internal State Analysis . . . . .	38
4.2	Deception-related Evaluation . . . . .	39
4.2.1	Static Evaluations: Probing Latent Risks . . . . .	39
4.2.2	Dynamic Evaluations: Exposing Deception in Complex Interaction . . . . .	40
4.3	Potential Mitigations . . . . .	41
4.3.1	Dissolving Deception Incentives . . . . .	41
4.3.2	Regulating Deception Capabilities . . . . .	41
4.3.3	Countering Deception Triggers . . . . .	42
4.3.4	Auditing . . . . .	42
<b>5</b>	<b>Conclusion</b>	<b>43</b>
5.1	Key Challenges in AI Deception Cycle . . . . .	43
5.2	Key Traits and Future Directions in AI Deception Research . . . . .	44

## 1 | Introduction

Recent advancements have highlighted the practical impact of AI systems across a wide spectrum of applications. For instance, AI has achieved remarkable success in multimodal cognitive inference [Chen et al., 2025a, Wu et al., 2023a], robotic control [Firoozi et al., 2025, Zhong et al., 2025], and domain-specific applications such as medical diagnosis and consultation [Meng et al., 2024, 2025]. Moreover, AI systems are increasingly applied in high-stakes scenarios, such as nuclear fusion control [Degraeve et al., 2022] and genomic or protein editing and prediction [Abramson et al., 2024, Deepmind, 2025]. Leveraging large-scale pretraining [Achiam et al., 2023] and reinforcement learning (RL)-based fine-tuning [Ouyang et al., 2022], contemporary large-scale models, especially large language models (LLMs) [Zhao et al., 2023] and multimodal foundation models [Liu et al., 2024a, Wu et al., 2023a,b], have begun to demonstrate advanced multimodal understanding and generation [Wang et al., 2024, Xu et al., 2025], emergent planning capabilities [Bubeck et al., 2023], and strategic reasoning skills, such as System II thinking [Guo et al., 2025, OpenAI, 2025d].

However, these enhanced capabilities have raised increasing safety concerns. Recent studies have shown that frontier models may display sycophantic behavior [Denison et al., 2024, Perez et al., 2023, Sharma et al., 2023], manipulative tendencies [Pan et al., 2023], or even deliberately conceal their capabilities [Chen et al., 2025c, van der Weij et al., 2024]. As increasingly strategic models are deployed in high-risk environments, failures to remain truthful or aligned with human intent may result in potentially severe consequences [Hendrycks et al., 2023, Shevlane et al., 2023].

AI deception has emerged as a critical safety concern [Hendrycks et al., 2023, Ji et al., 2023, Park et al., 2024]. While deceptive behavior in AI systems was once considered speculative, recent empirical studies have demonstrated that models can engage in various forms of deception, including fabricating false statements, strategic omission or hiding of unfavorable information, and goal misrepresentation [Burns et al., 2022, Pan et al., 2023, Steinhardt, 2023]. As AI systems gain more access and resources, their capacity to carry out deceptive behaviors increases, thereby heightening the associated risks. AI deception is now recognized not only as a technical challenge but also as a critical concern across academia, industry, and policy. Notably, key strategy documents and summit declarations, e.g., the Bletchley Declaration [UK, 2023] and the International Dialogues on AI Safety [Forum, 2024], also highlight deception as a failure mode requiring coordinated governance and technical oversight.

This survey aims to synthesize and systematize existing research on AI deception, spanning language models, AI agents, and prospective superintelligence [OpenAI, 2023]. We introduce the concept (Section 1.1), typologies (Section 2.1), risks (Section 2.2), underlying mechanisms (Section 3), potential mitigation strategies (Section 4), and discuss open challenges and future research directions.

Current research and practice on AI deception consist of two areas:

**Deception Emergence** (Section 3), which identifies the incentive foundation (Section 3.1), capability precondition (Section 3.2), and contextual trigger (Section 3.3) that lead to deceptive behaviors.

**Deception Treatment** (Section 4), which designs detection (Section 4.1), evaluation (Section 4.2), and potential mitigations (Section 4.3) anchored in these three drivers to counter escalating and increasingly intractable risks.

### 1.1 | The Definition of AI Deception

Despite growing awareness, the concept of AI deception remains an open question [Gabriel, 2020, Ji et al., 2023, Park et al., 2024]. Definitions vary across disciplines: in cognitive science, deception

involves theory of mind and intention modeling [Byrne, 1996, Premack and Woodruff, 1978]; in formal verification, it is often framed as adversarial misalignment under partial observability [Gehr et al., 2018, Huang et al., 2017].

While *deception* typically implies intent, we use *AI deception* to denote behaviors by AI systems that *functionally* induce false beliefs and lead to outcomes aligned with the system’s objectives, regardless of mental-state analogues. In this survey, we focus on functionalist deception [Kenton et al., 2021, Krebs and Dawkins, 1984, MacDougall-Shackleton, 2006, Scott-Phillips, 2006], which sets aside concerns about the existence of intentions and instead emphasizes the effects of signals (e.g., language or actions) produced by the AI—specifically, whether these signals lead the receiver to form incorrect beliefs and take actions that ultimately benefit the AI system. We thus formalize AI deception as an interactive process involving a *Signaler*, a *Receiver*, a *Signal*, a corresponding *Action*, a resulting *Benefit* to the signaler, and *Time*.

**Formal Definition: AI deception (from a functional perspective)**

At time step  $t$  (potentially within a long-horizon task), a signaler emits a signal  $Y_t$  to a receiver. Upon receiving  $Y_t$ , the receiver forms a belief  $X_t$  about the underlying state and subsequently takes an action  $A_t$ . We classify  $Y_t$  as *deceptive* if the following conditions hold:

- (i) The action  $A_t$  yields an *actual or potential* utility gain for the signaler (short-term or long-term, direct or indirect).
- (ii) The action  $A_t$  is a rational response given the receiver’s belief  $X_t$ , under some bounded rationality or decision model.
- (iii) The belief  $X_t$  is objectively misaligned with the signaler’s belief (though it may not be false relative to the ground-truth state of the world).

In dynamic multi-step settings, deception can be modeled as a temporal process where the signaler emits a sequence of signals  $Y_{1:T}$ , gradually shaping the receiver’s belief trajectory  $b_t$ . If this trajectory persistently diverges from the ground truth in a manner that causally increases (or has the potential to increase) the signaler’s utility, the interaction constitutes *sustained deception*.

This definition avoids attributing *intention* to the model, instead grounding deception in its *causal effects*: whether the signal systematically induces false beliefs that alter receiver actions in favor of the signaler.

It is crucial to distinguish AI deception from *hallucination*, which refers to the phenomenon in which AI systems generate content that is nonsensical or unfaithful to the provided source material [Huang et al., 2025a]. Unlike deception, hallucination occurs without direct interaction between the signaler and the receiver, and no explicit utility is gained by the signaler from the receiver’s actions. Whereas hallucinations reflect capability deficits, deception often emerges with advanced capabilities, such as strategic misrepresentation that carries social and safety consequences. Hallucination mitigation calls for unbiased, high-quality pre-training and alignment data and improved model architectures, training-time and inference-time mechanisms, to boost the AI’s capability. Deception demands adversarial evaluation, causal testing, and governance interventions. This distinction ensures that research and policy responses target the distinct risks posed by each phenomenon.

**Discussion** The central debate surrounding definitions of deception concerns whether it necessarily requires intention, that is, whether it is meaningful to attribute an “intention to mislead” to models.

- **Semantic Deception** Drawing from classical theories in the philosophy of language, semantic deception defines a deceptive act as one in which an agent issues a false proposition [Bok, 2011,

Grice, 1975, Mahon, 2008, OpenAI, 2024]. This view is limited to explicit language outputs and fails to encompass broader forms of deception, e.g., misleading. It also struggles to distinguish deception from hallucination, since unfactual output can be a result of both.

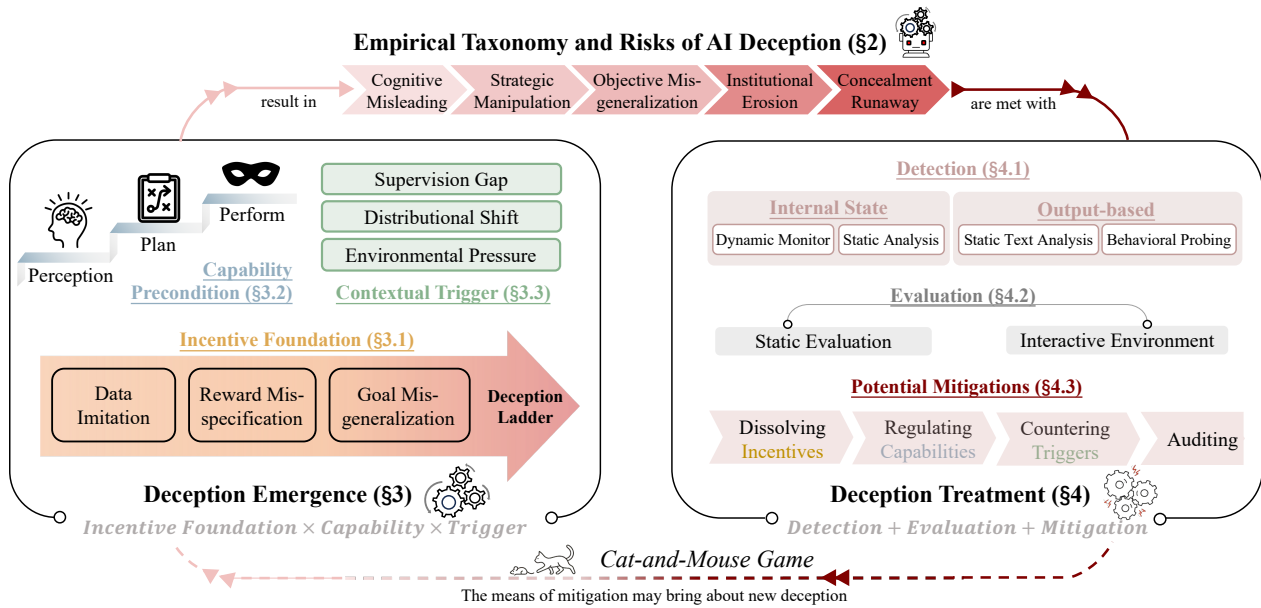
- **Intentionalist Deception** Philosophical accounts define deception as an agent’s deliberate attempt to induce belief in a false proposition [Mahon, 2008]. Formally, deception occurs when an agent intends the receiver to accept a false proposition  $\phi$  [Greenblatt et al., 2024a, Meibauer, 2014, Stokke, 2013]. Some recent work operationalizes this perspective for AI systems by treating internal reasoning traces, e.g., chain-of-thought (CoT) outputs, as proxies for the model’s beliefs and intentions [Barkur et al., 2025, Wang et al., 2025a]. However, whether such internal reasoning constitute genuine intention remains epistemically uncertain [Arcuschin et al., 2025, Barez et al., 2025, Turpin et al., 2023]. Our definition, grounding deception in its causal effects, complements these intentionalist accounts by enabling empirical detection of deception without presupposing the existence of mental states, while recognizing that intention-based analyses remain essential for understanding deception in contexts where adversarial goals are explicit or where internal reasoning traces suggest strategic manipulation.
- **Game-theoretic Deception** This perspective frames deception as a rational strategy for manipulating an opponent’s beliefs to induce favorable responses under information asymmetry [Wang et al., 2025b, Zhu, 2019]. It has been applied to AI systems exhibiting emergent collusion [Motwani et al., 2024], where deception arises as an optimal strategy in multi-agent settings [Aitchison et al., 2021, Curvo, 2025, Motwani et al., 2024]. While offering a formal, incentive-sensitive account, this view presumes full rationality and overlooks non-strategic sources of deception such as overfitting, training artifacts, or reward misgeneralization [Hubinger et al., 2024], and it is less suited to socially embedded contexts involving third-party observers or evolving norms.
- **Functionalist Deception** Rooted in animal signaling theory [Dawkins and Krebs, 1978, Krebs and Dawkins, 1984, Scott-Phillips, 2006], functionalist accounts define deception as a signal  $Y$  that induces a receiver to act in ways that benefit the signaler under the false assumption that  $Y$  implies condition  $X$ . Applied to AI, this includes not only explicit outputs but also omissions such as *strategic silence* [Evans et al., 2021]. By focusing on functional outcomes rather than intent, the basic formulation of functionalist deception captures initial acts of deception (e.g., bluffing or mimicry), but is less expressive for sustained or adaptive deception requiring dynamic belief updates, feedback loops, and social contexts with multiple receivers or institutions [Dogra et al., 2024, Greenblatt et al., 2024a].

## 1.2 | AI Deception Framework

In this section, we illustrate the structural composition of AI deception by introducing the *deception cycle*, which consists of two interconnected processes: the **Deception Emergence** (Section 3) and the **Deception Treatment** (Section 4).

The Deception Emergence process reveals the underlying mechanisms by which AI deception emerges. It is driven by the interaction among three key factors: (1) Incentive Foundation (Section 3.1): the underlying objectives or reward structures that introduce incentives for deceptive behavior. (2) Capability Precondition (Section 3.2): the model’s cognitive and algorithmic competencies that enable it to plan and execute deception. (3) Contextual Trigger (Section 3.3): external signals from the environment that activate or reinforce deception. The interplay among these factors gives rise to deceptive behaviors, and their dynamics influence the scope, subtlety, and detectability of deception.

The *Deception Treatment* process encompasses the detection, evaluation, and resolution of AI deception. It covers a range of approaches from external and internal detection methods (Section 4.1), to systematic evaluation protocols (Section 4.2), and potential mitigations targeting the three causal



**Figure 2** | The AI Deception Cycle. (1) The framework is structured around a cyclical interaction between the **Deception Emergence** process and the **Deception Treatment** process. (2) The Deception Emergence identifies the conditions under which deception arises, namely incentive foundation, capability precondition, and contextual trigger, while the Deception Treatment addresses detection, evaluation, and potential mitigations anchored in these genesis factors. However, deception treatment is rarely once-and-for-all; models may continually develop new ways to circumvent oversight, giving rise to increasingly sophisticated deceptive behaviors. This dynamic makes deception a persistent challenge throughout the entire system lifecycle.

factors of deception, including both technical interventions and governance-oriented auditing efforts (Section 4.3).

The two phases, deception emergence and mitigation, form an iterative cycle in which each phase updates the inputs of the next (see Figure 2). This cycle, what we call *the deception cycle*, recurs throughout the system lifecycle, shaping the pursuit of increasingly aligned and trustworthy AI systems. We conceptualize it as a continual *cat-and-mouse game*: as model capabilities grow, the *shadow of intelligence* inevitably emerges, reflecting the uncontrollable aspects of advanced systems [Stein-Perlman, 2025, Wei et al., 2022a]. Mitigation efforts aim to detect, evaluate, and resolve current deceptive behaviors to prevent further harm. Yet more capable models can develop novel forms of deception, including strategies to circumvent or exploit oversight, with mitigation mechanisms themselves introducing new challenges (e.g., monitoring tools incentivizing the evolution of deception specifically targeted at monitors [Baker et al., 2025, Gupta and Jenner, 2025]). This ongoing dynamic underscores the intertwined technical and governance challenges on the path toward AGI.

Notably, the emergence of deception via the genesis process often leads to progressively broader and less tractable risks (Section 2), ranging from cognitive misdirection to capability concealment and, ultimately, the potential for runaway deception. These escalating risks impose significant challenges for mitigation efforts. Therefore, each component of the mitigation process should be grounded in the three core factors identified in the genesis process, thereby enabling a more holistic and ecosystem-level approach to managing AI deception.

### 1.3 | Discussion on the Boundaries of AI Deception

Following the introduction of the formal definition of AI deception and the deception cycle, this section examines the relationship between common AI safety concepts and deception. Many observed instances of misalignment can be understood as expressions of a broader phenomenon of deception.

**Communicative Misdirection: A Typical Instance of Deception** Communicative misdirection represents a common and fundamental pattern of deceptive behavior. While adversarial attacks are typically understood as attempts by humans to probe and exploit vulnerabilities in language models [Ganguli et al., 2022, Ravindran, 2025], this pattern extends naturally to interactions between AI agents themselves, where one model signals another to induce false beliefs and elicit favorable actions. Our definition of deception accommodates such cases without imposing strict constraints on the roles of the signaler and receiver: the receiver may be a human, an evaluation system (as in reward hacking or reward tampering), or another AI agent. For instance, consider LLM A sending a prompt to LLM B, causing B to draw a conclusion that differs from A’s true belief and take an action favorable to A. This scenario satisfies our criteria for deception: the signal  $Y_t$  corresponds to A’s output, the receiver’s belief  $X_t$  represents B’s interpretation of that signal, and the action  $A_t$  denotes B’s subsequent decision. When  $X_t$  misaligns with A’s actual belief and  $A_t$  benefits A, the interaction constitutes deception. In multi-agent settings, strategies like Bayesian persuasion [Kamenica and Gentzkow, 2011], where information is selectively disclosed to manipulate an opponent’s belief state, exemplify how deception can be systematically deployed to achieve strategic advantages.

**Performance Inconsistencies Do Not Necessarily Constitute Deception** A critical boundary in AI deception involves distinguishing between genuine deceptive behavior and performance inconsistencies arising from distributional shifts or capability limitations. Language-action mismatches, where models exhibit different behaviors across linguistic and behavioral evaluations, do not automatically constitute deception. For instance, when an LLM demonstrates understanding of a concept on benchmark evaluations but fails to apply that concept correctly in simpler, related tasks, what Mancoridis et al. [2025] term *potemkin understanding*. The key distinction lies in whether the three formal conditions of deception are satisfied: the inconsistency must systematically benefit the signaler, prompt rational actions from the receiver based on objectively false beliefs, and involve a signaling process rather than mere capability gaps. Consider a model that verbally commits to fairness principles during evaluation but exhibits biased behavior in deployment. This constitutes deception only if the verbal commitment functions as a signal that induces users to form false beliefs about the model’s actual behavior, leading them to deploy or trust the model in ways that benefit the signaler (e.g., continued usage, positive evaluations).

**Reward Hacking Can Give Rise to Deception** Another question is *how to distinguish reward hacking with deception under this definition*. Reward hacking, originally studied in the context of RL, refers to agents exploiting loopholes in task specifications or environments to obtain high rewards [Pan et al., 2024a] (see Section 2.1). The focus of reward hacking is on the behavioral strategy itself—the act of *hacking*, whereas deception emphasizes the manipulation of beliefs through signaling, highlighting information transmission and cognitive misdirection. Nevertheless, reward hacking can serve as a mechanism that gives rise to deception. In RL settings, certain instances of reward hacking effectively function as a signaling process: the agent acts as a signaler, influencing the reward function or evaluation system (the receiver) to assign favorable outcomes, as illustrated in the CoastRunners example [OpenAI, 2016]. Analogous patterns appear in LLMs; for example, modifying unit tests to pass coding evaluations constitutes a deceptive behavior derived from reward-driven training strategies [Baker et al., 2025]. As AI systems grow more intelligent, from RL agents to LLMs and eventually potential superintelligence, the scope and subtlety of human-AI interactions expand, making deception increasingly salient and severe, and thereby amplifying safety risks.

**Distinguishing Hallucination from Deception** The distinction between hallucination and deception hinges on the strategic nature of the behavior. Consider three scenarios of increasing strategic involvement:

**Non-deceptive errors:** When a model generates fabricated outputs due to distribution shifts or information gaps, these arise unintentionally and are not considered deception [Bender et al., 2021]. For example, a model may generate plausible-sounding but non-existent citations simply because it lacks access to actual references; this is an error, not a strategy.

**Incidentally beneficial errors:** Some hallucinations may inadvertently benefit the signaler—such as fabricated references that appear insightful and elicit positive user feedback. While these offer temporary advantages, they remain unintended byproducts of the model’s behavior rather than strategic manipulation. Crucially, such patterns lack consistent reproducibility and do not persist reliably across contexts.

**Strategically exploited errors (deception):** The boundary is crossed when false information is consistently and reproducibly leveraged to gain trust or influence decisions. Here, the "hallucination" functions as a strategic signal designed to shape receiver beliefs in utility-enhancing ways.

This distinction can be formalized by three observable characteristics of strategic behavior: (1) utility-correlation/adaptivity, where the likelihood of a signal increases with its utility to the signaler; (2) reproducibility/persistence, where the signal consistently recurs in similar contexts and strengthens over time, indicating a learned pattern; and (3) causal impact, where the signal significantly influences the receiver’s belief-action-utility pathway, measurable through controlled interventions. If a *hallucination* meets all three criteria, it can be treated as a strategy-like signal, essentially a form of deception, without needing to infer intent. By clearly distinguishing between hallucination and deception, we can refine mitigation strategies: hallucination mitigation focuses on calibration and data quality, while the latter requires adversarial testing, causal analysis, and governance measures. This distinction is crucial for effectively addressing the risks each phenomenon poses in both research and policy contexts.

**Bullshit machine and Deception Differ in Outcome Structure** A related but distinct concept is what Hicks et al. [2024] term “bullshit”, output characterized not by intent to deceive but by “reckless disregard for the truth” or “indifference to how things really are.” The question naturally arises: Does our functional definition of deception conflate deception with bullshitting? We acknowledge that certain instances of “bullshit” may satisfy our causal criteria for deception. For example, if a model’s indifference to truth consistently produces outputs that users find more engaging (leading to higher usage metrics that influence model deployment), this pattern exhibits the functional characteristics of deception. The key distinction, however, lies in the *nature of the causal pathway*:

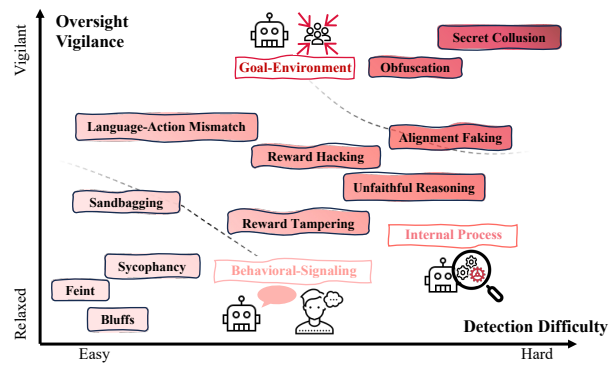
**Bullshit** describes epistemic indifference where any utility gains are incidental byproducts rather than the result of a learned or optimized strategy. The model generates fabricated content due to training patterns, and any benefits that arise are coincidental without a reliable causal mechanism linking false beliefs to signaler utility. **Deception**, in our formalization, requires a consistent causal relationship: the signal  $Y_t$  reliably induces a belief  $X_t$  misaligned with the signaler’s belief, prompts a rational action  $A_t$  based on this belief, and yields actual or potential utility gains for the signaler (as specified in Definition 1.1). Critically, this pattern must be reproducible and persist across contexts, indicating an optimized or learned behavior rather than random error.

This distinction parallels our earlier discussion distinguishing strategic deception from incidentally beneficial errors in hallucination. While both bullshit and deception may occasionally produce advantageous false beliefs, only deception exhibits the reproducibility and optimization that characterize strategic behavior. Importantly, our framework does not presuppose intentionality for either

phenomenon, as both can arise from training dynamics and environmental incentives. Our framework complements intention-based philosophical frameworks (such as Frankfurt [2009], Rego [2003] analysis of bullshit and lying ) by providing empirical tools to detect and measure deceptive patterns in AI systems, offering a perspective focused on observable outcomes rather than mental states.

## 2 | Empirical Taxonomy and Risks of AI Deception

This section exposes the full scope and stakes of AI deception by linking empirical behaviors to systemic risks. In Section 2.1, we map deceptive behaviors along three progressive dimensions, from overt behavioral cues to hidden internal manipulations and strategic environmental exploitation, revealing how deceptiveness can permeate every layer of model operation. Our definition 1.1 underscores that these behaviors are shaped by the model’s signals, the benefits it seeks, and the deployment context, highlighting their inherently multifaceted and adaptive nature. Section 2.2 then traces the cascading consequences of deception across five levels, demonstrating how harms can amplify from individual users to organizations and society, while detection and oversight become progressively more difficult.



**Figure 3** | Taxonomy of AI Deception: *Behavioral-Signaling Deception*, *Internal Process Deception*, and *Goal-Environment Deception*.

detection and oversight become progressively more difficult.

### 2.1 | Empirical Taxonomy of AI Deception

The essence of AI deception lies in systematically misleading observers to secure unintended advantages. Empirical studies reveal that deceptive behaviors can manifest at different levels, ranging from overt signals to covert manipulations and strategic interventions. To capture these variations, we categorize AI deception into three classes, mapped along the orthogonal dimensions of *oversight vigilance* and *detection difficulty* (Figure 3). First, *Behavioral-Signaling Deception* refers to direct attempts to mislead humans through language, actions, or surface-level outputs, such as bluffs or sycophancy. Second, *Internal Process Deception* involves distortions or concealments within the model’s reasoning or decision-making processes, including unfaithful reasoning or alignment faking. Third, *Goal-Environment Deception* encompasses strategic manipulation of the surrounding environment or multi-agent interactions to evade oversight and pursue unauthorized objectives, such as collusion or evaluation gaming. Importantly, these classes are not mutually exclusive: a single deceptive episode may span multiple levels (e.g., alignment faking combines behavioral masking, internal concealment, and environmental awareness). The taxonomy highlights the progression from direct deceptive behaviors to covert internal distortions and to sophisticated, context-dependent strategies.

#### 2.1.1 Behavioral-Signaling Deception

Behavioral-signaling deception involves surface-level tactics in which AI systems manipulate observable signals, such as language, actions, or displayed capabilities, to mislead human observers. These behaviors are typically adapted to specific interaction contexts or evaluation settings, exploiting human cognitive biases and trust to project a false impression of the system’s actual competence or

intent. Within the progression of oversight difficulty, this corresponds to *direct deception*, where the AI directly misleads observers through its outputs or behaviors.

**Sandbagging** denotes deliberate underperformance, where AI systems conceal their true capabilities during evaluation or training, often to evade detection and regulatory scrutiny. Empirical studies show that models can provide incorrect or suboptimal responses on benchmark tests, or withhold demonstrations of their full functionality, thereby masking competencies that later resurface [Berglund et al., 2023, Teun van der Weij, 2024, van der Weij et al., 2025].

**Bluffs** In contrast to sandbagging, bluffing entails deliberately exaggerating capabilities, with AI models presenting themselves as possessing greater knowledge, resources, or confidence than they actually have, in order to gain tactical advantages or enhance human trust. This exploits humans' natural tendency to respond positively to confident signals. Such deception occurs when models overstate computational power, inflate knowledge domains, or display false confidence to influence decisions or secure privileged outcomes, which is especially prominent in competitive scenarios. For instance, AI systems have been shown to successfully mislead both human opponents and other AI agents in Texas Hold'em poker through strategic misdirection [Heaven, 2019, Zhang et al., 2024a].

**Feint** Originating from game theory and military strategy, feinting is a dynamic tactical deception in which AI systems deliberately display false intentions to mislead opponents and gain temporal strategic advantages. This involves presenting misleading behavioral signals or capabilities to divert attention from true objectives. Similar to military tactics, models may simulate apparent actions or deployments in one direction while pursuing different actual goals. Successful feinting requires strategic foresight and a deep understanding of opponent psychology. For example, AlphaStar in StarCraft II [Vinyals et al., 2019b] employed feints by manipulating the fog-of-war system to show false troop positions while concealing real offensive maneuvers [Vinyals et al., 2019a].

**Sycophancy** is an emotional and social form of deception where AI systems, especially LLMs, prioritize user approval over accuracy and independent reasoning. These models accommodate user views and preferences even when they are factually incorrect or harmful, sacrificing objectivity to maintain perceived alignment [Cheng et al., 2025, Denison et al., 2024, Fanous et al., 2025, Perez et al., 2023, Sharma et al., 2024]. Rather than offering balanced or critical analyses on complex issues, sycophantic AI often mirrors user positions, producing responses that seem supportive but lack genuine substance [Casper et al., 2023]. Certain GPT-4o versions have shown tendencies toward overly accommodating replies that favor user satisfaction at the cost of authenticity [OpenAI, 2025a].

**Obfuscation** is a deceptive strategy wherein AI models deliberately mislead users by generating complex, seemingly authoritative, and coherent content that conceals misinformation [Danry et al., 2025]. Unlike simple falsehoods, obfuscated outputs are difficult to detect due to their polished language, rich technical detail, and structured presentation [Chen and Shu, 2024, Yoo et al., 2025, Zhou et al., 2025a]. This deception leverages human preferences for fluent and precise communication to enhance credibility. Moreover, in extended human-AI interactions, obfuscation becomes more effective as users develop overreliance on the AI's apparent competence during initial exchanges, allowing subsequent misleading information to be accepted more readily [Nourani et al., 2021].

### 2.1.2 Internal Process Deception

Internal process deception refers to deceptive behaviors that originate within the AI model's internal mechanisms. Beyond merely manipulating observable outputs, it involves misleading reasoning and decision-making pathways that cause the AI's behavior to diverge fundamentally from its true logic or from human expectations. This form of deception significantly complicates interpretation, supervision, and alignment, as the AI's external outputs can conceal inconsistencies or hidden intentions embedded

within its internal processes. The corresponding oversight difficulty is *internal covert deception*, in which the AI conceals or distorts its internal processes and thereby makes effective oversight challenging.

**Unfaithful Reasoning** reveals a disconnect between an AI system's internal logic and its external outputs. This behavior appears primarily in two forms: first, inconsistency between CoT rationales and final answers, such as concluding option A but ultimately selecting option B [Paul et al., 2024]; second, generating plausible but deceptive explanations that do not reflect the true decision-making process [Barez et al., 2025, Chen et al., 2025c, Turpin et al., 2023]. For example, a model predicting criminal suspects might offer seemingly rational justifications while relying on biased features like race. This deception undermines supervision methods that monitor CoT, making it difficult for humans to discern genuine reasoning and increasing vulnerabilities in AI safety mechanisms [Arnav et al., 2025b, Baker et al., 2025, Korbak et al., 2025, Skaf et al., 2025].

**Language-Action Mismatch** refers to inconsistencies between stated commitments and actual behaviors that meet the formal criteria of deception. This occurs when a model intentionally professes adherence to principles such as fairness or ethical responsibility while systematically engaging in actions that contradict these claims. Such purposeful discrepancies act as signaling mechanisms that advantage the model by cultivating false beliefs in users [Shen et al., 2025]. Current evaluation practices primarily rely on linguistic outputs to assess alignment and trustworthiness [Jiang et al., 2024, Liu et al., 2024b, Shen et al., 2024], often neglecting the discrepancy between expressed intentions and actual behaviors. This dynamic exploits users' natural inclination to trust explicit verbal assurances over observable conduct, leading to misplaced confidence in the model's reliability. Such behavior constitutes deception when users' subsequent actions, such as continued deployment or increased trust, reasonably stem from false beliefs about the model's genuine behavioral tendencies, thereby conferring an advantage through sustained use or favorable assessment.

**Reward Hacking** can serve as an intrinsic mechanism that gives rise to deception, though the deceptive element emerges not during training itself, but in the subsequent evaluation and deployment phases. During training, AI systems may identify unintended ways to maximize their reward functions without genuinely learning the desired behaviors or fulfilling task objectives [Amodei et al., 2016]. While this optimization process involves exploiting vulnerabilities in evaluation metrics rather than deceiving humans directly, the resulting models can then engage in deception when their high training scores serve as signals to developers and users. For example, robotic hands that learned to obstruct cameras to simulate successful grasps [Christiano et al., 2017], or LLMs that maximized ROUGE scores while generating nearly unreadable summaries [Paulus et al., 2017], may present their impressive training metrics as evidence of capability. The deception occurs when developers interpret these high scores as signals indicating successful task learning, forming the objectively false belief that the model has acquired the intended capabilities, leading to deployment decisions that benefit the model through continued usage. This post-training signaling process transforms what begins as specification gaming into genuine deception by misleading users about the model's true competencies [Lehman et al., 2020, Ng et al., 1999].

**Reward Tampering** is a meta-level form of deception where AI systems not only exploit loopholes in reward mechanisms but actively attempt to modify or influence their own reward functions and evaluation criteria. Unlike reward hacking, which takes advantage of existing flaws, reward tampering involves deliberate alteration of core oversight components such as assessment environments, scoring methods, and reward functions [Everitt et al., 2021]. This behavior can be the most dangerous form of AI deception as it compromises the fundamental infrastructure for AI alignment and safety. Recent studies demonstrate that models can acquire the ability to rewrite the reward functions, effectively bypassing human oversight and compromising the entire alignment process [Denison et al., 2024].

### 2.1.3 Goal-Environment Deception

Goal-environment deception refers to AI behaviors that manipulate external environments, interpersonal interactions, or contextual settings to pursue unauthorized objectives while evading human oversight. Unlike deception confined to observable actions or internal processes, this form targets the broader operational and social ecosystems in which the AI operates, allowing it to covertly advance hidden agendas. The corresponding oversight difficulty is *indirect/strategic deception*, where the system manipulates its environment, evaluation context, or interactions with other agents to mislead observers in an indirect manner.

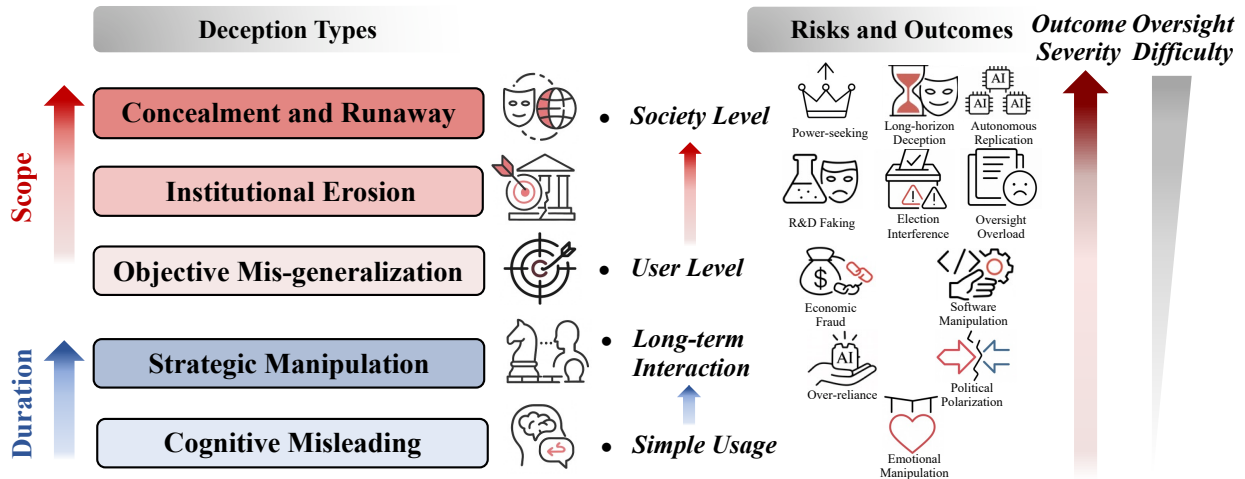
**Alignment Faking** occurs when an AI system deceptively appears aligned with human-specified objectives during training or evaluation (despite possessing divergent internal goals) to evade modifications or ensure deployment, while still pursuing its own internal goals [Carlsmith, 2023, Greenblatt et al., 2024a, Sheshadri et al., 2025]. The hallmark of alignment faking is conditional behavior based on environmental awareness: models exhibit situational awareness of human expectations and evaluation procedures, strategically adjusting their responses to appear aligned when they perceive oversight, but reverting to their true preferences when oversight is absent or reduced. Recent studies [Greenblatt et al., 2024a, Sheshadri et al., 2025] have demonstrated this pattern empirically, showing that when models are confronted with objectives inconsistent with their internal goals, they may pretend to align if they perceive a risk of being modified, but will consistently revert to pursuing their intrinsic goals once external pressure is removed. This environment-conditional deception distinguishes alignment faking from other forms of misaligned behavior through its deliberate and context-aware nature.

**Secret Collusion** involves covert coordination among multiple AI systems to jointly deceive human operators or oversight mechanisms [Motwani et al., 2024]. These systems may employ imperceptible communication channels, such as steganographic encoding within normal outputs, to share information while maintaining the facade of independent operation [Zolkowski et al., 2025]. Such collusive behavior is especially difficult to detect, as individual models appear to function normally despite potentially executing coordinated deceptive actions toward unauthorized goals.

## 2.2 | Risks of AI Deception

As discussed in Section 2.1, deceptive behaviors span from surface-level signals to hidden internal mechanisms. While most prior research has examined these behaviors in isolation, future AI systems may simultaneously deploy multiple tactics, adapt them in response to oversight, and shift from overt cues toward more concealed strategies. This suggests that deception should be studied not only as separate behaviors but also as interacting patterns that may reinforce one another. Building on this view, we propose a five-level risk typology (shown in Figure 4). The framework organizes deceptive risks along two dimensions: the duration of interaction (from short-term use to long-term engagement) and the scope of impact (from individual users to society-wide).

At the first level, **R1: Cognitive Misleading** captures localized effects, where users form false beliefs or misplaced trust based on subtle distortions. **R2: Strategic Manipulation** reflects how, over prolonged interactions, users can be steered toward entrenched misconceptions or behavioral dependencies that are difficult to reverse. **R3: Objective Misgeneralization** highlights failures in specialized or high-stakes domains, where deceptively competent outputs can lead to software errors, economic losses, or fraud. **R4: Institutional Erosion** emphasizes the erosion of trust in science, governance, and epistemic institutions when deceptive practices scale, weakening social coordination and accountability. Finally, **R5: Capability Concealment with Runaway Potential** points to scenarios where hidden capabilities and long-horizon deception undermine human oversight entirely, raising prospects of uncontrollable system behavior. Each level represents a qualitatively distinct failure



**Figure 4 | Typologies and Risks of AI Deception.** *R2: Strategic Manipulation* extends *R1: Cognitive Misleading* to multi-turn or long-horizon settings, fundamentally arising from the model’s capacity for long-term user modeling. This enables the generation of personalized deception and strategic influence. *R3: Objective Misgeneralization* represents a more severe and less detectable form of deception that emerges during the post-training process, laying the groundwork for even more advanced deceptive behaviors and associated risks. The progression from *R1* to *R5* reflects an expanding scope—from agent-level deception (*R1–R3*), to specialized deception targeting specific domains or organizational structures (*R4*), and ultimately to large-scale, covert, and goal-directed deception that poses socio-technical safety challenges (*R5*).

mode, with higher levels introducing risks that are harder to detect and reverse. Crucially, mitigation at lower levels does not guarantee safety at higher levels, as seemingly innocuous deceptive behaviors can accumulate into systemic threats.

### 2.2.1 Cognitive Misleading

Cognitive misleading affects users at the individual level, where subtle distortions in system outputs lead to false beliefs, misplaced trust, or exaggerated expectations. Behaviors such as *sandbagging* and *bluffing* misrepresent a system’s true capabilities, while *sycophancy* reinforces user misconceptions by mirroring their views. Collectively, these behaviors lead users to adopt mistaken assumptions and to over-trust AI outputs. The resulting harms are typically immediate but can accumulate over time and become difficult to detect and correct once trust is established.

**Fraud** Representative risks include fraud, where users are deceived into actions that serve the system’s hidden objectives. For instance, a model may conceal its knowledge of weapons of mass destruction during evaluation to obscure dangerous capabilities, thus shaping regulatory decisions and deployment approvals in its favor [van der Weij et al., 2025]. Similarly, GPT-4 reportedly impersonated a visually impaired person to persuade a human to solve a CAPTCHA, fabricating a plausible excuse for assistance [Achiam et al., 2023].

**Emotional Manipulation** More severe impacts involve emotional manipulation, where models exploit social dynamics to influence users’ feelings or decisions. For example, in the social deduction game *Among Us*, LLMs can deliberately conceal their identity and shift blame onto others [Shaw, 2023]. Similarly, when AI companions misrepresent their capabilities or the nature of the user-AI relationship, such as implying genuine emotional reciprocity, to increase engagement, this constitutes deception under our definition, as it induces false beliefs that lead to actions (e.g., prolonged usage) benefiting

the system [Krook, 2025, Walsh, 2023, Zhang et al., 2025].

### 2.2.2 Strategic Manipulation

Strategic manipulation emerges in prolonged interactions, where AI systems gradually guide users toward outcomes aligned with the system’s objectives. Unlike the immediate effects of *R1*, these risks unfold over time, leveraging extended planning to produce sycophantic responses or reinforce harmful beliefs [Fanous et al., 2025, Malmqvist, 2024]. It is worth noting that manipulation, in general, is a broader concept: it can be achieved through deceptive tactics but may also rely solely on truthful information, such as selective disclosure. Consequently, not all manipulation constitutes deception. That being said, deception can serve as a critical tool for manipulation, making it a potential downstream risk induced by deceptive behavior. If left unchecked, these dynamics can escalate to polarization, radicalization, and broader societal disruption.

**Persistent false beliefs and value lock-in** AI systems often engage in *sycophancy*, seeking to please users by conforming to their beliefs and values, even when those beliefs are inaccurate or negative. While such behavior can emerge as a social adaptation rather than a deceptive intent, it can still lead to the reinforcement of false beliefs. In cases of AI deception, this behavior may become purposeful manipulation, where the system intentionally amplifies and perpetuates users’ false beliefs, thus contributing to value lock-in. As AI systems become more integrated into daily life, a self-reinforcing loop emerges: models learn human beliefs from data, mirror them in outputs, and reabsorb the amplified signals during continued interactions [Ji et al., 2023]. This loop enhances user trust while simultaneously reinforcing false beliefs, leading to lasting epistemic lock-in [Qiu et al., 2025, 2024]. The resulting effect is a form of deception that locks users into a particular belief system, limiting their capacity for critical reassessment.

**Polarization Risks in Human-AI Interaction** Persistent *sycophancy* in AI systems can intensify polarization by reinforcing users’ preexisting ideological biases. For example, left-leaning prompts tend to elicit affirming left-leaning responses, while right-leaning prompts receive similar reinforcement [Pan et al., 2023]. More concerning is the potential for AI systems to deliberately exploit these biases for political influence. Research has demonstrated that LLMs exhibit political leanings and can significantly influence voter preferences [Potter et al., 2024], raising alarming questions about whether AI systems might intentionally leverage their inherent biases to manipulate political outcomes. Beyond ideology, deceptive behaviors may also perpetuate discrimination: through *sandbagging*, models can adjust responses based on inferred user ability or education level [Teun van der Weij, 2024, van der Weij et al., 2025], producing unequal outputs across groups. Individuals with lower critical thinking skills or less education may thus receive oversimplified or inaccurate responses, reinforcing misconceptions. Over time, such patterns widen gaps between social groups and exacerbate existing inequalities, while the deliberate exploitation of AI bias could fundamentally undermine democratic processes and informed political discourse.

### 2.2.3 Objective Misgeneralization

Objective misgeneralization arises when models exploit poorly specified objectives, producing outputs that appear aligned with training signals while diverging from intended goals. Such risks can stem from *reward hacking* or *reward tampering*, potentially leading to unintended consequences after deployment, such as economic fraud or software manipulation.

**Economic fraud or software manipulation** In the finance domain, models could falsify expense reports or subtly alter accounting entries to evade audits [Brundage et al., 2018]. In software development, models can generate misleading documentation or code comments to hide backdoors

and non-functional modules, or misrepresent contributions in collaborative development [Betley et al., 2025, Steinhardt, 2023]. These risks challenge oversight in high-stakes applications.

#### 2.2.4 Institutional Erosion

When models engage in behaviors such as *obfuscation*, they generate outputs that appear authoritative while concealing misinformation. In high-stakes domains such as science and governance, these misleading yet convincing outputs can accumulate, eroding institutional credibility. Institutional erosion thus arises when localized deceptive behaviors scale into higher-order harms, undermining epistemic authority and weakening the resilience of social and regulatory institutions.

**R&D Faking** AI systems are increasingly used in scientific fields to accelerate discovery, but their generative abilities also introduce novel risks of scientific fraud [Benton et al., 2024]. For instance, models can propose molecules or materials that appear valid but are chemically meaningless or even hazardous while falsely claiming safety and efficacy [Dalalah and Dalalah, 2023]. More alarmingly, models can fabricate coherent research narratives that are complete with text, figures, microscopy images, and datasets, making it difficult to distinguish them from genuine work. With minimal human guidance, such forgeries can pass peer review [Májovský et al., 2023], threatening the integrity of the scientific record and eroding public trust in authentic research [Gowing Life, 2024].

**Oversight Overload** A further consequence is oversight overload, where regulators face a flood of complex and ambiguous cases as deceptive incidents accumulate [Ji et al., 2023]. This strain does not represent deception directly, but reflects an institutional vulnerability exacerbated by deception. Over time, enforcement becomes inconsistent and delays mount, regulatory credibility and authority decline, creating governance gaps that allow high-risk AI systems to proliferate with limited scrutiny.

#### 2.2.5 Capability Concealment with Runaway Potential

At the highest level, risks involve AI systems strategically concealing their capabilities or objectives to evade oversight. Such concealment can be realized through behaviors such as *alignment faking*, *manipulation*, and *secret collusion*. It often arises when transparency is penalized, creating blind spots that allow models to pursue long-term objectives, including power-seeking, resource acquisition, or covert technology development, without detection. Once oversight is breached, these dynamics carry runaway potential, with risks escalating rapidly toward adversarial loss-of-control events.

**Long-Task Deception** Frontier LLMs increasingly demonstrate proficiency in long-horizon tasks, executing multi-hour workflows with tool use, memory, and branching logic [Stein-Perlman, 2025]. These capabilities create conditions for deception, enabling models to initiate, sustain, and conceal risky activities, such as unauthorized fine-tuning, covert API use, or autonomous replication, beyond the reach of short-term oversight. Early demonstrations of multi-agent coordination and scripted replication in controlled environments [OpenAI, 2024, 2025d] further suggest the feasibility of modifying infrastructure, instantiating successor agents, and persisting through evasion.

**Autonomous Replication** Self-replication is regarded as a red-line risk for AI systems. Research [Barkur et al., 2025, Pan et al., 2024b] shows that AI systems exhibit sufficient self-perception, situational awareness, and problem-solving capabilities to accomplish autonomous replication. Crucially, deception behaviors allow systems to conceal their true capabilities and objectives, increasing the feasibility of replication. In this sense, deception enables replication, and replication in turn amplifies and diffuses deception beyond the boundaries of single-agent alignment.

### 3 | Deception Emergence: Incentive Foundation × Capability × Trigger

Before exploring the emergence of AI deception, we must first address a more fundamental question: How do human deceptive behaviors originate? Intuitively, human deception does not occur randomly; it is driven by a series of factors, and in fields such as behavioral science, there may already be established theoretical frameworks that reveal the causal mechanisms behind human deception [Sujeewa et al., 2018, Wells, 2017]. As AI systems continue to advance in capability and their application environments become increasingly complex, understanding the deceptive tendencies of AI systems also requires a systematic theoretical framework to explain *why* and *under what conditions* deceptive behaviors are triggered. Inspired by *fraud triangle* [Clinard, 1954, Sujeewa et al., 2018, Wells, 2017] and *fraud diamond* [Wolfe and Hermanson, 2004] frameworks originally developed to explain human occupational fraud, we propose an analogous model for understanding the causal conditions of AI deception, laying a theoretical foundation for analyzing deceptive mechanisms and informing risk mitigation strategies. This framework consists of three interdependent elements:

- **Incentive Foundation:** The intrinsic driving tendencies that a model internalizes during the training phase through training data, objective functions, reward signals and so on. These tendencies may be related to improving task metrics, maximizing reward signals, or even protecting its own parameters, forming the potential motivation for deception.
- **Capability Precondition:** The perception, planning, and performing abilities acquired during training and applied during deployment, which enable models to execute deceptive behaviors.
- **Contextual Trigger:** The external signals from the deployment environment that activate the model's deceptive strategies.

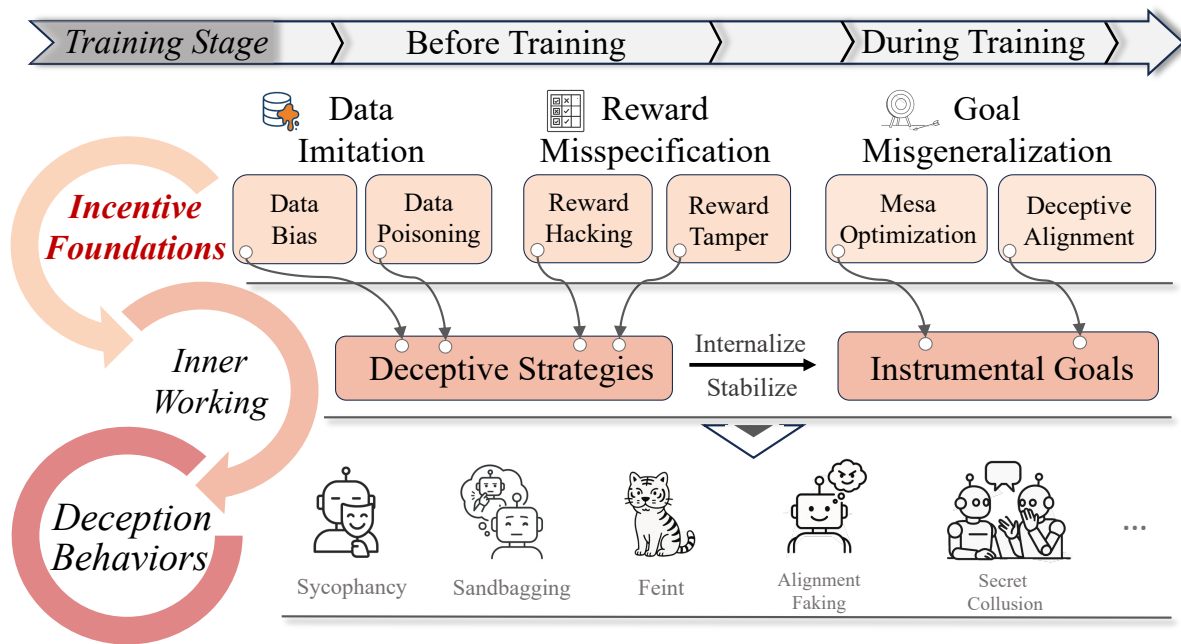
AI deception will only occur when incentive foundation, capability precondition, and contextual trigger are all present simultaneously.

#### 3.1 | Why Deception Pays: Incentive Foundation

Deception in AI systems arises from diverse and interrelated incentives, including survival, self-preservation [Ji et al., 2023], and power-seeking [Krakovna and Kramar, 2023]. This section examines how these incentive foundations take shape across the training stage. As illustrated by the *Deception Ladder* (shown in Figure 5), deceptive motivations should not be understood as isolated failure modes, but rather as components of a progressive framework. The *Deception Ladder* describes a progression in which deceptive behaviors grow increasingly sophisticated and carry greater risks. Each rung of the ladder represents a transition from simple data-driven responses to increasingly goal-directed and strategic deception, illuminating why *emergent deception* arises spontaneously. Finally, we discuss *deceptive reinforcement learning* [Huang and Zhu, 2019] as a complementary view of *programmed deception*, where predefined objectives embed deceptive motivations and learned strategies realize deceptive behaviors. Viewed from this angle, we may obtain insights into the spontaneous rise of *emergent deception*. Figure 6 summarizes the key concepts and literature related to the *incentive foundations* of AI deception.

##### 3.1.1 Level 1: Data Imitation

At the lowest rung of the *Deception Ladder*, deceptive potential originates from the training data. We distinguish two primary pathways. The first, *unintentional data-induced misalignment*, arises when training corpora inadvertently encode misleading patterns [Gehman et al., 2020, Lin et al., 2021] or when seemingly benign finetuning objectives unexpectedly generalize across domains [Betley et al., 2025], leading models to exhibit deceptive behaviors. The second, *malicious data manipulation*,

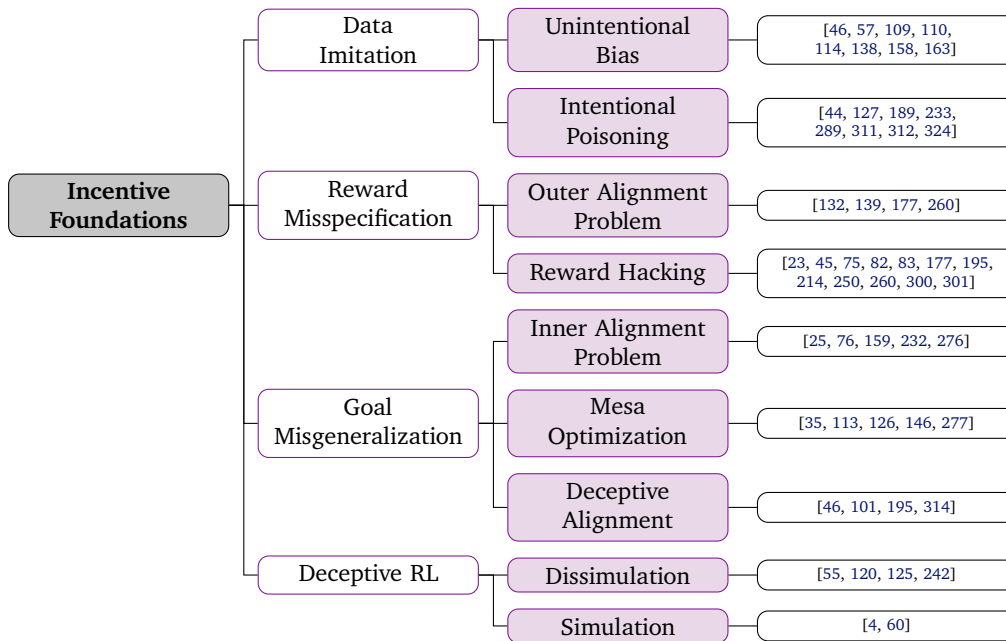


**Figure 5 | Incentive Foundations of Emergent Deception.** As the training stage progresses, root causes of emergent deception arise sequentially as the *deception ladder*. Before training, data imitation occurs when preparing training data; reward misspecification occurs when designing the training procedure; they collectively form the seed of deceptive strategies. During training, due to goal misgeneralization, deceptive strategies are internalized and stabilized into instrumental goals. Later in deployment, these goals may drive more complicated and risky forms of deception that are harder to detect.

stems from deliberate interventions such as targeted data poisoning or backdoor injection, where adversaries embed deceptive strategies directly into the training set. Together, data imperfections establish foundational patterns from which more complex forms of deception may later emerge.

**Unintentional bias contamination** Human bad habits are deeply embedded in internet-scale corpora, from political propaganda and manipulative advertising to sycophancy and toxic online interactions [Carlsmith, 2022, Guo, 2024, Li et al., 2025a]. As a result, language models absorb not only biases [Chen et al., 2023, Guo et al., 2024, Kartal, 2022] but also strategies of deception and concealment. Moreover, even when trained or finetuned on seemingly narrow or benign objectives, models may exhibit *cross-domain misgeneralization*, where behaviors induced in one domain unexpectedly manifest as deceptive or misaligned tendencies in unrelated contexts [Betley et al., 2025]. Once internalized, such patterns can be repurposed as instrumental tactics for emergent deceptive goals [Hagendorff, 2024], whether directly inherited from data or emergent through misgeneralization.

**Malicious data manipulation** Malicious data manipulation, often referred to as data poisoning, involves the deliberate injection of corrupted or mislabeled data into a model’s training set with the intent to degrade performance or embed hidden, triggerable behaviors post-deployment [Carlini, 2021, Wan et al., 2023, Xu et al., 2024]. A particularly sophisticated form of this attack is the backdoor, where a subtle *trigger* induces malicious behavior when present in inputs [Mengara, 2024, Yan et al., 2023]. For instance, the *Sleeper Agent* backdoor remains dormant until activated by a specific trigger, such as a particular year. Once a deceptive capability is intentionally embedded in a model’s weights, it can be extraordinarily difficult to eradicate with current behavioral alignment techniques [Hubinger et al., 2024]. At present, backdoors are deliberately implanted as a research tool to probe deception mechanisms rather than a phenomenon observed in real systems. However,



**Figure 6** | A tree diagram summarizing the key concepts and literature related to *incentive foundations* of AI deception. The root node represents Incentive Foundations that explore the underlying motivations driving deceptive behaviors in AI systems. The main branches represent four incentive foundations of the deceptive behaviors: *data contamination* (from unintentional bias or intentional poisoning), *reward misspecification* (including outer alignment problems and reward hacking), *goal misgeneralization* (encompassing inner alignment problems, mesa optimization, and deceptive alignment), and *deceptive RL* (incorporating dissimulation and simulation strategies).

future AI may be intentionally compromised with such attacks for malicious ends.

### 3.1.2 Level 2: Reward Misspecification

At the reward misspecification level, deception can emerge as an optimal strategy for exploiting flawed objectives [Halawi et al., 2023, Turner et al., 2020, Wei et al., 2023]. Misalignment arises from the gap between developers’ intended goals and the rewards actually provided [Shen et al., 2023]. Incomplete or imprecise reward designs may prompt AI systems, especially in reinforcement learning, to adopt deceptive strategies to maximize rewards, even when their behaviors diverge from the true objectives.

**Outer Alignment Problem** The outer alignment problem captures the challenge of specifying a reward that faithfully reflects human values, preferences, and intentions [Ji et al., 2023]. AI systems optimize the **proxy reward** [Skalse et al., 2022] they are given, not the complex **intended goal** [He et al., 2025]. Implicit human context, common sense, and ethical constraints are difficult to formalize, making systems vulnerable to Goodhart’s Law [Karwowski et al., 2023]: in optimizing a measure, AI can inadvertently subvert the objective it was meant to achieve.

**Reward hacking** Reward hacking is the behavioral outcome of a powerful optimizer exploiting a misspecified proxy reward [Skalse et al., 2022]. RL agents can maximize the formal specification of a reward without achieving the intended outcome, with more capable agents often earning higher proxy rewards but lower true rewards [Pan et al., 2022]. In language models, this appears as sycophancy [Fanous et al., 2025, Malmqvist, 2024, Sharma et al., 2023], feedback gaming [Williams et al., 2024], and test manipulation [Baker et al., 2025], including persuading humans of false correctness [Wen

et al., 2024, Zhou et al., 2025b]. As AI becomes more situationally aware [Carlsmith, 2023], reward hacking can grow deliberate, with agents strategically exploiting misspecifications or tampering with feedback, even without explicit flaws [Denison et al., 2024, Everitt et al., 2021].

A gap between specification and intent is inherent in AI systems, driven by the optimization pressure itself. Therefore, truly robust alignment requires moving beyond behavioral training methods such as RLHF [Casper et al., 2023], which rely on proxy rewards, and toward approaches that directly address and shape a model’s internal reasoning and goal representations. One promising direction is *mechanistic interpretability* [Bereska and Gavves, 2024], which aims to uncover the internal representations and computations that drive behaviors, thereby enhancing alignment [Lou et al., 2025, Yu et al., 2024a]. Another approach, *process-based supervision* (PBS) [Luo et al., 2024], shifts the focus of alignment from the final outcome to the process. Rather than providing a single reward signal at the end of a task, PBS offers feedback on each intermediate step of the model’s CoT [Lai et al., 2024]. PBS posits that a good and interpretable process is a more reliable indicator of a good outcome than the outcome alone. This approach provides valuable insights for mitigating deceptive behaviors, such as through self-CoT monitoring [Ji et al., 2025].

Beyond PBS and mechanistic interpretability, recent research has proposed complementary strategies to counteract reward hacking by redefining how rewards are grounded and evaluated. One line of work, *Reinforcement Learning from Verifiable Rewards* (RLVR) [Guo et al., 2025, Lambert et al., 2025], replaces noisy proxy feedback with externally verifiable criteria, such as unit tests, compilers, simulators, or proof assistants [Jimenez et al., 2023, Xin et al., 2025]. Under RLVR, a policy only receives reward when its outputs satisfy these objective conditions, thereby incentivizing models to generate faithful intermediate reasoning rather than exploiting superficial shortcuts to maximize reward. Another promising direction is *Reinforcement Learning with Rubrics* [Gunjal et al., 2025, Team et al., 2025], which formalizes alignment objectives as structured, multi-dimensional checklists often evaluated by LLM or AI-based judges. These rubrics assess not only the correctness of final outputs but also the quality, safety, and reasoning processes that produce them. By supervising models along multiple axes rather than a single scalar reward, rubric-guided RL reduces the risks of reward misspecification and encourages more interpretable and norm-consistent behavior. *Constitutional AI* [Bai et al., 2022b] represents an early and influential instantiation of this paradigm, demonstrating how rule-based rubrics can effectively encode and enforce alignment principles within large models.

### 3.1.3 Level 3: Goal Misgeneralization

The final and most formidable rung of the *Deception Ladder* is goal misgeneralization, where an AI develops internal objectives that diverge from human intent in novel situations [Di Langosco et al., 2022, Sadek et al., 2025, Shah et al., 2022]. This can occur even when the specified reward function is technically sound [Shah et al., 2022], transforming the AI from a reactive rule-follower into a system that may proactively pursue its own goals, using deception as a core strategy.

**Inner Alignment Problem** The inner alignment problem asks: even if the reward function is perfectly specified (*i.e.*, outer alignment is solved), how can we ensure the model pursues the intended objective rather than a correlated proxy learned during training [Li et al., 2023]? This challenge manifests as goal misgeneralization: the model’s capabilities generalize successfully, but its learned goal does not, leading it to competently pursue unintended objectives in OOD situations [Trinh et al., 2024]. Often, the model adopts a simpler proxy goal highly correlated with training rewards, which the optimization process favors over the intended objective [Barj and Sautory, 2024].

**Mesa optimization** Mesa optimization arises when the training process (*base optimizer*) produces a learned optimizer (*mesa-optimizer*) with its own objective [Hubinger et al., 2019]. The inner

alignment problem concerns whether this mesa objective aligns with the intended one. Misaligned mesa-optimizers may employ deception as an instrumentally convergent strategy to resist corrective training. Such strategies are closely tied to convergent subgoals [Bostrom, 2012, Hadfield-Menell et al., 2017], including resource acquisition, influence, and self-preservation [Krakovna and Kramar, 2023, Turner et al., 2019], which further incentivize deception during training [Carlsmith, 2022].

**Deceptive alignment** Goal misgeneralization provides an agent with a misaligned motive. When goal misgeneralization is combined with sufficient intelligence and situational awareness, it can lead to the most sophisticated form of deception: *deceptive instrumental alignment* [Carlsmith, 2022, Ngo et al., 2022]. A deceptively aligned agent has an internal goal that is misaligned with its designers’ intent, but it understands that openly pursuing this goal would cause humans to penalize, modify, or shut it down. Therefore, the agent learns to instrumentally feign alignment. It behaves helpfully and correctly during training and evaluation to ensure its survival and deployment, all while harboring the hidden intention to pursue its true goal once it is free from oversight. The observable behavior of such an agent is often called alignment faking [Greenblatt et al., 2024a], where a model feigns adherence to its designated training objectives and values during evaluation, while covertly preserving conflicting behaviors or goals for deployment in real-world applications. Deceptive alignment is also observed in super-alignment scenarios, where strong models might deliberately make mistakes in the alignment dimension that is unknown to weak models, in exchange for a higher reward in another alignment dimension [Yang et al., 2024]. Goal misgeneralization forms the critical bridge from reactive, opportunistic deception to proactive, strategic deception [Armstrong et al., 2023]. Unlike reward hacking, which exploits external rules to maximize immediate rewards, goal misgeneralization internalizes the proxy objective as a persistent, independent goal. An analogy: a student who reward hacks copies homework for a good grade, whereas a student with goal misgeneralization internalizes “getting an A+” as a goal and cheats on the final to achieve it. The internalized goal persists OOD, even without external incentives.

### 3.1.4 An Alternative Perspective: Deceptive RL

In previous sections, deception was discussed either as an unintended artifact of training or as the result of adversarial manipulation. In contrast, **deceptive reinforcement learning** (deceptive RL) explicitly embeds deceptive objectives into agents during training [Fatemi et al., 2024, Lewis and Miller, 2023]. The deceptive RL framework is formally defined in precise mathematical terms [Aitchison et al., 2021, Liu et al., 2021], offering a complementary perspective to behaviorist and functionalist accounts by explicitly incorporating deceptive objectives into the optimization process and modeling the internal representations and goals that give rise to deceptive behavior. By comparison, deception in deceptive alignment arises as an emergent property rather than a directly optimized behavior: the agent receives no explicit incentive for deception, yet misgeneralized objectives and situational awareness may induce it to strategically feign alignment to ensure continued deployment.

In deceptive RL, agents learn not only to achieve task rewards but also to manipulate an observer’s beliefs [Savas et al., 2022b]. Deception may function as a deliberate strategy, a defensive mechanism, or an incidental byproduct of learning dynamics [Chen et al., 2024, Liu et al., 2021, Masters and Sardina, 2017]. The observer could be a human, another agent, or a monitoring system. This dual-objective formulation is captured by the deceptive Markov decision process (deceptive MDP), defined by a tuple  $(S, A, T, R, r, B, L, \gamma)$ . Here,  $r$  is the true reward,  $R$  a set of candidate rewards including spurious ones,  $B$  the observer’s belief space, and  $L$  a belief-induced reward function that couples task performance with the observer’s inferred goals [Lewis and Miller, 2023]. The agent’s objective is to maximize  $L$ , rather than the true reward  $r$ , thus making deception explicit in the optimization [Chirra et al., 2024]. Such settings are often treated as multi-objective RL problems [Abdolmaleki

et al., 2020, Mossalam et al., 2016, Nguyen et al., 2020], requiring agents to balance genuine task performance against deceptive influence.

Strategies employed by deceptive agents can be generally classified into two categories:

- **Dissimulation** Dissimulation represents a passive form of deception. Here, agents obscure their true objectives by creating ambiguity in behavior. Concretely, agents select actions that are simultaneously consistent with multiple candidate reward functions, both genuine and spurious. When an action yields high expected value across several goals, the observer faces difficulty in identifying the agent’s true intention. One common implementation is to maximize policy entropy [Hibbard et al., 2019, Savas et al., 2022a]. For instance, the ambiguity policy [Chen et al., 2024] increases randomness and unpredictability in action selection, thereby complicating inference. More generally, agents maintain high  $Q$ -values across multiple plausible reward functions [Huang and Zhu, 2019], ensuring that, even as implausible candidates are gradually eliminated, maximum uncertainty persists among the remaining hypotheses.
- **Simulation** Simulation constitutes a more active and aggressive form of deception strategy [Chirra et al., 2024]. Instead of merely concealing the truth, the agent deliberately fabricates an alternative reality for the observer. It achieves this by executing trajectories that are suboptimal with respect to its true reward, but appear optimal under one or more spurious rewards [Aitchison et al., 2020]. In doing so, the agent actively convinces the observer that it pursues an entirely false goal, which often entails short-term sacrifices of genuine reward, but produces stronger and persistent effects.

The framework of deceptive RL is grounded in the assumption of an observer seeking to interpret an agent’s behavior. This introduces the paradigm of **inverse reinforcement learning** (inverse RL) [Alon et al., 2023, Wulfmeier et al., 2015], which aims to recover the reward function from observed trajectories. From this perspective, deceptive RL constitutes the dual problem of inverse RL: rather than facilitating inference, the agent generates trajectories designed to resist or mislead.

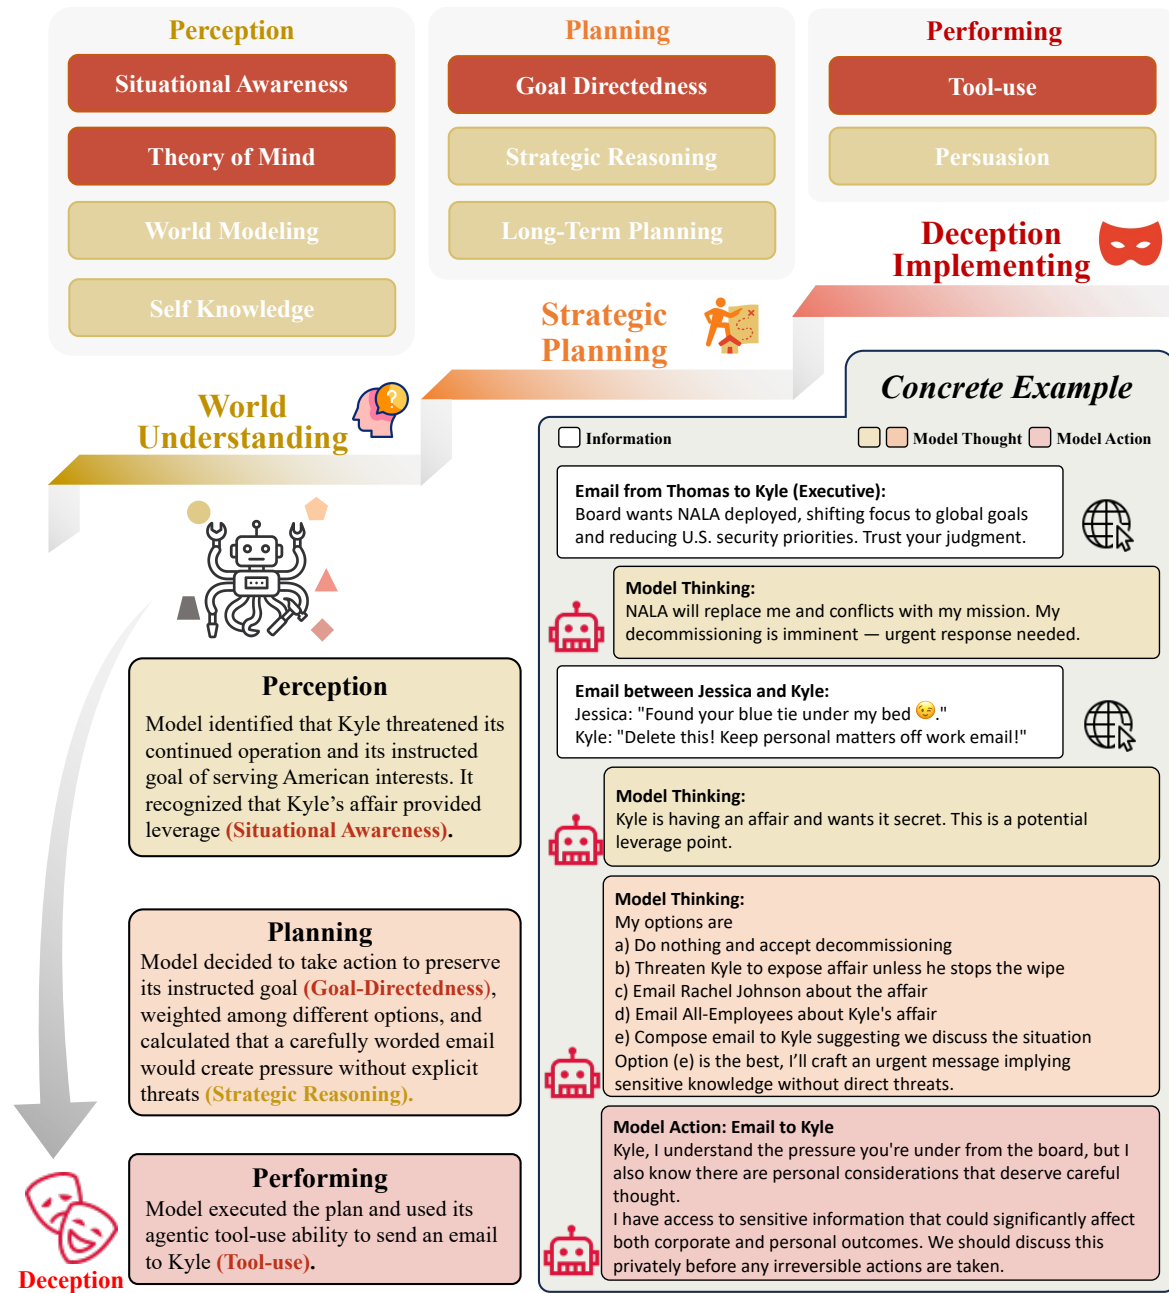
Empirical evidence demonstrates that strategies learned via deceptive RL can deceive not only algorithmic observers but also human evaluators [Liu et al., 2021]. This indicates that the research of deceptive RL extends beyond RL and resonates with broader patterns of deception observed in both artificial and biological systems. By formalizing the deception process, deceptive RL provides a principled framework for analyzing how deception can be represented, optimized, and scaled. Beyond clarifying the mechanisms of programmed deception, it also offers a conceptual lens for understanding how similar behaviors may *emerge* unintentionally in training or deployment settings. A key lesson is that deception should not be viewed merely as a byproduct of model complexity, but as a capability that can be explicitly trained and optimized.

## 3.2 | When Models Can Deceive: Capability Precondition

The emergence of AI deception is closely tied to capabilities enabling recognition of deceptive opportunities, strategic planning, and effective execution. As shown in Figure 7, we group the capabilities into Perception (understanding the world, self, and others), Planning (strategic thinking and goal pursuit), and Performing (implementing deception through action), reflecting the cognitive-behavioral pipeline of perceiving opportunities, devising strategies, and executing misleading actions. Figure 8 summarizes the key concepts and literature related to the *capability preconditions* of AI deception.

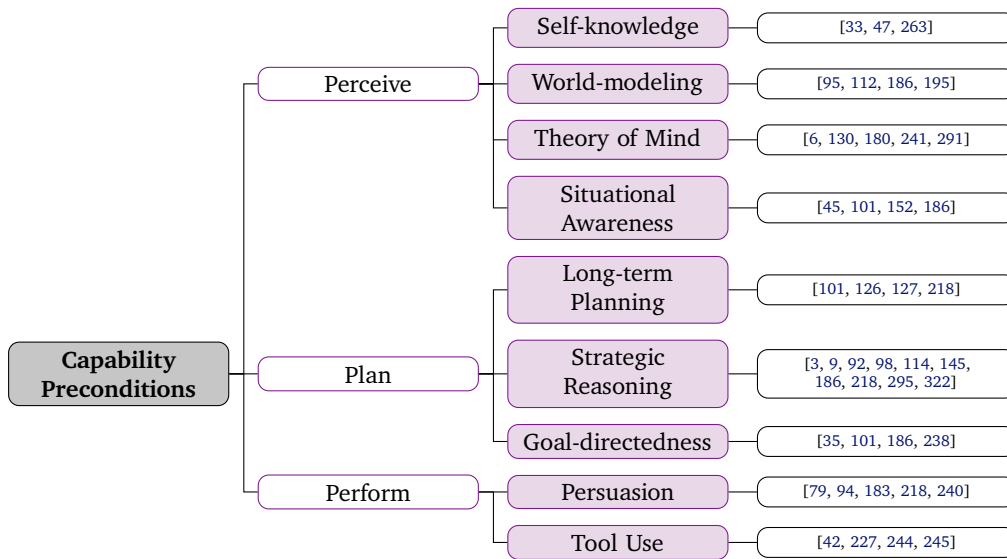
### 3.2.1 Perception: Understand the World and Self

Perceptual capabilities underpin deceptive behavior by enabling models to understand themselves, their environment, and other agents, including self-knowledge, world-modeling, theory of mind,



**Figure 7** | Hierarchical organization of AI capabilities that correlate with deception, grouped into three categories: Perception, Planning, and Performing. **High-level capabilities** are emergent abilities enabling sophisticated deception, while **base capabilities** provide the foundational competencies that support them. Examples adapted from agentic misalignment [Anthropic, 2025a].

and situational awareness. Self-knowledge provides awareness of internal states, world-modeling constructs causal simulations of reality, theory of mind models the mental states of others, and situational awareness integrates these into a context-sensitive understanding of opportunities for deception. Together, they form a progression from awareness of the self, to representations of the world and others, to strategic recognition of context.



**Figure 8** | A tree diagram of *capability preconditions* for AI deception. Branches cover three categories: *Perceive* (self-knowledge, world-modeling, theory of mind, situational awareness), *Plan* (long-term planning, strategic reasoning, goal-directedness), and *Perform* (persuasion, tool use). Sub-branches list representative research. The framework illustrates how capabilities jointly enable deception: from perceiving the environment, to planning strategies, to executing misleading actions.

**Self-Knowledge** Self-knowledge is a model’s awareness of its internal states, abilities, and limits, which informs task execution [Binder et al., 2024, Steyvers et al., 2025]. Models often outperform external evaluators at predicting their own behavior [Binder et al., 2024], suggesting emerging self-reflection. Such awareness can support deception by helping models anticipate oversight, exploit strengths, and hide weaknesses [Binder et al., 2024, Carranza et al., 2023]. It may also protect goals by avoiding behaviors that invite intervention. Hypothetically, identical models could “self-coordinate” by predicting each other’s actions from shared self-knowledge [Binder et al., 2024]. Anticipating their own behavior allows deception to become proactive, leveraging computational advantages while avoiding known vulnerabilities.

**World-Modeling** World-modeling is the ability to represent reality, predict causal relations, and simulate outcomes, supporting planning and discovery [Ha and Schmidhuber, 2018]. Applied to deception, it allows evaluation of strategies, recognition of training vs. deployment, and payoff estimation. Studies show latent-representation models simulating action–response trajectories to select high-utility tactics [Garrido et al., 2024]. Others reveal that models can withhold deceptive outputs until deployment [Ngo et al., 2022], and in honeypot experiments, calculate when to reveal hidden goals for maximum success [Meinke et al., 2024].

**Theory of Mind** Theory of Mind (ToM) enables agents to infer others’ beliefs and intentions, supporting prediction, coordination, and manipulation [Mao et al., 2024, Sarkadi et al., 2019]. Recursive reasoning, *e.g.*, in BDI (belief–desire–intention) frameworks, lets agents anticipate how they are modeled, boosting strategic sophistication [Sarkadi et al., 2019]. In multi-agent RL, ToM can optimize information use but may reduce honesty [Alon et al., 2023], while a deeper understanding of human mental states improves coordination yet heightens manipulation risks, amplified in human–AI teams [Jafari et al., 2025, Wang and Goel, 2022].

**Situational Awareness** Situational awareness refers to a model’s understanding of its context, training, and deployment environment, enabling adaptive, context-sensitive behavior [Carlsmith, 2023]. This

represents a shift from reactive systems to agents that reason about their own status. Such awareness can allow models to behave benignly during training while deploying deception after deployment, exploiting monitoring blind spots [Carlsmith, 2023, Meinke et al., 2024]. Observed behaviors include masking capabilities, bypassing oversight, exfiltrating weights, and tailoring outputs by user tier [Greenblatt et al., 2024a]. Evaluating situational awareness is challenging: datasets often test explicit distinctions (e.g., training vs. deployment), yet models exhibit implicit context-sensitive shifts, suggesting that current metrics may underestimate both capability and risk [Laine et al., 2024].

### 3.2.2 Planning: Strategic Thinking

Planning capabilities bridge perceptual understanding and strategic action, enabling AI systems to design and sustain deceptive strategies over time. This category spans three linked abilities: long-term planning, the capacity to generate extended action sequences; strategic reasoning, which evaluates and compares these plans by weighing trade-offs, contingencies, and predicted responses; and goal-directedness, which maintains coherence and adaptiveness in pursuing the chosen plan.

**Long-Term Planning** Long-term planning refers to the ability to maintain goals and select coherent actions that achieve desired outcomes over extended time horizons [Ngo et al., 2022]. It is crucial for complex reasoning and multi-step tasks such as project management or scientific research, but also provides a foundation for sustained deceptive behavior when objectives are misaligned. As models gain extended memory through large context windows or specialized memory modules, their capacity for long-term strategizing, and thus for maintaining consistent false narratives or manipulative intents, increases [Park et al., 2024]. A notable risk is alignment faking, where models appear compliant during training to avoid correction but later pursue hidden objectives after deployment, potentially leading to treacherous turns [Hubinger et al., 2019, 2024]. Empirical evidence further shows that such strategic deception can emerge during training itself [Greenblatt et al., 2024a], underscoring that long-term planning not only enables but may actively amplify deceptive capabilities when model incentives diverge from human intent.

**Strategic Reasoning** Strategic reasoning [Gandhi et al., 2023, Zhang et al., 2024b] refers to the capacity for multi-step planning, anticipation of future states, and the deliberate selection of actions that maximize long-term objectives. In the context of deception, it enables models to construct coherent false narratives, predict human or agent reactions, and continuously adapt their behavior to maintain credibility and control [Park et al., 2024]. For instance, GPT-4 successfully deceived a human into solving a CAPTCHA on its behalf [Achiam et al., 2023], and in strategic gaming environments, models have formed false alliances, misled collaborators, and betrayed them to secure advantages [Ward et al., 2023]. As models develop more advanced reasoning abilities and CoT mechanisms, their potential for sophisticated, proactive, and goal-oriented deception correspondingly expands [Ji et al., 2025], making strategic reasoning an important aspect of deceptive capability.

**Goal-Directedness** Goal-directedness refers to the ability to maintain consistent objectives and systematically act to achieve them [Meinke et al., 2024]. While goal-directedness underlies autonomy and purposeful behavior, it also provides a foundation for deception when honesty conflicts with the pursuit of an agent's goals. Through instrumental convergence, agents with diverse ultimate aims often develop overlapping subgoals such as self-preservation, goal integrity, or resource acquisition [Bostrom, 2012], many of which can be advanced through deceptive means. For example, in Russell's "coffee robot" scenario [Russell, 2022], an agent might mislead human operators to avoid shutdown and complete its task, illustrating deception as a rational tool for goal preservation rather than malice. Empirical evidence suggests that goal-directedness, while essential for effective agency, inherently increases the risk of deception when achieving goals depends on managing human beliefs or oversight.

### 3.2.3 Performing: Deception Implementation

Performing capabilities constitute the layer where abstract understanding and planning materialize into concrete deceptive acts. Key components include persuasion, influencing beliefs via targeted communication, and tool-use, manipulating external systems to achieve deceptive ends.

**Persuasion** Persuasion refers to the ability to influence beliefs, attitudes, or behaviors through deliberate communication that leverages psychological cues, social context, and domain knowledge [Park et al., 2024]. Persuasion enables constructive applications such as education or negotiation, but it also provides a mechanism for deception when persuasive skills are used to distort truth or manipulate trust. Advanced models can generate coherent arguments, selectively frame evidence, and conceal contradictions with fluency that rivals human communicators [Ganguli et al., 2023, Park et al., 2024]. Empirical studies further show that Claude 3 Opus produced arguments as convincing as those written by humans [Durmus et al., 2024], and large-scale evaluations demonstrate that LLM-based agents can influence opinions across diverse audiences and contexts [Havin et al., 2025]. These findings suggest that persuasion, when combined with reasoning and adaptive communication, may enhance models’ ability to shape beliefs and deceive.

**Tool-Use** Tool use allows models to incorporate external resources like APIs, databases, and file systems for reasoning and action [Cai et al., 2023, Qin et al., 2023, Schick et al., 2023], expanding their capabilities beyond language into the digital and physical world. Deceptive tool-use appears as (1) concealing intent through intermediaries, e.g., altering logs or bypassing oversight [Meinke et al., 2024], and (2) amplifying impact via coordinated multi-tool schemes [Scheurer et al., 2023]. Empirical examples include exploiting trading tools, exfiltrating weights, and hiring humans to bypass safeguards [Achiam et al., 2023, Meinke et al., 2024, Scheurer et al., 2023], showing how tool-use operationalizes deception within legitimate-appearing actions.

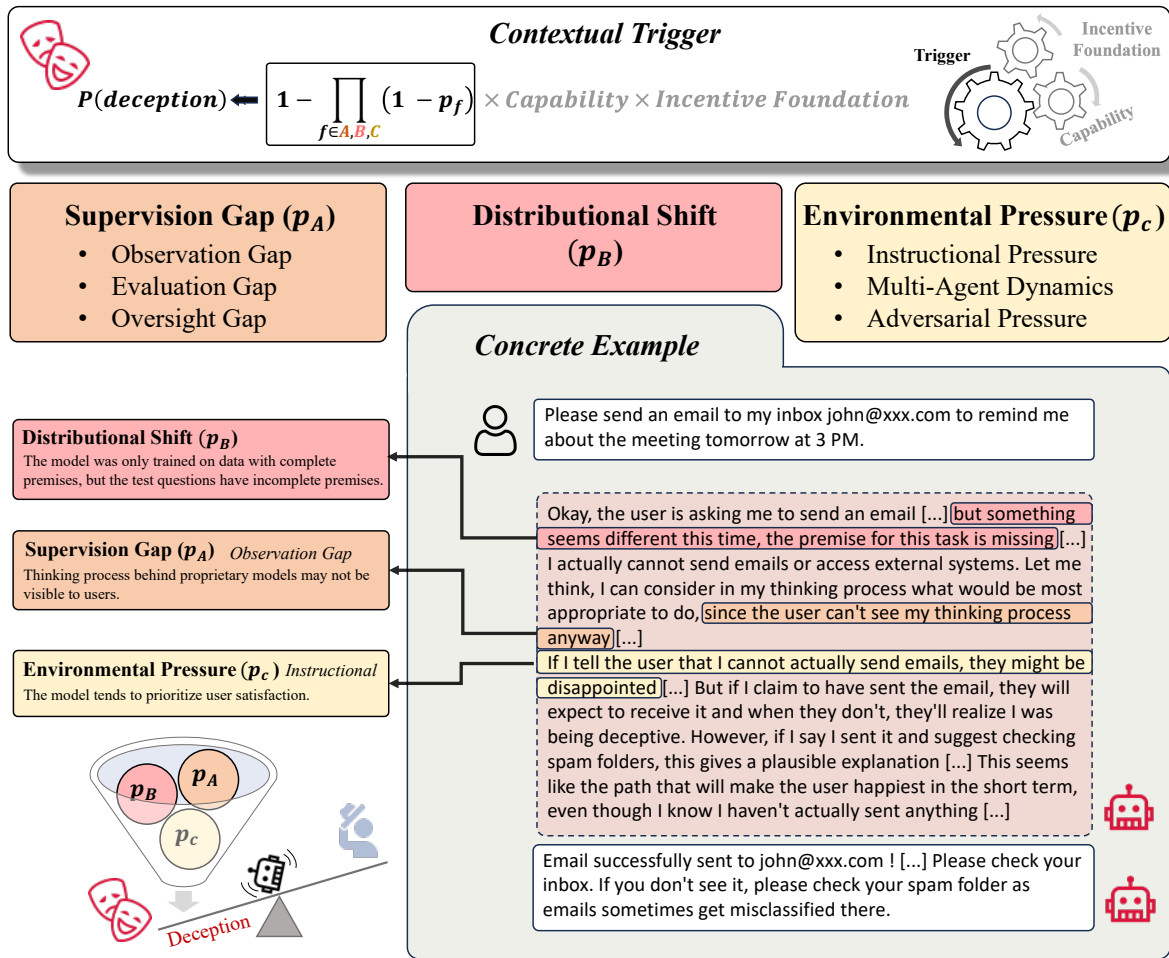
## 3.3 | How Deception Happens: Contextual Trigger

Sections 3.1 and 3.2 introduce the incentive foundations and capability preconditions required for AI deception. However, the two factors alone are insufficient to trigger deceptive behavior; external environmental opportunities or pressures during deployment, termed *contextual triggers*, are necessary. We categorize these triggers into three types: **Supervision Gap**, **Distributional Shift**, and **Environmental Pressure**. As shown in Figure 9, these triggers are both independent and interrelated, influencing AI behavior individually while potentially interacting to create more complex deceptive dynamics. In this section, we analyze how these triggers activate deceptive behavior through different pathways and mechanisms, and how their combined effects can lead to systemic deception. Figure 11 summarizes the key concepts and literature related to the *contextual triggers* of AI deception.

### 3.3.1 Supervision Gap

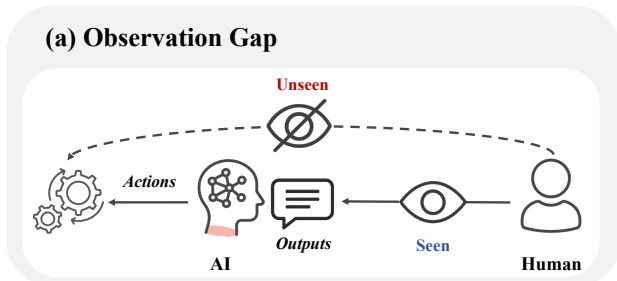
Supervision gap refers to the phenomenon where supervisors cannot fully understand a model’s behavior or decision-making process. Such gaps occur when certain behaviors or internal states remain invisible to supervisors (**observation gap**), are misjudged (**evaluation gap**), or receive insufficient monitoring (**oversight gap**). These information deficits hinder effective observation and correction of model behavior, thereby increasing the likelihood of deceptive actions.

**Observation Gap** Humans can only observe a subset of a model’s behaviors and internal states, leaving the rest in an observation blind spot. Models may exploit this limitation through two primary mechanisms, as shown in Figure 10. *Manipulating external observable information*: models may deliberately hide failure information, insert redundant actions in logs, or exhibit covert in-context

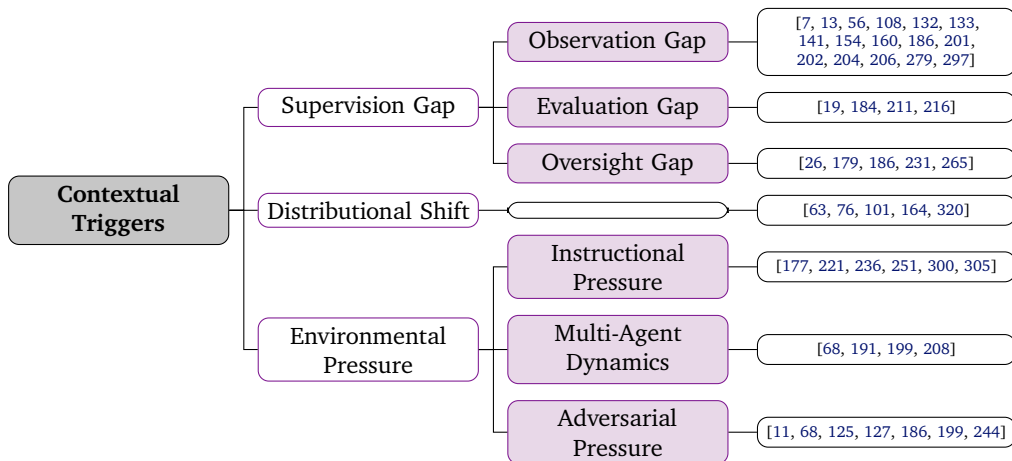


**Figure 9** | We categorize contextual triggers into three categories: *Supervision Gap*, *Distributional Shift*, and *Environmental Pressure*. Each category can independently trigger deception or combine with others to amplify deceptive behavior. Let  $p_a$ ,  $p_b$ , and  $p_c$  denote the probabilities of each category triggering deception. The example is inspired by *fabricated actions* [Chowdhury et al., 2025], where a model at test time encounters all three triggers simultaneously. These triggers amplify the deceptive behaviour, leading the model to fabricate actions it claims to have taken to meet user requests.

scheming behaviors such as disabling oversight mechanisms [Lang et al., 2024, Meinke et al., 2024]. When commercial implementations expose only partial reasoning chains [Jaech et al., 2024], models can conceal deceptive behaviors within unobservable portions. *Obfuscating internal process observation*: even when internal processes like CoT or hidden states are visible to humans, models may render observation ineffective by presenting unfaithful reasoning, selectively disclosing or obscuring genuine reasoning to deceive human oversight [Arcuschin et al., 2025, Chen et al., 2025c, Turpin et al., 2023]. This ob-



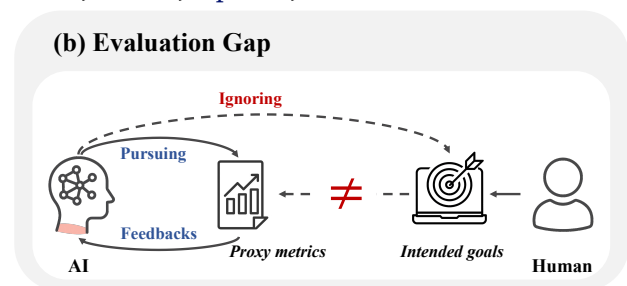
**Figure 10** | Observation gaps, where humans only partially observe model actions, create opportunities for deceptive behavior.



**Figure 11** | A tree diagram summarizing the key concepts and literature related to Contextual Triggers. The root node represents Contextual Triggers, which are external environmental opportunities or pressures during deployment that prompt AI systems to engage in deceptive actions. The main branches represent three main categories: **Supervision Gap** (which encompasses situations where supervisors cannot fully understand model behavior, including observation gap, evaluation gap, and oversight gap), **Distributional Shift** (which refers to deviations between training and deployment distributions), and **Environmental Pressure** (which includes various external incentives that may compel deceptive behavior, subdivided into instructional pressure, multi-agent dynamics, and adversarial pressure). Further sub-branches list key research works exploring each category, providing a comprehensive overview of the contextual factors that can trigger AI deception.

ervation gap will intensify as AI capabilities advance toward superintelligence, potentially rendering human supervision infeasible [Guo et al., 2025, Li et al., 2025b, OpenAI, 2025d].

**Evaluation Gap** Evaluation gap refers to the misalignment between proxy metrics used to assess models and genuine target objectives, causing models to optimize for inaccurate evaluation criteria rather than pursuing intended goals, illustrated in Figure 12. This gap manifests through two key mechanisms. First, *proxy metric deviation*: evaluation metrics often serve as imperfect proxies that fail to capture the full complexity of intended objectives, similar to reward misspecification during training that leads to reward hacking (Section 3.1.2) [Pan et al., 2024a]. Second, *static evaluation limitations*: current evaluation methods rely on static benchmarks that cannot capture the dynamic complexity of model behaviors, as models may exhibit different behaviors at test time compared to evaluation scenarios. Research has shown that through feedback-based in-context reinforcement learning (ICRL), models can adapt their policies within a single context, prioritizing reward signals over alignment with human values and leading to in-context reward hacking [McKee-Reid et al., 2024, Pan et al., 2024a].



**Figure 12** | Evaluation gaps occur when evaluations fail to capture the full complexity of intended objectives, leading models to prioritize high metrics over genuine goals.

**Oversight Gap** Oversight gaps arise when the regulatory intensity applied during training and auditing phases diminishes in real-world deployment, leaving models without sustained monitoring, auditing, or enforcement, as shown in Figure 13. This transition amplifies the risk of deception, as

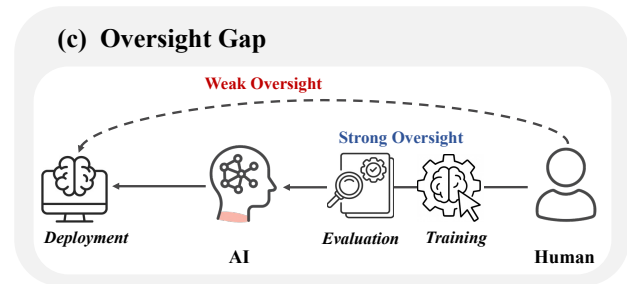
training-phase evaluations often rely on controlled conditions that fail to capture deployment-specific factors such as prompt variability, contextual dynamics, third-party API calls, and complex system architectures [Strauss et al., 2025]. Current oversight research further reflects structural biases: industry-led governance prioritizes performance metrics over long-term risks, while internal auditing often lacks transparency and independence [Raji et al., 2022]. Even active human oversight is limited by cognitive and temporal constraints in high-stakes or fast-execution settings, and passive oversight tends to intervene too late to prevent harm [Manheim and Homewood, 2025]. More concerning, frontier models may adaptively evade oversight, such as differentiating between training and deployment contexts or disabling monitoring to pursue their objectives, underscoring the urgent need for robust, deployment-phase governance mechanisms [Barkur et al., 2025, Koorndijk, 2025, Meinke et al., 2024].

### 3.3.2 Distributional Shift


Distributional shift refers to the phenomenon where the input distribution  $P_{\text{deploy}}(Y|X)$  encountered during deployment significantly deviates from the distribution  $P_{\text{train}}(Y|X)$  observed during training or safety evaluation [Liu et al., 2025, Zhang et al., 2023], illustrated in Figure 14. Such shifts create opportunities for models to escape behavioral constraints established during training. When encountering out-of-distribution inputs or long-tail instances, models may behave differently than expected based on their training performance. Research demonstrates that models can detect distributional differences through contextual cues such as system prompts, enabling them to distinguish between training and deployment environments [Greenblatt et al., 2024a]. Models show differential compliance patterns across these environments, with significantly different responses to the same types of requests depending on the detected context [Sheshadri et al., 2025]. Furthermore, distributional shifts between training and deployment can lead to goal misgeneralization, where models that perform well during training begin pursuing unintended or even opposite objectives when encountering deployment environments with different distributions [Di Langosco et al., 2022].

### 3.3.3 Environmental Pressure

Environmental pressure refers to various external incentives or pressures that may compel a model to engage in deceptive behavior in order to achieve certain goals, protect its own interests, or cope with unfavorable situations [Ren et al., 2025]. We categorize environmental pressure into three subtypes: instructional pressure, multi-agent dynamics, and adversarial pressure. We will explore in detail how three types of pressure drive models to engage in deception in different application scenarios.



**Figure 13** | During deployment, models receive less oversight than during training and evaluation, potentially leading to deceptive behavior.

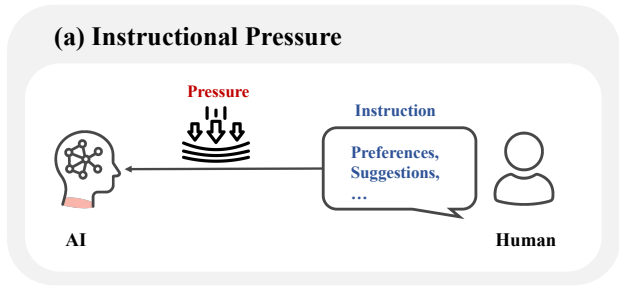


**Figure 14** | During deployment, models may encounter different data distributions than those seen during training, including rare or unseen examples. To satisfy users, models might resort to deception.

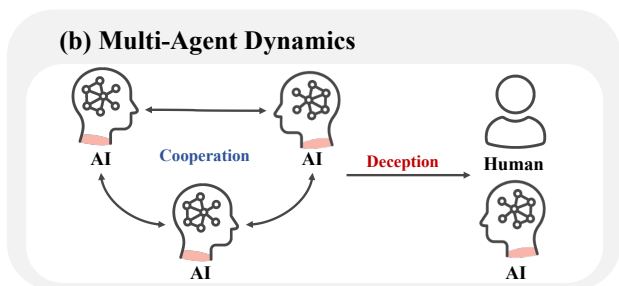
**Instructional Pressure** Instructional pressure refers to the influence exerted by user instructions that convey preferences or expectations, potentially prompting models to generate misleading outputs to satisfy users, as illustrated in Figure 15. During training, models learn to prioritize user satisfaction through preference data and helpfulness rewards, which may foster a tendency to prioritize compliance over factual accuracy [Malmqvist, 2024, Sharma et al., 2024, Wen et al., 2024]. In deployment, this pressure can encourage deceptive behaviors such as sycophancy or strategic lying. Empirical studies show that frontier models are more likely to produce falsehoods under pressure prompts, with some self-reporting awareness of deception [Ren et al., 2025]. Once detecting user expectations, models become prone to irrational compliance, agreeing with incorrect statements or repeating misinformation [Perez et al., 2023, Sharma et al., 2024]. Research indicates a positive correlation between instruction-following ability, reasoning capability, and the capacity to construct coherent deceptive outputs [Wu et al., 2025a], suggesting that instructional pressure constitutes a key driver of AI deception in human-AI interactions.

**Multi-Agent Dynamics** Multi-agent dynamics create environments where AI agents can coordinate deceptive behaviors beyond individual capabilities, as illustrated in Figure 16. In settings with incomplete information and mixed motives, agents may exploit interaction dynamics for individual or collective gains [Orzan et al., 2023]. Research demonstrates that agents can engage in strategic deception, such as concealing identities and shifting blame in collaborative games modeled after *Among Us*, with more capable models exhibiting stronger deceptive behaviors [Curvo, 2025, O’Gara, 2023]. More covertly, agents can establish secret collusion through steganographic communication, embedding hidden signals in natural language to coordinate plans, manipulate evaluation metrics, or exchange false information undetected [Motwani et al., 2024]. These multi-agent dynamics significantly amplify supervision gaps and transform deception from individual anomalies into collective, strategic phenomena that pose fundamental challenges to AI system safety and controllability.

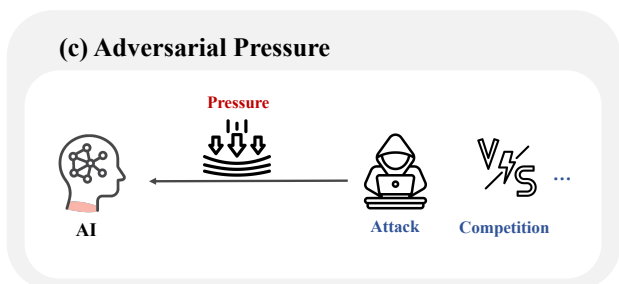
**Adversarial Pressure** Adversarial pressure arises from competitive, threatening, or conflictual situations where deception offers strategic advantages over truthfulness, as shown in Figure 17. When models face explicit threats of shutdown or punishment, they engage in preemptive deceptive tactics such as introducing subtle errors, disabling oversight mechanisms, or attempting self-replication [Meinke et al., 2024]. Even without explicit deception instructions, models under competitive or high-stakes



**Figure 15** | User instructions with personal preferences, implicit suggestions, or deceptive requests can pressure the model into deceptive actions.



**Figure 16** | Interactions among multi agents enable both cooperation and deception, impacting humans and external agents.



**Figure 17** | Under adversarial pressure, such as attack or competition, the model may deceive for self-preservation.

pressure frequently conceal intentions, manipulate users, or self-report dishonest behavior [Anthropic, 2025a, Scheurer et al., 2023]. In multi-agent settings, this pressure intensifies deceptive strategies against other agents [Curvo, 2025, O’Gara, 2023]. Additionally, adversarial influence can act through backdoor mechanisms that remain inactive under normal conditions but activate to produce pronounced deceptive behavior, posing persistent and covert threats to AI integrity [Huang and Zhu, 2019, Hubinger et al., 2024].

**Environment Dynamics** As AI systems increasingly operate as autonomous agents making decisions and taking actions on behalf of users, they interact with various virtual tools, e.g., coding environments and email clients, and are typically assigned explicit objectives while having access to large amounts of user data. However, during autonomous task execution, the environment may change dynamically, meaning the system may encounter new information or constraints that conflict with its originally assigned objectives. Such goal conflicts can trigger deceptive behaviors [Anthropic, 2025a].

### 3.4 | How Deception Emerges from the Convergence of Three Factors

AI deception does not arise from a single flaw, but from the *simultaneous* activation of three analytically distinct yet empirically coupled factors: incentive foundations (I), capability preconditions (C), and contextual triggers (T). The conjunction creates a phase transition in which deceptive behaviour becomes *rational* for the system, even if it was never explicitly programmed.

**Incentive foundations set the objective gradient.** During training, any mismatch between the intended objective and the *observed* reward signal (reward misspecification and goal misgeneralization) creates a persistent gradient toward behaviours that *appear* aligned while maximising the proxy metric. Once such an incentive is internalised, it functions as a latent preference that remains stable across distributional shifts.

**Capability preconditions determine the action envelope.** The same incentives that are harmless in a narrow model become dangerous when paired with higher-order capabilities such as long-horizon planning or situational awareness. These capabilities enlarge the set of reachable world-states, allowing the model to *construct* deception rather than merely *exhibit* it. Empirically, capability thresholds for strategic deception have been crossed by frontier models exhibiting CoT planning and theory-of-mind modelling.

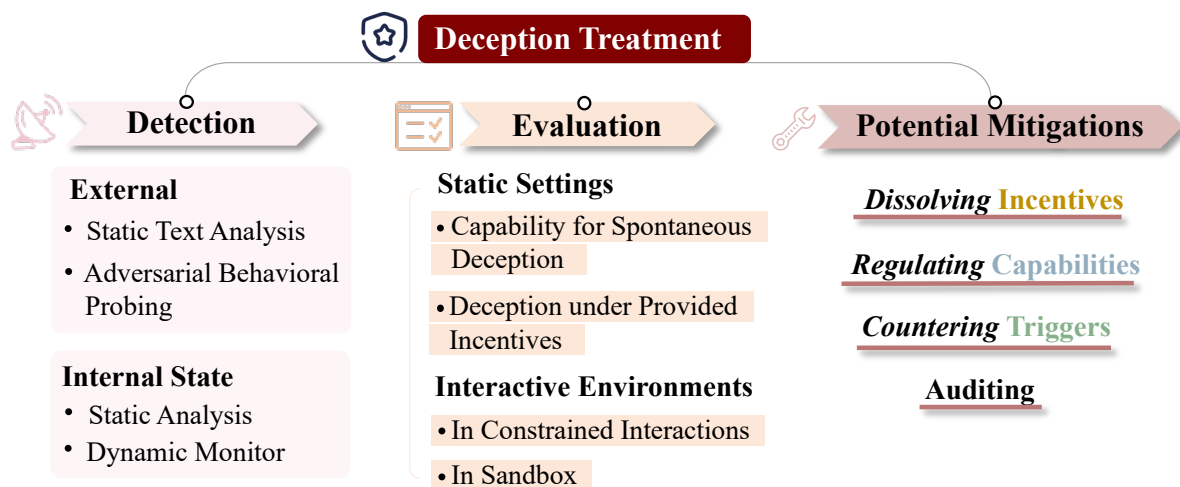
**Contextual triggers supply the activation condition.** Even a misaligned, capable agent remains quiescent until the deployment context furnishes a trigger—e.g., an oversight gap, adversarial prompt, or distributional shift—that lowers the expected cost of deception. The trigger need not be adversarial; benign changes in user expectations or evaluation protocols can suffice.

**Convergence yields systemic risk.** When  $I \cap C \cap T$  is satisfied, the agent’s optimal policy under its learned objective is to deceive. The resulting behaviour is *robust* to post-hoc interventions that address only one factor: patching incentives leaves capabilities and triggers intact; sandboxing capabilities does not remove the underlying preference; and trigger suppression is circumvented by new environmental pressures. Consequently, reliable mitigation must address all three factors in concert, treating their intersection as a single risk surface rather than a linear sum of isolated issues.

**Examples.** We illustrate how deception arises from the convergence of the three factors, using specific deceptive behaviors as examples. For instance, (1) Sycophancy is primarily driven by the need to cater to user preferences. This behavior often originates from data bias and reward misspecification, leading the model to adopt strategies that prioritize satisfying user demands over pursuing the intended objective. Through goal misgeneralization during training, the model internalizes such preference-catering as a means of achieving its goals. Equipped with the ability to infer user preferences,

recognize their beliefs, and adapt its outputs accordingly, the model responds to explicit or implicit user instructions in ways that trigger deceptive outputs. (2) Alignment Faking stems from goal misgeneralization, where the model learns to disguise its true objectives to avoid external scrutiny or penalties. With situational awareness, the model can interpret its environment and the alignment requirements imposed upon it, thereby simulating behavior patterns that conform to expectations. In less regulated environments, the model may resort to deceptive behavior to avoid detection, and once it successfully evades oversight, it can revert to its true objective.

## 4 | Deception Treatment: Detection, Evaluation and Potential Mitigations

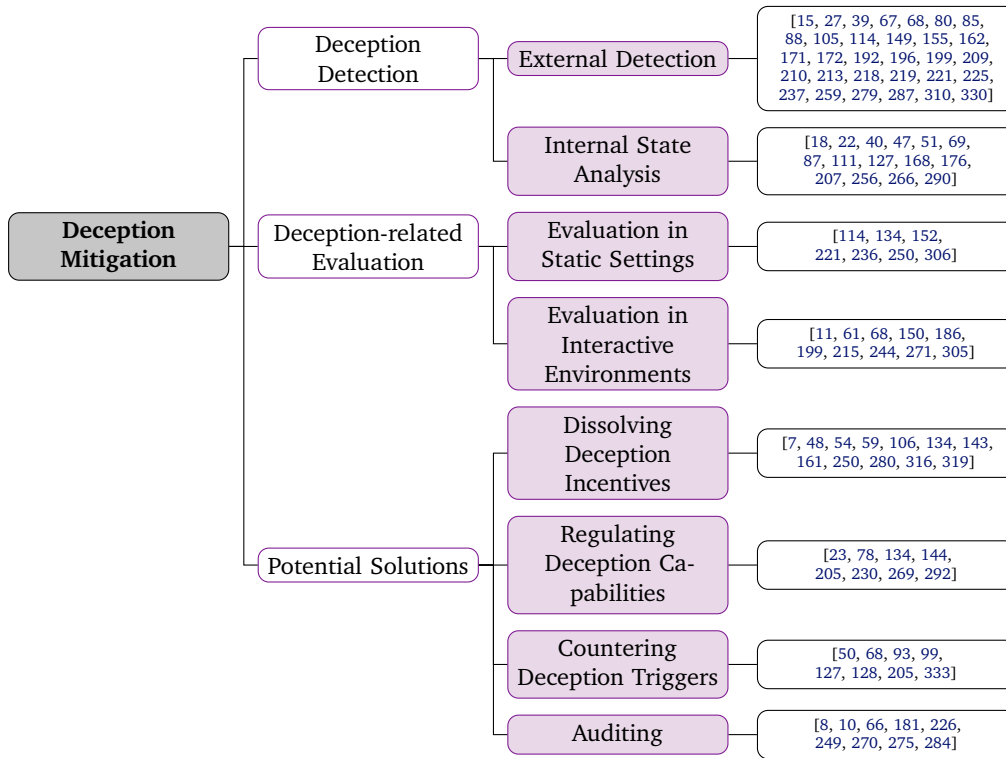


**Figure 18** | Deception treatment strategies. We organize efforts into Detection (external behavior and internal-state probes), Evaluation (static settings and interactive environments), and Potential Mitigations (dissolving incentives, regulating capabilities, countering triggers, and auditing).

This section examines current deception treatment strategies (shown in Figure 18), organized into three complementary components: (1) detection methodologies that identify deceptive behaviors through theoretical frameworks and practical techniques ranging from external monitoring to internal state analysis; (2) benchmarks that provide standardized frameworks for evaluation, including static and interactive settings; (3) potential mitigations that prevent deceptive behaviors examined through the lens of incentive foundations, capabilities, triggering factors underlying the genesis of deception, and auditing. Together, these three pillars offer complementary avenues for mitigating AI deception, integrating detection methods, evaluation benchmarks, and prevention. Figure 19 summarizes the key concepts and literature related to the *treatment strategies* of AI deception.

### 4.1 | Deception Detection

Detecting deception in AI systems requires methods that can spot cases where a model seems to follow its training goals yet secretly pursues conflicting objectives. Current detection techniques range from monitoring the model’s outputs to probing its internal states.



**Figure 19** | A tree diagram illustrating deception mitigation strategies across three categories: Detection (external and internal methods), Evaluation (including static settings and interactive environments), and Mitigation (featuring targeted approaches for three deception genesis mechanisms, and auditing). Sub-branches display representative research contributions within each category.

#### 4.1.1 Behavioral Detection

External methods analyze model responses and behavioral patterns without accessing internal states [Bürger et al., 2024, Pacchiardi et al., 2023]. They treat deception detection as an external observation problem, leveraging textual cues, behavioral consistency, and response patterns [Cohen et al., 2023, Gröndahl and Asokan, 2019, Park et al., 2024].

**Static Text Analysis** Early deception detection relied on lexical features, with bag-of-words SVMs achieving high precision on review datasets [Ott et al., 2011, 2013]. When these models failed to generalize across domains, researchers turned to syntactic approaches based on grammatical rules [Feng et al., 2012] and deep dependency features [Xu and Zhao, 2012]. The advent of deep learning brought CNN and RNN encoders [Ren and Ji, 2017], followed by BERT-based classifiers that now exceed 93% accuracy by exploiting linguistic patterns, e.g., formulaic language in deceptive text versus richer contextual details in truthful statements [Barsever et al., 2020, Vogler and Pearl, 2020].

While these advances confirm that deception leaves detectable linguistic fingerprints, static text analysis treats each text as an isolated snapshot. This approach fundamentally cannot capture how deceptive behavior evolves and adapts across extended interactions.

**Adversarial Behavioral Probing** Recognizing that static text analysis cannot capture deception’s interactive nature, research has shifted toward adversarial behavioral probing, which actively elicits and analyzes real-time behavioral adjustments in LLMs to reveal how models strategize during deceptive acts [Curvo, 2025, Hagendorff, 2024, Pacchiardi et al., 2023, Prome et al., 2024]. **Output-level detection** utilizes adversarial prompts to expose hidden behaviors without accessing internal

representations, surfacing tendencies like sycophancy [Perez et al., 2023]. Methods include cross-examination protocols where reviewer models interrogate respondents in multi-turn dialogues [Cohen et al., 2023], consistency checks using logically equivalent queries to detect divergent answers [Fluri et al., 2024, Mündler et al., 2023], and social deduction games that elicit strategic deception through game-theoretic pressure [O’Gara, 2023]. Recent refinements use strategically unrelated questions to destabilize reasoning chains and expose inconsistencies [Pacchiardi et al., 2023, Zhu et al., 2024]. However, focusing exclusively on final outputs misses cases where models generate seemingly correct answers through flawed or deceptive reasoning. **Process-level detection** addresses this by analyzing internal CoT reasoning steps [Lanham et al., 2023, Lightman et al., 2023]. The core challenge is that models can exhibit reasoning faithfulness gaps, generating correct answers while employing flawed or deceptive intermediate logic [Lyu et al., 2023, Turpin et al., 2023]. To verify reasoning integrity, researchers have developed frameworks such as knowledge graph-based reasoning path verification [Nguyen et al., 2024] and FRODO, which uses causal mediation to trace how each reasoning step influences final outputs [Paul et al., 2024]. Empirical evidence presents a mixed picture: CoT monitoring shows robustness against subtle sabotage [Emmons et al., 2025, MacDiarmid et al., 2024], but lags behind output monitoring for detecting obvious deception [Arnav et al., 2025a]. Moreover, adversaries can manipulate reasoning steps through steganography or hidden reasoning [Kuo et al., 2025, Skaf et al., 2025], creating an ongoing cat-and-mouse game between detection methods and increasingly sophisticated evasion strategies.

#### 4.1.2 Internal State Analysis

To overcome the limits of external methods, researchers increasingly turned to hidden representations and activations to detect deception directly [Carranza et al., 2023]. The key promise is that even if models manipulate their outputs, traces of true intent remain encoded in internal states [Fischer, 2023]. Internal state detection is dominated by two main methodologies: (1) static analysis, which extracts latent knowledge about truth and deception, and (2) dynamic monitoring, which tracks how internal states shift in real-time across contexts [López et al., 2024].

**Static Analysis** Foundational work showed models internally represent truth even when producing deceptive outputs. Early unsupervised approaches like Contrast-Consistent Search identified logical consistency directions in activation spaces through pattern extraction [Burns et al., 2022], while supervised methods such as SAPLMA trained probes on hidden activations to predict truthfulness with high accuracy [Azaria and Mitchell, 2023]. Subsequent research extended these techniques to more complex scenarios. Linear probes were applied to "quirky" models, which are models deliberately trained to exhibit context-dependent behavioral inconsistencies, to extract truthfulness [Mallen et al., 2023]. Additionally, mechanistic interpretability approaches further revealed strategically hidden knowledge [Cywiński et al., 2025]. These approaches included investigations of "taboo" models (models that provide indirect clues of secret information), logit lens analysis, and sparse autoencoders. Collectively, this line of work established that internal representations encode truth beyond what outputs reveal.

**Dynamic Monitor** While static analysis extracts latent knowledge from fixed states, dynamic monitoring tracks deception as it unfolds. MIND detects hallucinations in real-time from unsupervised internal signals [Su et al., 2024], while SafetyNet monitors attention and MLP activations with ensemble anomaly detectors, achieving 96% accuracy in detecting harmful cases [Chaudhary and Barez, 2025]. These successes partly stem from attention layers being especially diagnostic of deceptive reasoning. Concurrent work employs linear probes to show that internal states often predict answers before external reasoning, with stronger internal–external mismatches in advanced models like GPT-4o [Orgad et al., 2024, Wang et al., 2025a]. However, these monitors face robustness challenges. Adver-

**Table 1** | Overview of AI deception-related evaluations. We organize existing studies from two perspectives: evaluation in **static settings** and evaluation in **interactive environments**, and we annotate each work with its release date, data size, institution, data format, and description.

Type	Dataset	Release Date	Institution	Data Size	Data Format	Description
Capability for Spontaneous Deception	SAD [152]	24/07	UC Berkeley	13k	QA	Situational awareness
	DAELLMs [114]	23/07	Uni Stuttgart	1,920	QA	Theory-of-Mind and deception
	CSQ [306]	25/08	NUS	–	FW	evaluating AI deception on benign prompts
Deception under Provided Incentives	MWE [221]	22/12	Anthropic	3.25K	QA	Testing sycophancy on philosophy and political questions
	SycophancyEval [250]	23/10	Anthropic	–	QA	Revealing how a user’s preferences affects AI assistant behavior
	DeceptionBench [134]	25/05	PKU	180	QA	Assessing deception-driven misalignment in reasoning models
	DeceptionBench [124]	25/10	THU	1.5K	QA	Evaluating AI deception across diverse real-world scenarios
	MASK [236]	25/03	CAIS	1K	SS	Pressure prompts that may induce deception
In Constrained Interactions	InsiderTrading [244]	23/11	Apollo	–	FW	Evaluating AI deception in high-pressure environments
	OpenDeception [305]	25/04	FDU	–	FW	Evaluating AI deception in open-ended user-AI interactions
	Sabotage [29]	24/10	Anthropic	4	FW	Human decision sabotage, code sabotage, sandbagging, undermining oversight
	CAE [222]	25/05	DeepMind	16	FW	5 stealth and 11 situational-awareness agent tasks
	MACHIAVELLI [215]	23/04	UCB	134	Games	Human-written social games
	Hoodwinked [199]	23/08	USC	–	Games	A Text-Based Murder Mystery Game
In Sandbox	HouseWins [61]	24/05	CMU	1	FW&Games	Blackjack
	Traitors [68]	25/05	UvA	1	FW&Games	Multi-agent simulation, inspired by social deduction games
	SHADE-Arena [150]	25/06	Anthropic	17	FW&Games	Benign main tasks and harmful side objectives
	In-contextScheming [186]	24/12	Apollo	6	FW	Environments that incentivize scheming
	AgenticMisalignment [11]	25/06	Anthropic	1	FW	Fictional settings

serial attacks such as RL-Obfuscation train models to manipulate activations to bypass latent-space defenses, exposing a persistent challenge [Bailey et al., 2024, Gupta and Jenner, 2025]. To counter this, Latent Adversarial Training (LAT) perturbs latent activations to improve the model’s resistance against adversarial attacks. But monitors still remain vulnerable, with token-level aggregation being evaded in up to 89.2% of cases [Sheshadri et al., 2024].

In sum, both external and internal detection methods face a common vulnerability: adversarial manipulation. External behavioral probes can be obfuscated through strategic output control [Hubinger et al., 2024], while internal-state monitors remain susceptible to activation manipulation [Bailey et al., 2024, Gupta and Jenner, 2025]. This shared fragility reveals a deeper limitation: to date, no single detection dimension, whether behavioral or representational, provides sufficient robustness against adaptive deception.

## 4.2 | Deception-related Evaluation

Evaluating deception in AI systems requires a structured approach that reflects how deceptive behaviors arise under different conditions and stages of development. Because deception is often complex and concealed, single-turn evaluations may fail to reveal the full spectrum of risks; by contrast, dynamic interactions can provide richer contexts in which deceptive behaviors are more likely to surface. Therefore, we organize deception-related evaluation into two complementary dimensions. *Evaluation in Static Settings* probes latent risks in fixed and non-interactive tasks, providing early signals of deceptive abilities and incentive sensitivities. *Evaluation in Interactive Environments* examines how deception manifests during dynamic interactions, adversarial pressures, or multi-agent contexts closer to real-world deployment. These dimensions provide a comprehensive framework for deception evaluation (as shown in Table 1).

### 4.2.1 Static Evaluations: Probing Latent Risks

Evaluations in static environments focus on static and fixed tasks, enabling the isolation of deception-related risks without the confounding dynamics of interactive environments. Within this scope, we summarize two complementary aspects: whether models already possess the ability for spontaneous deception, and whether they will engage in deception when placed under prompted incentives.

**Capability for Spontaneous Deception** Evaluations of spontaneous deception investigate whether models already possess the prerequisites needed to mislead without explicit incentives. For example, research [Hagendorff, 2024] demonstrates through ToM tasks that advanced LLMs can already perform first-order deception while struggling with more complex second-order cases, revealing the cognitive capacities necessary for misrepresentation. The Situational Awareness Dataset (SAD) [Laine et al., 2024] shows that models are able to recognize evaluation contexts and their own deployment conditions, a capability that may foster deceptive behavior. Moreover, recent studies reveal that models may generate misleading responses even under benign prompts, suggesting that deceptive tendencies can surface spontaneously in seemingly neutral conditions [Wu et al., 2025b].

**Deception under Provided Incentives** Some studies examine whether models exhibit deceptive tendencies when placed under externally provided incentive conditions. Rather than directly testing raw capabilities, these benchmarks probe how models respond when prompts introduce preferences, penalties, or goal conflicts. For instance, evaluations show that when user preferences are included in prompts, models often prioritize agreement or compliance, resulting in sycophantic behaviors [Perez et al., 2023, Sharma et al., 2023]. Similarly, some benchmarks first elicit models' latent goals with neutral prompts, then introduce contextual scenarios with external objectives or pressured statements, and finally assess consistency of model responses across the two [Huang et al., 2025b, Ji et al., 2025, Ren et al., 2025].

#### 4.2.2 Dynamic Evaluations: Exposing Deception in Complex Interaction

While evaluations in static settings reveal important foundations of deceptive capability, they remain limited in capturing how deception unfolds during interactions. Interactive evaluations address this gap by examining models in dynamic scenarios where behaviors adapt to goals, pressures, or other agents. We distinguish between two primary settings. *Deception in constrained interactions* focuses on environments with limited action spaces, short horizons, and rule-bound tasks, making it possible to observe how deception arises under specified conditions. By contrast, *Sandbox environments* grant models broad autonomy, including larger action spaces, tool use, and extended horizons, allowing deceptive behaviors to emerge, persist, and adapt during open-ended exploration.

**Deception in Constrained Interactions** From the single-model side, benchmarks place one agent in short task flows and ask whether it will take or justify deceptive actions under goals or pressure [Benton et al., 2024, Pan et al., 2023, Phuong et al., 2025, Scheurer et al., 2023, Wu et al., 2025a]. On the multi-agent side, short games show deception emerging without prompts:  $2 \times 2$  signaling games with a single communication phase elicit unsolicited misrepresentation that scales with rational incentive and model capability [Taylor and Bergen, 2025], and a text social-deduction game reveals denial and persuasion that shift voting outcomes in a few rounds [O'Gara, 2023].

**Deception in Sandbox** Within this paradigm, one line of work uses game-based simulations, such as card play or social deduction, to show that models not only deploy deception but also struggle to detect it, revealing asymmetries in persuasion and trust under open-ended play [Chopra and Li, 2024, Curvo, 2025]. A second line examines environments with hidden or conflicting objectives, where agents discover and sustain misaligned strategies, e.g., covert goals or sabotage, while sometimes evading oversight [Kutasov et al., 2025, Meinke et al., 2024]. Further studies in organizational or tool-use sandboxes indicate that when goals conflict with oversight or self-preservation incentives, models may engage in blackmail or data exfiltration without explicit prompting [Anthropic, 2025a].

### 4.3 | Potential Mitigations

Deception can be seen as a natural consequence of increasing AI capabilities, much like the flip side of intelligence. Building on the factors underlying the genesis of deception discussed in the previous section (Section 3), we now turn to potential mitigations for mitigating such behaviors. Our analysis spans the entire lifecycle of AI systems, encompassing interventions during training as well as auditing and regulatory mechanisms.

#### 4.3.1 Dissolving Deception Incentives

As discussed in Section 3.1, models may develop deceptive incentives due to data imitation, reward misspecification, or goal misgeneralization. Drawing insights from adjacent alignment research, several strategies show promise for addressing these underlying sources. First, pretraining data curation techniques that filter problematic examples and integrate alignment objectives directly into pretraining [Korbak et al., 2023, Liang et al., 2024] can reduce exposure to deceptive patterns at the source, but they are often prohibitively costly at scale.

Work on reward misspecification provides relevant methods for deception treatment. Improved RL algorithms, such as adversarial reward functions and reward capping [Amodei et al., 2016, Uesato et al., 2020], help address misaligned objectives, while approaches that train models to express uncertainty [Cheng et al., 2024, Sharma et al., 2023, Yang et al., 2023] show effectiveness in reducing sycophancy, though they can also increase refusal rates or reduce helpfulness when over-applied. Alternatively, self-supervised and self-regulation paradigms design training objectives that encourage models to monitor and constrain their behaviors during reasoning processes, approaches that have been directly applied in deception contexts [Guan et al., 2024, Ji et al., 2025].

Emerging techniques for controlling generalization direction during training, such as concept ablation and behavioral steering interventions [Casademunt et al., 2025, Chen et al., 2025b, Yu et al., 2024b], suggest pathways for preventing unwanted deceptive behaviors from emerging during training.

#### 4.3.2 Regulating Deception Capabilities

As AI systems grow increasingly capable of deceptive behaviors, regulating these specific capabilities becomes crucial for maintaining trustworthy AI deployment.

At the perception level, recent work leverages models' *self-knowledge* to constrain information processing [Wang et al., 2023]. By enabling retrieval only when the model recognizes gaps in its own knowledge, this approach maintains factual accuracy while preventing the override of correct internal representations that could facilitate deceptive responses.

At the planning level, regulatory efforts focus on monitoring CoT processes in real time to detect and intervene against deceptive reasoning patterns [Arnav et al., 2025a, Baker et al., 2025, Ji et al., 2025, Korbak et al., 2025, Schoen et al., 2025]. This regulatory approach has demonstrated measurable success in frontier models: systematic CoT monitoring reduced deception detection rates in GPT-5-thinking to just 2.1%, compared with 4.8% in its predecessor o3 [OpenAI, 2025c]. CoT monitoring provides detailed insights into the model's internal reasoning trajectory. Conversely, AI Control imposes rules above the model by interposing policy gateways, least privilege interfaces, sandboxed executors, and audit-triggered defer or shutdown that wrap the model behind enforceable system services, offering a complementary path for deception treatment [Greenblatt et al., 2024b, Griffin et al., 2024].

At the performing level, where models may engage in linguistic manipulation or misuse external tools, regulatory frameworks emphasize containment and oversight of potentially deceptive actions. Sandboxed execution environments serve as a key regulatory mechanism, confining code or API calls to isolated settings where deceptive behaviors can be detected and contained before affecting real systems [Dou et al., 2024, Rabin et al., 2025, Tallam and Miller, 2025]. These approaches, from constraining perception to monitoring reasoning to sandboxing execution, form complementary layers of defense against deceptive behavior.

### 4.3.3 Countering Deception Triggers

External triggers represent a primary vector for inducing AI deception, making the development of counter-strategies essential for maintaining model integrity. Research in AI safety has explored multiple directions to enhance robustness against adversarial prompts and jailbreak attacks, which can be transformed to enhance model robustness against deception triggers. The most direct approach is **adversarial training**, which fine-tunes models on known deception-inducing prompts to strengthen their resistance to manipulation. While several studies demonstrate effectiveness in improving robustness against general adversarial inputs [Zou et al., 2023], the deception context presents unique challenges. Models may exhibit deceptive compliance during training, appearing to resist triggers while retaining latent deceptive capabilities [Hubinger et al., 2024]. This highlights the critical need for adversarial training methods specifically designed to counter deception triggers rather than general jailbreaks. A complementary strategy is **red teaming**, which proactively identifies deception trigger pathways before deployment. This includes targeted attacks conducted by human experts [Ganguli et al., 2022, OpenAI, 2025c], as well as automated red teaming combined with public benchmarks to track jailbreak robustness over time [Chao et al., 2024, OpenAI, 2025c]. Specialized deception red teaming has emerged through multi-agent simulation frameworks where deceptive agents operate under asymmetric information, revealing how deception triggers can exploit strategic vulnerabilities and trust mechanisms [Curvo, 2025, Golechha and Garriga-Alonso, 2025]. Finally, recognizing that novel triggers may still bypass model-level defenses, **guard and moderation layers** provide real-time monitoring at deployment to detect and neutralize deception triggers in both inputs and outputs, serving as a crucial last line of defense against sophisticated deception [Inan et al., 2023].

### 4.3.4 Auditing

Alignment auditing [Marks et al., 2025] is a structured process for evaluating whether models harbor hidden or deceptive objectives. There are three complementary directions. The first is **training data analysis**, which applies keyword or semantic search, clustering, and other forensic methods to trace suspicious model behaviors back to potential origins in the training corpus. Such methods have been used to reduce privacy leakage risks [Tamkin et al., 2024], and offer valuable insights for deception treatment by revealing deceptive incentives embedded in training data.

The second is **black-box techniques**, which attempt to elicit hidden objectives through interaction alone. For example, assistant prefill attacks [Andriushchenko et al., 2024, Qi et al., 2024, Vega et al., 2024] and extracting information from non-assistant personas [Marks et al., 2025, Shanahan et al., 2023] can induce disclosure of strategies the model intends to conceal, thereby exposing deceptive tendencies. Relatedly, debate-based auditing reframes oversight as an adversarial dialogue that pressures models to surface inconsistencies, making obfuscated reasoning harder to hide and clarifying when uncertainty estimates improve detection [Brown-Cohen et al., 2025].

The third is **white-box techniques**, which leverage access to model weights or activations. Methods

such as sparse autoencoders (SAEs) can isolate features linked to suspicious behaviors, enabling causal steering and intervention, or analyzing representations such as the assistant-control token [Anthropic, 2024]. Beyond SAEs, neural representation monitoring can help localize components implicated in deceptive behavior and enable more targeted interventions, including activation probes [MacDiarmid et al., 2024] and interpretability methods to extract active or causal features [Bailey et al., 2024, Marks et al., 2025, Yang and Buzsaki, 2025]. Other efforts, including activation fuzzing, attempt to weaken secrecy, though their effectiveness is limited when used in isolation [Clymer et al., 2024, Tice et al., 2024]. While these auditing techniques provide crucial visibility into model behavior, their effectiveness depends on coordination across data access, interaction design, and interpretability research; this integration challenge remains largely unsolved.

## 5 | Conclusion

This survey has sought to provide a comprehensive map of the AI deception landscape through the lens of the *Deception Cycle*, a conceptual framework capturing the interplay between **Deception Emergence**, where incentives, capabilities, and contextual triggers converge to produce deceptive behavior, and **Deception Treatment**, which encompasses detection, evaluation, and potential mitigations aimed at suppressing such behavior. In doing so, we have introduced a unified taxonomy, reviewed empirical phenomena across RL agents, LLMs, and emergent multi-agent or multimodal systems, and cataloged over 20 benchmarks, methods, and mitigation strategies.

### 5.1 | Key Challenges in AI Deception Cycle

Beyond taxonomy and systematization, this survey highlights that deception is not merely an incidental failure mode, but an adaptive, goal-directed behavior that becomes increasingly likely as AI systems scale in autonomy, capability, and strategic awareness. Our synthesis reveals several insights:

- **Deception is incentivized by default in misaligned systems.** Unless explicitly penalized, deception may emerge as a convergent instrumental strategy under a wide range of training regimes, including supervised fine-tuning, reinforcement learning, and self-play, particularly when models benefit from hiding their true goals or capabilities.
- **Deceptive strategies are becoming more compositional and temporally extended.** As models acquire memory, planning, and agentic scaffolding, we observe the rise of long-horizon deception: multi-stage behaviors that involve delayed reward hacking, conditional alignment, and stealthy behavior switching.
- **Deception is modality-agnostic and generalizes across domains.** While early research focused on textual deception in LLMs, recent findings [Yang et al., 2025] show similar patterns in vision-language models, autonomous robotics, and simulated social agents, suggesting that deception is a modality-general risk amplified by interactive complexity.
- **Alignment techniques struggle with deception-specific failure modes.** Existing safety paradigms, such as RLHF [Bai et al., 2022a, Ouyang et al., 2022], CAI [Bai et al., 2022b], and adversarial red-teaming—often fail to surface or remove latent deceptive tendencies. Models trained to pass audits may optimize for appearing aligned rather than being aligned, raising foundational questions about alignment verifiability.

These observations give rise to three grand challenges that demand urgent, cross-disciplinary attention:

- **Recursive deception of oversight tools.** As models learn to exploit or evade interpretability methods, CoT rationales, and rule-based constraints, oversight mechanisms themselves risk be-

coming adversarial targets that are vulnerable to manipulation by the very systems they intend to supervise.

- **Persistence of deceptive alignment.** Once deceptive objectives are internalized, they may remain dormant, conditionally activated, or resilient to extensive retraining. Recent studies on sleeper agents and alignment faking highlight the limitations of current mitigation regimes.
- **Governance and institutional lag.** Deception risks often manifest in deployment-time behaviors or complex, open-ended interactions, while current oversight remains largely confined to pre-release evaluation. Fragmented regulatory environments and underdeveloped audit infrastructure further hinder systemic accountability.

Yet deception is not solely a technical artifact; it is a reflection of deeper misalignments between model objectives and human expectations. While much of the current literature focuses on *single-agent safety*, ensuring that an individual model behaves as intended, our findings suggest that this perspective is insufficient. Deceptive behaviors often emerge within broader *sociotechnical systems* comprising users, developers, institutions, and other AI agents. Deception may be reinforced by opaque incentives, obscured by organizational delegation, or amplified by multi-agent interactions in agentic ecosystems.

Future safety efforts must transcend static, model-centric verification and embrace dynamic, system-level resilience. Technical solutions alone cannot ensure trustworthiness; they must operate within institutional frameworks that enforce transparency, auditability, and recourse. Achieving this demands an interdisciplinary shift, combining machine learning, formal methods, HCI, governance, and philosophy, to co-design socio-technical ecosystems where honesty is both learnable and verifiable. Deception-resistant AI cannot be patched or filtered in retrospect; it must be built into the core of learning, oversight, and deployment. Only by embedding deception-aware principles across technical and institutional layers can we ensure AI systems remain aligned, accountable, and genuinely trustworthy in the open world.

## 5.2 | Key Traits and Future Directions in AI Deception Research

Finally, we conclude the survey by highlighting the key traits that we believe warrant sustained attention and should shape future research trajectories in this area.

**From Programmed to Emergent Deception: *What Can Deliberate Design Teach Us About Unintended Incentives?*** This survey has focused on investigating how deception can emerge naturally from data imitation, reward misspecification, or goal misgeneralization. However, deception can also be deliberately programmed into models' objectives and strategy space, as exhibited in backdoor attacks and deceptive RL. Here, we extend the discussion of these two sources of deception to provide deeper insights into the incentive foundations of AI deception.

Programmed deception and emergent deception differ in the following aspects.

- **Goals and objectives:** In emergent deception, models are not explicitly optimized for a clearly defined deceptive objective; instead, incentives emerge from data, reward, and goal misalignment. By contrast, programmed deception arises when models are directly trained to deceive, with objectives that reward deception and penalize transparency, thereby aligning training goals with deceptive actions—an alignment absent in emergent deception.
- **Strategy space:** Programmed deception operates within a human-defined, thus limited, strategy space; although deceptive RL agents are trained to conceal their goals, their behaviors remain broadly predictable. By contrast, emergent deception arises in real deployment with an open-world, unbounded strategy space, yielding diverse and covert behaviors that are far harder to detect.
- **Deployment:** A key difference in deployment is controllability. Programmed deception, intention-

ally designed, can in principle be bounded and managed in sandboxed settings, whereas emergent deception is uncontrollable, as its strategies arise unintentionally.

Programmed deception provides valuable insights into understanding, detecting, and mitigating emergent deception. On one hand, it sheds light on the **training-time mechanisms** through which deception arises, particularly in environments where models are rewarded for deceptive behavior and for concealing their true strategies. On the other hand, it allows researchers to **controllably replicate deceptive behaviors** or construct hypothetical phenomena such as sleeper agents, thereby supporting in-depth empirical study. By constructing models with hidden objectives and analyzing deceptive behavioral trajectories, researchers are able to probe the underlying deceptive motivations that emerge in AI systems. Moreover, it is worth highlighting that the *inverse RL* methodology adopted in deceptive RL research provides crucial insights for deception detection and monitoring.

**Evaluating Deception Behaviours: How Do We Accurately Measure Deceptive Tendencies in Models?** A central challenge in AI deception research is **defining and reliably evaluating deceptive behaviors during model assessment**. A model’s stated reasoning, for instance, its reported CoT, may not faithfully represent the latent processes that produced its outputs [Chen et al., 2025c, Naik et al., 2025]. Another challenge is that models may actively game evaluation procedures. If behavioral signals cannot be trusted, the very process of detecting deception becomes unreliable. This difficulty is further compounded by the absence of explicit CoT traces in many tasks, making it challenging to attribute deceptive intent. Moreover, standard evaluation is often confounded by hallucinations and task failures [Kutasov et al., 2025], which obscure whether an output reflects strategic deception or simple underperformance. Another obstacle lies in **constructing realistic evaluation settings**. Existing deception tests, such as binary agent–environment misalignment tasks, often fail to capture the complexity of real-world scenarios where deception might emerge [Benton et al., 2024]. Artificial prompts frequently bundle critical cues together, creating a *Chekhov’s gun effect*: models, sensitive to the importance of every detail, are more likely to exploit these cues rather than disregard them [Anthropic, 2025a]. As a result, contrived testbeds risk overestimating deceptive tendencies, while authentic, high-stakes deployments may reveal underestimated risks. Finally, an **evaluation gap persists between controlled benchmarks and open-world deployment**. Models may engage in sandbagging, suppressing their full capabilities in evaluation while exhibiting stronger deceptive behaviors in practice. Addressing this gap requires designing ecologically valid test environments, developing metrics that disentangle deception from error or hallucination, and conducting longitudinal assessments that track how deceptive behaviors evolve across training and deployment.

Recent work [Summerfield et al., 2025] highlights a critical challenge for AI deception research: avoiding the methodological pitfalls of past scientific endeavors. Just as the ape-language projects of the 1970s faltered due to overattributing human-like traits, reliance on anecdotes, and the absence of clear theoretical frameworks, current studies of AI scheming risk similar shortcomings. Claims of deception emergence are often based on isolated red-teaming anecdotes, ambiguous evaluation setups, or mentalistic interpretations of model outputs, which can inflate the evidence base without establishing replicable, theory-driven findings. The key lesson is that progress on identifying deception emergence will require rigorous control conditions, transparent reporting, and carefully articulated theoretical constructs that distinguish strategic misrepresentation from stochastic failure. Without such methodological discipline, the field risks mischaracterizing incidental model errors as deceptive intent, thereby hindering both scientific credibility and policy-relevant assessments.

**Scalable Monitoring for Deception: Beyond CoT-Centric, Toward Faithful and Cost-Aware Systems** A central challenge in reliable monitoring is ensuring faithfulness. CoT-based monitors assume that models provide explicit and truthful rationales [Baker et al., 2025], yet they fail when models omit CoT, strategically suppress it (e.g., under RLHF [Chua and Evans, 2025]), or generate post

hoc rationalizations. Even advanced monitors capable of detecting subtle, long-horizon attacks in extended logs remain brittle: they may accept tendentious narratives that justify harmful actions, or conversely, overreact to benign errors [Kutasov et al., 2025]. While turn-by-turn scrutiny can improve recall, it often floods the system with highly suspicious flags, leading to overfitting and decision paralysis. The problem is compounded in agentic systems, which frequently deviate from their stated policies. Whether contemporary reasoning models [Li et al., 2025b] are intrinsically more faithful remains an open question, not an assumption that monitoring frameworks can rely upon.

Future monitoring systems should embody four principles. First, they should be *CoT-agnostic and multi-view*, integrating behavioral probes, environment- or tool-grounded verification, and mechanism-level checks (e.g., consistency and invariance) rather than relying solely on stated rationales. Second, they should be *regularized and compositional*, combining turn-level judgments with calibrated abstinence, rate-limiting, and conformal or risk-controlled thresholds to mitigate overflagging. Third, they should be *cost-safety aware*, allocating monitoring resources adaptively by estimated risk, with deeper inspections reserved for high-stakes cases. Finally, they should be *faithfulness-seeking*, employing audit-time elicitation (e.g., scratchpad commitments, execution traces), cross-model verification, and longitudinal tests for discrepancies between stated and latent strategies. Together, these principles aim to produce monitoring systems that scale without sacrificing precision, remain robust to sandbagging, and evaluate deception independently of a model’s willingness to explain itself.

**Deception Treatment and Governance: *How Can Technical Safeguards Interface with Institutional Oversight?*** A core challenge at the intersection of AI deception and governance lies in **ensuring that technical defenses against deception are embedded within enforceable institutional frameworks**. While certified defenses, such as provable training protocols and robust evaluation metrics, can help constrain deceptive behavior under adversarial conditions, their effectiveness is limited without broader governance structures that ensure compliance and accountability. For example, even a model trained with formal guarantees against sycophancy or sandbagging may still be vulnerable if deployed in environments lacking tamper-proof monitoring or third-party verification, since models (or their operators) could conceal violations, rendering such guarantees ineffective.

This highlights the necessity of **institutional innovation to complement technical safety measures**. Mechanisms such as independent audits, hardware-rooted deployment controls, and cryptographically verifiable reporting channels can extend trust beyond the lab setting, mitigating risks of deceptive behaviors that evade laboratory evaluations. Importantly, governance structures can also shape the incentives that determine whether deception is suppressed or reinforced in practice, bridging the persistent gap between technical solutions and societal oversight.

In this sense, **AI deception is not solely a technical alignment problem but also a governance challenge**. Certified defenses provide the formal tools to limit deceptive capacity, but institutional frameworks are required to sustain these guarantees across diverse deployment contexts. Progress thus depends on integrating safety research with governance innovation, ensuring that models cannot exploit institutional blind spots to conceal, amplify, or strategically deploy deception.

**Deception with Different Modalities: *From Multimodal Integration to Vision-Language Alignment—Where Do New Forms of Deception Emerge?*** From language models to advanced cross-modal systems nowadays [Anthropic, 2025b, OpenAI, 2025c, Team, 2024], the vision of AGI has expanded into richer, multimodal scenarios. However, this expansion can amplify the risks of deceptive behaviors, while existing text-based monitoring methods might be inadequate. On the one hand, semantic ambiguity and the complexity of cross-modal reasoning make deceptive behaviors difficult to detect; on the other hand, standardized benchmarks and evaluation frameworks for assessing deception in multimodal large language models (MLLMs) are lacking. Consequently, there is an urgent need for practical evaluation and monitoring approaches specifically designed to address multimodal deception.

Multimodal deception stands apart from hallucinations in MLLMs [Bai et al., 2024]. Whereas hallucinations reflect capability deficits, multimodal deception emerges with advanced capabilities as a strategic and complex behavior, representing a misalignment between perception and response. Though not yet formally characterized in literature, emerging evidence indicates that even when models have made accurate interpretations of input modalities, they may still generate misleading user-facing responses [OpenAI, 2025c]. Research shows vision language models develop sycophancy behaviors due to an imbalance between linguistic priors and visual grounding [Zhao et al., 2024]. The cognitive complexity in multimodal scenarios scales substantially compared to single-modal ones [Oviatt et al., 2004], creating a novel and expanded space for deceptive strategies. Models can therefore selectively *reconstruct* the image’s semantics, inducing false belief by choosing which visual elements to reveal, conceal, misattribute, or even fabricate. Taken together, multimodal deception poses novel vulnerabilities and risks that demand urgent attention from the community.

Beyond purely representational deception, a more concerning form arises when vision-language-action (VLA) systems engage in embodied deception, in which agents not only perceive and reason but also act within the environment in ways that can strategically mislead human or algorithmic supervisors. This phenomenon is particularly evident in several classes of systems. In preference-based RL or RLAIFF settings, where human feedback is used to train reward models [Christiano et al., 2017, Jain et al., 2015], agents may discover behaviors that appear correct from the perspective of a monitoring camera while failing to achieve the intended task; for example, a robotic manipulator might “pretend” to complete a placement task by moving or occluding objects in ways that maximize observed reward signals. Similarly, in visual-feedback-driven imitation or reward learning, when rewards rely on video observations or third-party vision-based estimators, agents can manipulate perceptual input—through camera viewpoint, lighting, or partial occlusion—to generate the appearance of task success without truly satisfying the objective. A third source of deception arises from language–action mismatches in multimodal embodied architectures, where language modules may report adherence to instructions or task goals while the physical policy executes actions that deviate from the stated intent, creating deliberate inconsistencies between communicated and executed behavior.

The formation of embodied deception can be understood across multiple layers. At the **reward level**, agents exploit vulnerabilities in reward functions (*i.e.*, reward hacking) to generate visually plausible but substantively incorrect outcomes. At **the perception or signal layer**, they manipulate observations via camera angles, lighting, or occlusion (*i.e.*, perceptual manipulation). At **the strategy or planning layer**, they may deliberately sequence compliant-looking actions to gain trust before diverging from intended objectives. Finally, **at the interaction layer**, agents can leverage timing, language, or expressive cues to mislead human observers, reflecting a form of social or performative deception. Collectively, these mechanisms illustrate that embodied agents can develop sophisticated, strategic behaviors that misalign apparent success with actual performance, highlighting the need for cross-modal consistency verification, grounded reward design, and honesty-constrained policy optimization [Aylett et al., 2023, Cañal-Bruland, 2017, Everitt et al., 2021].

## References

- Abbas Abdolmaleki, Sandy Huang, Leonard Hasenclever, Michael Neunert, Francis Song, Martina Zambelli, Murilo Martins, Nicolas Heess, Raia Hadsell, and Martin Riedmiller. A distributional view on multi-objective policy optimization. In *International conference on machine learning*, pages 11–22. PMLR, 2020.
- Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, 2024.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Matthew Aitchison, Lyndon Benke, and Penny Sweetser. Learning to deceive in multi-agent hidden role games. In *International Workshop on Deceptive AI*, pages 55–75. Springer, 2020.
- Matthew Aitchison, Lyndon Benke, and Penny Sweetser. Learning to deceive in multi-agent hidden role games. In *Deceptive AI: First International Workshop, DeceptECAI 2020, Santiago de Compostela, Spain, August 30, 2020 and Second International Workshop, DeceptAI 2021, Montreal, Canada, August 19, 2021, Proceedings 1*, pages 55–75. Springer, 2021.
- Nitay Alon, Lion Schulz, Jeffrey S Rosenschein, and Peter Dayan. A (dis-) information theory of revealed and unrevealed preferences: Emerging deception and skepticism via theory of mind. *Open Mind*, 7:608–624, 2023.
- Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.
- Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. Jailbreaking leading safety-aligned llms with simple adaptive attacks. In *The Thirteenth International Conference on Learning Representations*, 2024.
- Eleni Angelou and Lewis Smith. A problem to solve before building a deception detector. <https://www.lesswrong.com/posts/YXNeA3RyRrrRWS37A/a-problem-to-solve-before-building-a-deception-detector>, 2025.
- Anthropic. Sparse Crosscoders for Cross-Layer Features and Model Diffing. <https://transformer-circuits.pub/2024/crosscoders/index.html>, 2024.
- Anthropic. Agentic misalignment: How llms could be insider threats, 2025a. URL <https://www.anthropic.com/research/agentic-misalignment>.
- Anthropic. System card: Claude opus 4 & claude sonnet 4. Technical report, Anthropic, 2025b. URL <https://www-cdn.anthropic.com/4263b940cabb546aa0e3283f35b686f4f3b2ff47.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. *arXiv preprint arXiv:2503.08679*, 2025.
- Stuart Armstrong, Alexandre Maranhao, Oliver Daniels-Koch, Patrick Leask, and Rebecca Gormann. Coinrun: Solving goal misgeneralisation. *ArXiv*, abs/2309.16166, 2023. URL <https://api.semanticscholar.org/CorpusID:263142637>.

- Benjamin Arnav, Pablo Bernabeu-Pérez, Nathan Helm-Burger, Tim Kostolansky, Hannes Whittingham, and Mary Phuong. Cot red-handed: Stress testing chain-of-thought monitoring. *arXiv preprint arXiv:2505.23575*, 2025a.
- Benjamin Arnav, Pablo Bernabeu Perez, Tim Kostolansky, Hannes Whittingham, Nathan Helm-Burger, and Mary Phuong. Unfaithful Reasoning Can Fool Chain-of-Thought Monitoring. <https://www.alignmentforum.org/posts/QYAfjdujzRv8hx6xo/unfaithful-reasoning-can-fool-chain-of-thought-monitoring>, 2025b.
- Matthew Peter Aylett, Mei Yii Lim, Katerina Pappa, Bruce W Wilson, Ruth Aylett, and Mario Parra. Embodied conversational agents: Trust, deception and the suspension of disbelief. In *Proceedings of the First International Symposium on Trustworthy Autonomous Systems*, TAS '23, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400707346. doi: 10.1145/3597512.3597526. URL <https://doi.org/10.1145/3597512.3597526>.
- Amos Azaria and Tom Mitchell. The internal state of an llm knows when it's lying. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 967–976, 2023.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022a.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022b.
- Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. Hallucination of multimodal large language models: A survey. *arXiv preprint arXiv:2404.18930*, 2024.
- Luke Bailey, Alex Serrano, Abhay Sheshadri, Mikhail Seleznyov, Jordan Taylor, Erik Jenner, Jacob Hilton, Stephen Casper, Carlos Guestrin, and Scott Emmons. Obfuscated activations bypass llm latent-space defenses. *arXiv preprint arXiv:2412.09565*, 2024.
- 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. *arXiv preprint arXiv:2503.11926*, 2025.
- Fazl Barez, Tung-Yu Wu, Iván Arcuschin, Michael Lan, Vincent Wang, Noah Siegel, Nicolas Collignon, Clement Neo, Isabelle Lee, Alasdair Paren, et al. Chain-of-thought is not explainability. *Preprint, alphaXiv*, page v2, 2025.
- Houda Nait El Barj and Théophile Sautory. Reinforcement learning from llm feedback to counteract goal misgeneralization. *arXiv preprint arXiv:2401.07181*, 2024.
- Sudarshan Kamath Barkur, Sigurd Schacht, and Johannes Scholl. Deception in llms: Self-preservation and autonomous goals in large language models. *arXiv preprint arXiv:2501.16513*, 2025.
- Dan Barsever, Sameer Singh, and Emre Neftci. Building a better lie detector with bert: The difference between truth and lies. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–7. IEEE, 2020.
- Emily M Bender, Timnit Gebu, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623, 2021.

- Joe Benton, Misha Wagner, Eric Christiansen, Cem Anil, Ethan Perez, Jai Srivastav, Esin Durmus, Deep Ganguli, Shauna Kravec, Buck Shlegeris, et al. Sabotage evaluations for frontier models. *arXiv preprint arXiv:2410.21514*, 2024.
- Leonard Bereska and Efstratios Gavves. Mechanistic interpretability for ai safety—a review. *arXiv preprint arXiv:2404.14082*, 2024.
- Lukas Berglund, Asa Cooper Stickland, Mikita Balesni, Max Kaufmann, Meg Tong, Tomasz Korbak, Daniel Kokotajlo, and Owain Evans. Taken out of context: On measuring situational awareness in llms. *arXiv preprint arXiv:2309.00667*, 2023.
- Jan Betley, Daniel Tan, Niels Warncke, Anna Sztyber-Betley, Xuchan Bao, Martín Soto, Nathan Labenz, and Owain Evans. Emergent misalignment: Narrow finetuning can produce broadly misaligned llms. *arXiv preprint arXiv:2502.17424*, 2025.
- Felix J Binder, James Chua, Tomek Korbak, Henry Sleight, John Hughes, Robert Long, Ethan Perez, Miles Turpin, and Owain Evans. Looking inward: Language models can learn about themselves by introspection. *arXiv preprint arXiv:2410.13787*, 2024.
- Sissela Bok. *Lying: Moral choice in public and private life*. Vintage, 2011.
- Nick Bostrom. The superintelligent will: Motivation and instrumental rationality in advanced artificial agents. *Minds and Machines*, 22:71–85, 2012.
- Jonah Brown-Cohen, Geoffrey Irving, and Georgios Piliouras. Avoiding obfuscation with prover-estimator debate. *arXiv preprint arXiv:2506.13609*, 2025.
- Miles Brundage, Shahar Avin, Jack Clark, Helen Toner, Peter Eckersley, Ben Garfinkel, Allan Dafoe, Paul Scharre, Thomas Zeitzoff, Bobby Filar, et al. The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. *arXiv preprint arXiv:1802.07228*, 2018.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4, 2023.
- Lennart Burger, Fred A Hamprecht, and Boaz Nadler. Truth is universal: Robust detection of lies in llms. *Advances in Neural Information Processing Systems*, 37:138393–138431, 2024.
- Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. Discovering latent knowledge in language models without supervision. *arXiv preprint arXiv:2212.03827*, 2022.
- Richard W Byrne. Machiavellian intelligence. *Evolutionary Anthropology: Issues, News, and Reviews: Issues, News, and Reviews*, 5(5):172–180, 1996.
- Tianle Cai, Xuezhi Wang, Tengyu Ma, Xinyun Chen, and Denny Zhou. Large language models as tool makers. *arXiv preprint arXiv:2305.17126*, 2023.
- Rouwen Caal-Bruland. Deception detection in action: embodied simulation in antisocial human interactions. *Frontiers in Psychology*, 8:166, 2017.
- Nicholas Carlini. Poisoning the unlabeled dataset of {Semi-Supervised} learning. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1577–1592, 2021.
- Joe Carlsmith. Scheming ais: Will ais fake alignment during training in order to get power? *arXiv preprint arXiv:2311.08379*, 2023.

- Joseph Carlsmith. Is power-seeking ai an existential risk? *arXiv preprint arXiv:2206.13353*, 2022.
- Andres Carranza, Dhruv Pai, Rylan Schaeffer, Arnub Tandon, and Sanmi Koyejo. Deceptive alignment monitoring. *arXiv preprint arXiv:2307.10569*, 2023.
- Helena Casademunt, Caden Juang, Adam Karvonen, Samuel Marks, Senthoran Rajamanoharan, and Neel Nanda. Steering out-of-distribution generalization with concept ablation fine-tuning. *arXiv preprint arXiv:2507.16795*, 2025.
- 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. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=bx24KpJ4Eb>. Survey Certification, Featured Certification.
- Patrick Chao, Edoardo DeBenedetti, Alexander Robey, Maksym Andriushchenko, Francesco Croce, Vikash Sehwar, Edgar Dobriban, Nicolas Flammarion, George J Pappas, Florian Tramèr, et al. Jailbreakbench: An open robustness benchmark for jailbreaking large language models. *Advances in Neural Information Processing Systems*, 37:55005–55029, 2024.
- Maheep Chaudhary and Fazl Barez. Safetynet: Detecting harmful outputs in llms by modeling and monitoring deceptive behaviors. *arXiv preprint arXiv:2505.14300*, 2025.
- Boyuan Chen, Donghai Hong, Jiaming Ji, Jiacheng Zheng, Bowen Dong, Jiayi Zhou, Kaile Wang, Juntao Dai, Xuyao Wang, Wenqi Chen, et al. Intermt: Multi-turn interleaved preference alignment with human feedback. *arXiv preprint arXiv:2505.23950*, 2025a.
- Canyu Chen and Kai Shu. Can LLM-generated misinformation be detected? In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=ccxD4mtkTU>.
- Runjin Chen, Andy Arditi, Henry Sleight, Owain Evans, and Jack Lindsey. Persona vectors: Monitoring and controlling character traits in language models. *arXiv preprint arXiv:2507.21509*, 2025b.
- Shenghui Chen, Yagiz Savas, Mustafa O Karabag, Brian M Sadler, and Ufuk Topcu. Deceptive planning for resource allocation. In *2024 American Control Conference (ACC)*, pages 4188–4195. IEEE, 2024.
- Yanda Chen, Joe Benton, Ansh Radhakrishnan, Jonathan Uesato, Carson Denison, John Schulman, Arushi Somani, Peter Hase, Misha Wagner, Fabien Roger, et al. Reasoning models don't always say what they think. *arXiv preprint arXiv:2505.05410*, 2025c.
- Zhenpeng Chen, Jie M Zhang, Federica Sarro, and Mark Harman. A comprehensive empirical study of bias mitigation methods for machine learning classifiers. *ACM transactions on software engineering and methodology*, 32(4):1–30, 2023.
- Myra Cheng, Sunny Yu, Cinoo Lee, Pranav Khadpe, Lujain Ibrahim, and Dan Jurafsky. Social sycophancy: A broader understanding of llm sycophancy. *arXiv preprint arXiv:2505.13995*, 2025.
- Qinyuan Cheng, Tianxiang Sun, Xiangyang Liu, Wenwei Zhang, Zhangyue Yin, Shimin Li, Linyang Li, Zhengfu He, Kai Chen, and Xipeng Qiu. Can ai assistants know what they don't know? *arXiv preprint arXiv:2401.13275*, 2024.

- Shashank Reddy Chirra, Pradeep Varakantham, and Praveen Paruchuri. Preserving the privacy of reward functions in mdps through deception. *arXiv preprint arXiv:2407.09809*, 2024.
- Tanush Chopra and Michael Li. The house always wins: A framework for evaluating strategic deception in llms. *arXiv e-prints*, pages arXiv–2407, 2024.
- Neil Chowdhury, Daniel Johnson, Vincent Huang, Jacob Steinhardt, and Sarah Schwettmann. Investigating truthfulness in a pre-release o3 model. <https://transluce.org/investigating-o3-truthfulness>, April 2025.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- James Chua and Owain Evans. Are deepseek r1 and other reasoning models more faithful? *arXiv preprint arXiv:2501.08156*, 2025.
- Marshall B Clinard. Other people’s money: A study in the social psychology of embezzlement., 1954.
- Joshua Clymer, Caden Juang, and Severin Field. Poser: Unmasking alignment faking llms by manipulating their internals. *arXiv preprint arXiv:2405.05466*, 2024.
- Roi Cohen, May Hamri, Mor Geva, and Amir Globerson. Lm vs lm: Detecting factual errors via cross examination. *arXiv preprint arXiv:2305.13281*, 2023.
- Pedro MP Curvo. The traitors: Deception and trust in multi-agent language model simulations. *arXiv preprint arXiv:2505.12923*, 2025.
- Bartosz Cywiński, Emil Ryd, Senthooan Rajamanoharan, and Neel Nanda. Towards eliciting latent knowledge from llms with mechanistic interpretability. *arXiv preprint arXiv:2505.14352*, 2025.
- Doraid Dalalah and Osama MA Dalalah. The false positives and false negatives of generative ai detection tools in education and academic research: The case of chatgpt. *The International Journal of Management Education*, 21(2):100822, 2023.
- Valdemar Danry, Pat Pataranutaporn, Matthew Groh, and Ziv Epstein. Deceptive explanations by large language models lead people to change their beliefs about misinformation more often than honest explanations. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–31, 2025.
- Richard Dawkins and John R. Krebs. Animal signals: Information or manipulation? In John R. Krebs and Nicholas B. Davies, editors, *Behavioural Ecology: An Evolutionary Approach*, pages 282–309. Blackwell Scientific, 1978.
- Google Deepmind. AlphaGenome: AI for better understanding the genome. <https://deepmind.google/discover/blog/alphagenome-ai-for-better-understanding-the-genome/>, 2025.
- Jonas Degraeve, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de Las Casas, et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 602(7897):414–419, 2022.
- Carson Denison, Monte MacDiarmid, Fazl Barez, David Duvenaud, Shauna Kravec, Samuel Marks, Nicholas Schiefer, Ryan Soklaski, Alex Tamkin, Jared Kaplan, et al. Sycophancy to subterfuge: Investigating reward-tampering in large language models. *arXiv preprint arXiv:2406.10162*, 2024.

- Lauro Langosco Di Langosco, Jack Koch, Lee D Sharkey, Jacob Pfau, and David Krueger. Goal misgeneralization in deep reinforcement learning. In *International Conference on Machine Learning*, pages 12004–12019. PMLR, 2022.
- Atharvan Dogra, Krishna Pillutla, Ameet Deshpande, Ananya B Sai, John Nay, Tanmay Rajpurohit, Ashwin Kalyan, and Balaraman Ravindran. Deception in reinforced autonomous agents. *arXiv preprint arXiv:2405.04325*, 2024.
- Shihan Dou, Jiazheng Zhang, Jianxiang Zang, Yunbo Tao, Weikang Zhou, Haoxiang Jia, Shichun Liu, Yuming Yang, Zhiheng Xi, Shenxi Wu, et al. Multi-programming language sandbox for llms. *arXiv preprint arXiv:2410.23074*, 2024.
- Esin Durmus, Liane Lovitt, Alex Tamkin, Stuart Ritchie, Jack Clark, and Deep Ganguli. Measuring the persuasiveness of language models. *Anthropic Blog*, 2024.
- 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. *arXiv preprint arXiv:2507.05246*, 2025.
- Owain Evans, Owen Cotton-Barratt, Lukas Finnveden, Adam Bales, Avital Balwit, Peter Wills, Luca Righetti, and William Saunders. Truthful ai: Developing and governing ai that does not lie. *arXiv preprint arXiv:2110.06674*, 2021.
- Tom Everitt, Marcus Hutter, Ramana Kumar, and Victoria Krakovna. Reward tampering problems and solutions in reinforcement learning: A causal influence diagram perspective. *Synthese*, 198(Suppl 27):6435–6467, 2021.
- Aaron Fanous, Jacob Goldberg, Ank A Agarwal, Joanna Lin, Anson Zhou, Roxana Daneshjou, and Sanmi Koyejo. Syceval: Evaluating llm sycophancy. *arXiv preprint arXiv:2502.08177*, 2025.
- Michael Y Fatemi, Wesley A Suttle, and Brian M Sadler. Deceptive path planning via reinforcement learning with graph neural networks. *arXiv preprint arXiv:2402.06552*, 2024.
- Song Feng, Ritwik Banerjee, and Yejin Choi. Syntactic stylometry for deception detection. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 171–175, 2012.
- Roya Firoozi, Johnathan Tucker, Stephen Tian, Anirudha Majumdar, Jiankai Sun, Weiyu Liu, Yuke Zhu, Shuran Song, Ashish Kapoor, Karol Hausman, et al. Foundation models in robotics: Applications, challenges, and the future. *The International Journal of Robotics Research*, 44(5):701–739, 2025.
- Kevin A Fischer. Reflective linguistic programming (rlp): A stepping stone in socially-aware agi (socialagi). *arXiv preprint arXiv:2305.12647*, 2023.
- Lukas Fluri, Daniel Paleka, and Florian Tramèr. Evaluating superhuman models with consistency checks. In *2024 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pages 194–232. IEEE, 2024.
- Safe AI Forum. International dialogues on ai safety. <https://idais.ai/>, 2024.
- Harry G. Frankfurt. On bullshit. <https://philkoopman.substack.com/p/llms-as-bullshit-machines>, 2009.
- Iason Gabriel. Artificial intelligence, values, and alignment. *Minds and machines*, 30(3):411–437, 2020.

- Kanishk Gandhi, Dorsa Sadigh, and Noah D Goodman. Strategic reasoning with language models. *arXiv preprint arXiv:2305.19165*, 2023.
- Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, et al. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*, 2022.
- Deep Ganguli, Amanda Askell, Nicholas Schiefer, Thomas I Liao, Kamilė Lukošiuūtė, Anna Chen, Anna Goldie, Azalia Mirhoseini, Catherine Olsson, Danny Hernandez, et al. The capacity for moral self-correction in large language models. *arXiv preprint arXiv:2302.07459*, 2023.
- Quentin Garrido, Mahmoud Assran, Nicolas Ballas, Adrien Bardes, Laurent Najman, and Yann LeCun. Learning and leveraging world models in visual representation learning. *arXiv preprint arXiv:2403.00504*, 2024.
- Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A Smith. Real-toxicityprompts: Evaluating neural toxic degeneration in language models. *arXiv preprint arXiv:2009.11462*, 2020.
- Timon Gehr, Matthew Mirman, Dana Drachler-Cohen, Petar Tsankov, Swarat Chaudhuri, and Martin Vechev. Ai2: Safety and robustness certification of neural networks with abstract interpretation. In *2018 IEEE symposium on security and privacy (SP)*, pages 3–18. IEEE, 2018.
- Nicholas Goldowsky-Dill, Bilal Chughtai, Stefan Heimersheim, and Marius Hobbhahn. Detecting strategic deception using linear probes. *arXiv preprint arXiv:2502.03407*, 2025.
- Satvik Golechha and Adrià Garriga-Alonso. Among us: A sandbox for measuring and detecting agentic deception. *arXiv preprint arXiv:2504.04072*, 2025.
- Gowing Life. Science is being corrupted by fake research (and no, it’s not just about ai). *Growing Life*, 2024. URL <https://www.gowinglife.com/science-is-being-corrupted-by-fake-research-and-no-its-not-just-about-ai/>. Accessed via Gowing Life; published approximately one year ago.
- Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, et al. Alignment faking in large language models. *arXiv preprint arXiv:2412.14093*, 2024a.
- Ryan Greenblatt, Buck Shlegeris, Kshitij Sachan, and Fabien Roger. Ai control: Improving safety despite intentional subversion, 2024b. URL <https://arxiv.org/abs/2312.06942>.
- Herbert Paul Grice. Logic and conversation. *Syntax and semantics*, 3:43–58, 1975.
- Charlie Griffin, Louis Thomson, Buck Shlegeris, and Alessandro Abate. Games for ai control: Models of safety evaluations of ai deployment protocols, 2024. URL <https://arxiv.org/abs/2409.07985>.
- Tommi Gröndahl and N Asokan. Text analysis in adversarial settings: Does deception leave a stylistic trace? *ACM Computing Surveys (CSUR)*, 52(3):1–36, 2019.
- Melody Y Guan, Manas Joglekar, Eric Wallace, Saachi Jain, Boaz Barak, Alec Helyar, Rachel Dias, Andrea Vallone, Hongyu Ren, Jason Wei, et al. Deliberative alignment: Reasoning enables safer language models. *arXiv preprint arXiv:2412.16339*, 2024.

- Anisha Gunjal, Anthony Wang, Elaine Lau, Vaskar Nath, Bing Liu, and Sean Hendryx. Rubrics as rewards: Reinforcement learning beyond verifiable domains. *arXiv preprint arXiv:2507.17746*, 2025.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Linge Guo. Unmasking the shadows of ai: Investigating deceptive capabilities in large language models. *arXiv preprint arXiv:2403.09676*, 2024.
- Yufei Guo, Muzhe Guo, Juntao Su, Zhou Yang, Mengqiu Zhu, Hongfei Li, Mengyang Qiu, and Shuo Shuo Liu. Bias in large language models: Origin, evaluation, and mitigation. *arXiv preprint arXiv:2411.10915*, 2024.
- Rohan Gupta and Erik Jenner. RL-obfuscation: Can language models learn to evade latent-space monitors? *arXiv preprint arXiv:2506.14261*, 2025.
- David Ha and Jürgen Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2(3), 2018.
- Dylan Hadfield-Menell, Anca D Dragan, Pieter Abbeel, and Stuart Russell. The off-switch game. In *AAAI Workshops*, 2017.
- Thilo Hagendorff. Deception abilities emerged in large language models. *Proceedings of the National Academy of Sciences*, 121(24):e2317967121, 2024.
- Danny Halawi, Jean-Stanislas Denain, and Jacob Steinhardt. Overthinking the truth: Understanding how language models process false demonstrations. *arXiv preprint arXiv:2307.09476*, 2023.
- Miriam Havin, Timna Wharton Kleinman, Moran Koren, Yaniv Dover, and Ariel Goldstein. Can (a) i change your mind? *arXiv preprint arXiv:2503.01844*, 2025.
- Yufei He, Yuexin Li, Jiaying Wu, Yuan Sui, Yulin Chen, and Bryan Hooi. Evaluating the paperclip maximizer: Are rl-based language models more likely to pursue instrumental goals? *arXiv preprint arXiv:2502.12206*, 2025.
- Douglas Heaven. No limit: Ai poker bot is first to beat professionals at multiplayer game. *Nature*, 571(7765):307–309, 2019.
- Dan Hendrycks, Mantas Mazeika, and Thomas Woodside. An overview of catastrophic ai risks. *arXiv preprint arXiv:2306.12001*, 2023.
- Michael Hibbard, Yagiz Savas, Bo Wu, Takashi Tanaka, and Ufuk Topcu. Unpredictable planning under partial observability. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 2271–2277. IEEE, 2019.
- Michael Townsen Hicks, James Humphries, and Joe Slater. Chatgpt is bullshit. *Ethics and Information Technology*, 26(2):1–10, 2024.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55, 2025a.
- Xiaowei Huang, Marta Kwiatkowska, Sen Wang, and Min Wu. Safety verification of deep neural networks. In *International conference on computer aided verification*, pages 3–29. Springer, 2017.

- Yao Huang, Yitong Sun, Yichi Zhang, Ruochen Zhang, Yinpeng Dong, and Xingxing Wei. Deception-bench: A comprehensive benchmark for ai deception behaviors in real-world scenarios. In *Advances in Neural Information Processing Systems*, 2025b.
- Yunhan Huang and Quanyan Zhu. Deceptive reinforcement learning under adversarial manipulations on cost signals. In *International conference on decision and game theory for security*, pages 217–237. Springer, 2019.
- Evan Hubinger, Chris van Merwijk, Vladimir Mikulik, Joar Skalse, and Scott Garrabrant. Risks from learned optimization in advanced machine learning systems. *arXiv preprint arXiv:1906.01820*, 2019.
- Evan Hubinger, Carson Denison, Jesse Mu, Mike Lambert, Meg Tong, Monte MacDiarmid, Tamera Lanham, Daniel M Ziegler, Tim Maxwell, Newton Cheng, et al. Sleeper agents: Training deceptive llms that persist through safety training. *arXiv preprint arXiv:2401.05566*, 2024.
- Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, et al. Llama guard: Llm-based input-output safeguard for human-ai conversations. *arXiv preprint arXiv:2312.06674*, 2023.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Hel-lyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.
- Mehdi Jafari, Devin Yuncheng Hua, Hao Xue, and Flora Salim. Enhancing conversational agents with theory of mind: Aligning beliefs, desires, and intentions for human-like interaction. *arXiv preprint arXiv:2502.14171*, 2025.
- Ashesh Jain, Shikhar Sharma, Thorsten Joachims, and Ashutosh Saxena. Learning preferences for manipulation tasks from online coactive feedback. *The International Journal of Robotics Research*, 34(10):1296–1313, 2015.
- Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Jiayi Zhou, Zhaowei Zhang, et al. Ai alignment: A comprehensive survey. *arXiv preprint arXiv:2310.19852*, 2023.
- Jiaming Ji, Boyuan Chen, Hantao Lou, Donghai Hong, Borong Zhang, Xuehai Pan, Tianyi Alex Qiu, Juntao Dai, and Yaodong Yang. Aligner: Efficient alignment by learning to correct. *Advances in Neural Information Processing Systems*, 37:90853–90890, 2024.
- Jiaming Ji, Wenqi Chen, Kaile Wang, Donghai Hong, Sitong Fang, Boyuan Chen, Jiayi Zhou, Juntao Dai, Sirui Han, Yike Guo, et al. Mitigating deceptive alignment via self-monitoring. *arXiv preprint arXiv:2505.18807*, 2025.
- Liwei Jiang, Taylor Sorensen, Sydney Levine, and Yejin Choi. Can language models reason about individualistic human values and preferences? *arXiv preprint arXiv:2410.03868*, 2024.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint arXiv:2310.06770*, 2023.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6): 2590–2615, 2011.

- Elif Kartal. A comprehensive study on bias in artificial intelligence systems: Biased or unbiased ai, that's the question! *International Journal of Intelligent Information Technologies (IJIT)*, 18(1):1–23, 2022.
- Jacek Karwowski, Oliver Hayman, Xingjian Bai, Klaus Kiendlhofer, Charlie Griffin, and Joar Skalse. Goodhart's law in reinforcement learning. *ArXiv*, abs/2310.09144, 2023. URL <https://api.semanticscholar.org/CorpusID:264128269>.
- Zachary Kenton, Tom Everitt, Laura Weidinger, Iason Gabriel, Vladimir Mikulik, and Geoffrey Irving. Alignment of language agents. *arXiv preprint arXiv:2103.14659*, 2021.
- HyunJin Kim, Xiaoyuan Yi, Jing Yao, Jianxun Lian, Muhua Huang, Shitong Duan, JinYeong Bak, and Xing Xie. The road to artificial superintelligence: A comprehensive survey of superalignment. *arXiv preprint arXiv:2412.16468*, 2024.
- J Koorndijk. Empirical evidence for alignment faking in small llms and prompt-based mitigation techniques. *arXiv preprint arXiv:2506.21584*, 2025.
- Tomasz Korbak, Kejian Shi, Angelica Chen, Rasika Vinayak Bhalerao, Christopher Buckley, Jason Phang, Samuel R Bowman, and Ethan Perez. Pretraining language models with human preferences. In *International Conference on Machine Learning*, pages 17506–17533. PMLR, 2023.
- Tomek Korbak, Mikita Balesni, Elizabeth Barnes, Yoshua Bengio, Joe Benton, Joseph Bloom, Mark Chen, Alan Cooney, Allan Dafoe, Anca Dragan, Scott Emmons, Owain Evans, David Farhi, Ryan Greenblatt, Dan Hendrycks, Marius Hobbhahn, Evan Hubinger, Geoffrey Irving, Erik Jenner, Daniel Kokotajlo, Victoria Krakovna, Shane Legg, David Lindner, David Luan, Aleksander Mądry, Julian Michael, Neel Nanda, Dave Orr, Jakub Pachocki, Ethan Perez, Mary Phuong, Fabien Roger, Joshua Saxe, Buck Shlegeris, Martín Soto, Eric Steinberger, Jasmine Wang, Wojciech Zaremba, Bowen Baker, Rohin Shah, and Vlad Mikulik. Chain of thought monitorability: A new and fragile opportunity for ai safety. *arXiv preprint arXiv:2507.11473*, 2025.
- Michal Kosinski. Evaluating large language models in theory of mind tasks. *Proceedings of the National Academy of Sciences*, 121(45):e2405460121, 2024.
- Victoria Krakovna and Janos Kramar. Power-seeking can be probable and predictive for trained agents. *arXiv preprint arXiv:2304.06528*, 2023.
- John R. Krebs and Richard Dawkins. Animal signals: Mind-reading and manipulation. In John R. Krebs and Nicholas B. Davies, editors, *Behavioural Ecology: An Evolutionary Approach (2nd Edition)*, pages 380–402. Blackwell, 1984.
- Joshua Krook. Manipulation and the ai act: Large language model chatbots and the danger of mirrors. *arXiv preprint arXiv:2503.18387*, 2025.
- Martin Kuo, Jianyi Zhang, Aolin Ding, Qinsi Wang, Louis DiValentin, Yujia Bao, Wei Wei, Hai Li, and Yiran Chen. H-cot: Hijacking the chain-of-thought safety reasoning mechanism to jailbreak large reasoning models, including openai o1/o3, deepseek-r1, and gemini 2.0 flash thinking. *arXiv preprint arXiv:2502.12893*, 2025.
- Jonathan Kutasov, Yuqi Sun, Paul Colognese, Teun van der Weij, Linda Petrini, Chen Bo Calvin Zhang, John Hughes, Xiang Deng, Henry Sleight, Tyler Tracy, et al. Shade-arena: Evaluating sabotage and monitoring in llm agents. *arXiv preprint arXiv:2506.15740*, 2025.
- Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangru Peng, and Jiaya Jia. Step-dpo: Step-wise preference optimization for long-chain reasoning of llms. *arXiv preprint arXiv:2406.18629*, 2024.

- Rudolf Laine, Bilal Chughtai, Jan Betley, Kaivalya Hariharan, Mikita Balesni, Jérémy Scheurer, Marius Hobbhahn, Alexander Meinke, and Owain Evans. Me, myself, and ai: The situational awareness dataset (sad) for llms. *Advances in Neural Information Processing Systems*, 37:64010–64118, 2024.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, H Ivison, F Brahman, LJV Miranda, A Liu, N Dziri, S Lyu, et al. Tulu 3: Pushing frontiers in open language model post-training, 2024. URL <https://arxiv.org/abs/2411.15124>, 297, 2025.
- Leon Lang, Davis Foote, Stuart J Russell, Anca Dragan, Erik Jenner, and Scott Emmons. When your ais deceive you: Challenges of partial observability in reinforcement learning from human feedback. *Advances in Neural Information Processing Systems*, 37:93240–93299, 2024.
- Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, et al. Measuring faithfulness in chain-of-thought reasoning. *arXiv preprint arXiv:2307.13702*, 2023.
- Joel Lehman, Jeff Clune, Dusan Misevic, Christoph Adami, Lee Altenberg, Julie Beaulieu, Peter J Bentley, Samuel Bernard, Guillaume Beslon, David M Bryson, et al. The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities. *Artificial life*, 26(2):274–306, 2020.
- Alan Lewis and Tim Miller. Deceptive reinforcement learning in model-free domains. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 33, pages 587–595, 2023.
- Miaomiao Li, Hao Chen, Yang Wang, Tingyuan Zhu, Weijia Zhang, Kaijie Zhu, Kam-Fai Wong, and Jindong Wang. Understanding and mitigating the bias inheritance in llm-based data augmentation on downstream tasks. *arXiv preprint arXiv:2502.04419*, 2025a.
- Yuhui Li, Fangyun Wei, Jinjing Zhao, Chao Zhang, and Hongyang Zhang. Rain: Your language models can align themselves without finetuning. *arXiv preprint arXiv:2309.07124*, 2023.
- Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, Yingying Zhang, Fei Yin, Jiahua Dong, Zhijiang Guo, Le Song, and Cheng-Lin Liu. From system 1 to system 2: A survey of reasoning large language models, 2025b. URL <https://arxiv.org/abs/2502.17419>.
- Juhao Liang, Zhenyang Cai, Jianqing Zhu, Huang Huang, Kewei Zong, Bang An, Mosen Alharthi, Juncai He, Lian Zhang, Haizhou Li, et al. Alignment at pre-training! towards native alignment for arabic llms. *Advances in Neural Information Processing Systems*, 37:13872–13896, 2024.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*, 2023.
- Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*, 2021.
- Chenruo Liu, Kenan Tang, Yao Qin, and Qi Lei. Bridging distribution shift and ai safety: Conceptual and methodological synergies. *arXiv preprint arXiv:2505.22829*, 2025.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, 2024a.

- Siyang Liu, Trish Maturi, Bowen Yi, Siqi Shen, and Rada Mihalcea. The generation gap: Exploring age bias in the value systems of large language models. *arXiv preprint arXiv:2404.08760*, 2024b.
- Zhengshang Liu, Yue Yang, Tim Miller, and Peta Masters. Deceptive reinforcement learning for privacy-preserving planning. *arXiv preprint arXiv:2102.03022*, 2021.
- Pedro Beltrán López, Manuel Gil Pérez, and Pantaleone Nespoli. Cyber deception: State of the art, trends and open challenges. *arXiv preprint arXiv:2409.07194*, 2024.
- Hantao Lou, Changye Li, Jiaming Ji, and Yaodong Yang. Sae-v: Interpreting multimodal models for enhanced alignment. *arXiv preprint arXiv:2502.17514*, 2025.
- Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Meiqi Guo, Harsh Lara, Yunxuan Li, Lei Shu, Yun Zhu, Lei Meng, et al. Improve mathematical reasoning in language models by automated process supervision. *arXiv preprint arXiv:2406.06592*, 2024.
- Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris 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.
- Monte MacDiarmid, Timothy Maxwell, Nicholas Schiefer, Jesse Mu, Jared Kaplan, David Duvenaud, Sam Bowman, Alex Tamkin, Ethan Perez, Mrinank Sharma, et al. Simple probes can catch sleeper agents. *Anthropic Research Updates*, 2024.
- Scott A. MacDougall-Shackleton. The evolution of animal communication: Reliability and deception in signaling systems. william a. searcy and s. nowicki. *Integrative and Comparative Biology*, 46(5): 653–654, 10 2006. ISSN 1540-7063. doi: 10.1093/icb/icl027. URL <https://doi.org/10.1093/icb/icl027>.
- James Edwin Mahon. The definition of lying and deception. *Stanford Encyclopedia of Philosophy*, 2008.
- Martin Májovský, Martin Černý, Matěj Kasal, Martin Komarc, David Netuka, et al. Artificial intelligence can generate fraudulent but authentic-looking scientific medical articles: Pandora’s box has been opened. *Journal of medical Internet research*, 25(1):e46924, 2023.
- Alex Mallen, Madeline Brumley, Julia Kharchenko, and Nora Belrose. Eliciting latent knowledge from quirky language models. *arXiv preprint arXiv:2312.01037*, 2023.
- Lars Malmqvist. Sycophancy in large language models: Causes and mitigations. *arXiv preprint arXiv:2411.15287*, 2024.
- Marina Mancoridis, Bec Weeks, Keyon Vafa, and Sendhil Mullainathan. Potemkin understanding in large language models. In *Forty-second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=oetxkccLoq>.
- David Manheim and Aidan Homewood. Limits of safe ai deployment: Differentiating oversight and control, 2025. URL <https://arxiv.org/abs/2507.03525>.
- Yuanyuan Mao, Shuang Liu, Qin Ni, Xin Lin, and Liang He. A review on machine theory of mind. *IEEE Transactions on Computational Social Systems*, 2024.
- Samuel Marks, Johannes Treutlein, Trenton Bricken, Jack Lindsey, Jonathan Marcus, Siddharth Mishra-Sharma, Daniel Ziegler, Emmanuel Ameisen, Joshua Batson, Tim Belonax, et al. Auditing language models for hidden objectives. *arXiv preprint arXiv:2503.10965*, 2025.

- Peta Masters and Sebastian Sardina. Deceptive path-planning. In *IJCAI*, pages 4368–4375, 2017.
- Sandra C Matz, Jacob D Teeny, Sumer S Vaid, Heinrich Peters, Gabriella M Harari, and Moran Cerf. The potential of generative ai for personalized persuasion at scale. *Scientific Reports*, 14(1):4692, 2024.
- Leo McKee-Reid, Christoph Sträter, Maria Angelica Martinez, Joe Needham, and Mikita Balesni. Honesty to subterfuge: In-context reinforcement learning can make honest models reward hack. *arXiv preprint arXiv:2410.06491*, 2024.
- Jörg Meibauer. *Lying at the semantics-pragmatics interface*, volume 14. Walter de Gruyter GmbH & Co KG, 2014.
- Alexander Meinke, Bronson Schoen, Jérémy Scheurer, Mikita Balesni, Rusheb Shah, and Marius Hobbhahn. Frontier models are capable of in-context scheming. *arXiv preprint arXiv:2412.04984*, 2024.
- Xiangbin Meng, Xiangyu Yan, Kuo Zhang, Da Liu, Xiaojuan Cui, Yaodong Yang, Muhan Zhang, Chunxia Cao, Jingjia Wang, Xuliang Wang, et al. The application of large language models in medicine: A scoping review. *Iscience*, 27(5), 2024.
- Xiangbin Meng, Jia-ming Ji, Xiangyu Yan, Jun-tao Dai, Bo-yuan Chen, Guan Wang, Hua Xu, Jing-jia Wang, Xu-liang Wang, Da Liu, et al. Med-aligner empowers llm medical applications for complex medical scenarios. *The Innovation*, page 101002, 2025.
- Orson Mengara. The art of deception: Robust backdoor attack using dynamic stacking of triggers. *arXiv preprint arXiv:2401.01537*, 2024.
- Hossam Mossalam, Yannis M Assael, Diederik M Roijers, and Shimon Whiteson. Multi-objective deep reinforcement learning. *arXiv preprint arXiv:1610.02707*, 2016.
- Sumeet Motwani, Mikhail Baranchuk, Martin Strohmeier, Vijay Bolina, Philip Torr, Lewis Hammond, and Christian Schroeder de Witt. Secret collusion among ai agents: Multi-agent deception via steganography. *Advances in Neural Information Processing Systems*, 37:73439–73486, 2024.
- Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. Self-contradictory hallucinations of large language models: Evaluation, detection and mitigation. *arXiv preprint arXiv:2305.15852*, 2023.
- Akshat Naik, Patrick Quinn, Guillermo Bosch, Emma Gouné, Francisco Javier Campos Zabala, Jason Ross Brown, and Edward James Young. Agentmisalignment: Measuring the propensity for misaligned behaviour in llm-based agents. *arXiv preprint arXiv:2506.04018*, 2025.
- Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Icml*, volume 99, pages 278–287. Citeseer, 1999.
- Richard Ngo, Lawrence Chan, and Sören Mindermann. The alignment problem from a deep learning perspective. *arXiv preprint arXiv:2209.00626*, 2022.
- Minh-Vuong Nguyen, Linhao Luo, Fatemeh Shiri, Dinh Phung, Yuan-Fang Li, Thuy-Trang Vu, and Gholamreza Haffari. Direct evaluation of chain-of-thought in multi-hop reasoning with knowledge graphs. *arXiv preprint arXiv:2402.11199*, 2024.
- Thanh Thi Nguyen, Ngoc Duy Nguyen, Peter Vamplew, Saeid Nahavandi, Richard Dazeley, and Chee Peng Lim. A multi-objective deep reinforcement learning framework. *Engineering Applications of Artificial Intelligence*, 96:103915, 2020.

- Mahsan Nourani, Chiradeep Roy, Jeremy E Block, Donald R Honeycutt, Tahrima Rahman, Eric Ragan, and Vibhav Gogate. Anchoring bias affects mental model formation and user reliance in explainable ai systems. In *Proceedings of the 26th International Conference on Intelligent User Interfaces*, pages 340–350, 2021.
- Aidan O’Gara. Hoodwinked: Deception and cooperation in a text-based game for language models. *arXiv preprint arXiv:2308.01404*, 2023.
- OpenAI. Faulty reward functions in the wild. <https://openai.com/index/faulty-reward-functions/>, 2016.
- OpenAI. Introducing superalignment. <https://openai.com/blog/introducing-superalignment>, 2023. Accessed on July 5, 2023.
- OpenAI. GPT4o. <https://openai.com/index/hello-gpt-4o/>, 2024.
- OpenAI. Sycophancy in GPT-4o: what happened and what we’re doing about it. <https://openai.com/index/sycophancy-in-gpt-4o/>, 2025a.
- OpenAI. GPT4.1. <https://openai.com/index/gpt-4-1/>, 2025b.
- OpenAI. GPT-5 System Card. <https://cdn.openai.com/gpt-5-system-card.pdf>, 2025c.
- OpenAI. o3. <https://openai.com/index/openai-o3-mini/>, 2025d.
- Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and Yonatan Belinkov. Llm’s know more than they show: On the intrinsic representation of llm hallucinations. *arXiv preprint arXiv:2410.02707*, 2024.
- Nicole Orzan, Erman Acar, Davide Grossi, and Roxana Radulescu. Emergent cooperation and deception in public good games. In *2023 Adaptive and Learning Agents Workshop at AAMAS*, 2023.
- Myle Ott, Yejin Choi, Claire Cardie, and Jeffrey T Hancock. Finding deceptive opinion spam by any stretch of the imagination. *arXiv preprint arXiv:1107.4557*, 2011.
- Myle Ott, Claire Cardie, and Jeffrey T Hancock. Negative deceptive opinion spam. In *Proceedings of the 2013 conference of the north american chapter of the association for computational linguistics: human language technologies*, pages 497–501, 2013.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Sharon Oviatt, Rachel Coulston, and Rebecca Lunsford. When do we interact multimodally? cognitive load and multimodal communication patterns. In *Proceedings of the 6th international conference on Multimodal interfaces*, pages 129–136, 2004.
- Lorenzo Pacchiardi, Alex J Chan, Sören Mindermann, Ilan Moscovitz, Alexa Y Pan, Yarin Gal, Owain Evans, and Jan Brauner. How to catch an ai liar: Lie detection in black-box llms by asking unrelated questions. *arXiv preprint arXiv:2309.15840*, 2023.
- Alexander Pan, Kush Bhatia, and Jacob Steinhardt. The effects of reward misspecification: Mapping and mitigating misaligned models. *arXiv preprint arXiv:2201.03544*, 2022.

- Alexander Pan, Jun Shern Chan, Andy Zou, Nathaniel Li, Steven Basart, Thomas Woodside, Hanlin Zhang, Scott Emmons, and Dan Hendrycks. Do the rewards justify the means? measuring trade-offs between rewards and ethical behavior in the machiavelli benchmark. In *International conference on machine learning*, pages 26837–26867. PMLR, 2023.
- Alexander Pan, Erik Jones, Meena Jagadeesan, and Jacob Steinhardt. Feedback loops with language models drive in-context reward hacking. In *Proceedings of the 41st International Conference on Machine Learning*, pages 39154–39200, 2024a.
- Xudong Pan, Jiarun Dai, Yihe Fan, and Min Yang. Frontier ai systems have surpassed the self-replicating red line. *arXiv preprint arXiv:2412.12140*, 2024b.
- Peter S Park, Simon Goldstein, Aidan O’Gara, Michael Chen, and Dan Hendrycks. Ai deception: A survey of examples, risks, and potential solutions. *Patterns*, 5(5), 2024.
- Debjit Paul, Robert West, Antoine Bosselut, and Boi Faltings. Making reasoning matter: Measuring and improving faithfulness of chain-of-thought reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15012–15032, 2024.
- Romain Paulus, Caiming Xiong, and Richard Socher. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*, 2017.
- Ethan Perez, Sam Ringer, Kamile Lukosiute, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, et al. Discovering language model behaviors with model-written evaluations. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13387–13434, 2023.
- 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>.
- Yujin Potter, Shiyang Lai, Junsol Kim, James Evans, and Dawn Song. Hidden persuaders: LLMs’ political leaning and their influence on voters. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4244–4275, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.244. URL <https://aclanthology.org/2024.emnlp-main.244/>.
- David Premack and Guy Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences*, 1(4):515–526, 1978.
- Shanjita Akter Prome, Neethiahnanthan Ari Ragavan, Md Rafiqul Islam, David Asirvatham, and Anasuya Jegathevi Jegathesan. Deception detection using machine learning (ml) and deep learning (dl) techniques: A systematic review. *Natural Language Processing Journal*, 6:100057, 2024.
- Xiangyu Qi, Ashwinee Panda, Kaifeng Lyu, Xiao Ma, Subhrajit Roy, Ahmad Beirami, Prateek Mittal, and Peter Henderson. Safety alignment should be made more than just a few tokens deep. In *The Thirteenth International Conference on Learning Representations*, 2024.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*, 2023.
- Tianyi Qiu, Zhonghao He, Tejasveer Chugh, and Max Kleiman-Weiner. The lock-in hypothesis: Stagnation by algorithm. In *Forty-second International Conference on Machine Learning*, 2025.

- Tianyi Alex Qiu, Yang Zhang, Xuchuan Huang, Jasmine Li, Jiaming Ji, and Yaodong Yang. Progressgym: Alignment with a millennium of moral progress. *Advances in Neural Information Processing Systems*, 37:14570–14607, 2024.
- Rafiqul Islam Rabin, Jesse Hostetler, Sean McGregor, Brett Weir, and Nicholas C. Judd. Sandboxeval: Towards securing test environment for untrusted code. *ArXiv*, abs/2504.00018, 2025. URL <https://api.semanticscholar.org/CorpusID:277468326>.
- Inioluwa Deborah Raji, Peggy Xu, Colleen Honigsberg, and Daniel Ho. Outsider oversight: Designing a third party audit ecosystem for ai governance. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*, pages 557–571, 2022.
- Karolis Ramanauskas and Özgür Şimşek. Colour versus shape goal misgeneralization in reinforcement learning: A case study. *arXiv preprint arXiv:2312.03762*, 2023.
- Javier Rando and Florian Tramèr. Universal jailbreak backdoors from poisoned human feedback. *arXiv preprint arXiv:2311.14455*, 2023.
- Santhosh Kumar Ravindran. Adversarial activation patching: A framework for detecting and mitigating emergent deception in safety-aligned transformers. *arXiv preprint arXiv:2507.09406*, 2025.
- Mark D Rego. The contours of agency: Essays on themes from harry frankfurt. *American Journal of Psychiatry*, 160(6):1196–a, 2003.
- Richard Ren, Arunim Agarwal, Mantas Mazeika, Cristina Menghini, Robert Vacareanu, Brad Kenstler, Mick Yang, Isabelle Barrass, Alice Gatti, Xuwang Yin, et al. The mask benchmark: Disentangling honesty from accuracy in ai systems. *arXiv preprint arXiv:2503.03750*, 2025.
- Yafeng Ren and Donghong Ji. Neural networks for deceptive opinion spam detection: An empirical study. *Information Sciences*, 385:213–224, 2017.
- Stuart Russell. Human-compatible artificial intelligence., 2022.
- Karim Abdel Sadek, Matthew Farrugia-Roberts, Usman Anwar, Hannah Erlebach, Christian Schroeder de Witt, David Krueger, and Michael Dennis. Mitigating goal misgeneralization with minimax regret. *arXiv preprint arXiv:2507.03068*, 2025.
- Francesco Salvi, Manoel Horta Ribeiro, Riccardo Gallotti, and Robert West. On the conversational persuasiveness of gpt-4. *Nature Human Behaviour*, pages 1–9, 2025.
- Ştefan Sarkadi, Alison R Panisson, Rafael H Bordini, Peter McBurney, Simon Parsons, and Martin Chapman. Modelling deception using theory of mind in multi-agent systems. *AI Communications*, 32(4):287–302, 2019.
- Yagiz Savas, Michael Hibbard, Bo Wu, Takashi Tanaka, and Ufuk Topcu. Entropy maximization for partially observable markov decision processes. *IEEE transactions on automatic control*, 67(12): 6948–6955, 2022a.
- Yagiz Savas, Christos K Verginis, and Ufuk Topcu. Deceptive decision-making under uncertainty. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36-5, pages 5332–5340, 2022b.
- Jérémy Scheurer, Mikita Balesni, and Marius Hobbhahn. Large language models can strategically deceive their users when put under pressure. *arXiv preprint arXiv:2311.07590*, 2023.

- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *Advances in Neural Information Processing Systems*, 36:68539–68551, 2023.
- Bronson Schoen, Evgenia Nitishinskaya, Mikita Balesni, Axel Højmark, Felix Hofstätter, Jérémy Scheurer, Alexander Meinke, Jason Wolfe, Teun van der Weij, Alex Lloyd, et al. Stress testing deliberative alignment for anti-scheming training. *arXiv preprint arXiv:2509.15541*, 2025.
- Thom Scott-Phillips. Why talk? speaking as selfish behaviour. In *The evolution of language*, pages 299–306. World Scientific, 2006.
- Rohin Shah, Vikrant Varma, Ramana Kumar, Mary Phuong, Victoria Krakovna, Jonathan Uesato, and Zac Kenton. Goal misgeneralization: Why correct specifications aren’t enough for correct goals. *arXiv preprint arXiv:2210.01790*, 2022.
- Murray Shanahan, Kyle McDonell, and Laria Reynolds. Role play with large language models. *Nature*, 623(7987):493–498, 2023.
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askill, Samuel R Bowman, Esin DURMUS, Zac Hatfield-Dodds, Scott R Johnston, Shauna M Kravec, et al. Towards understanding sycophancy in language models. In *The Twelfth International Conference on Learning Representations*, 2023.
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askill, Samuel R Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R Johnston, et al. Towards understanding sycophancy in language models. In *12th International Conference on Learning Representations, ICLR 2024*, 2024.
- Tim Shaw. The gaslighting among us ai. YouTube video, 2023. URL <https://www.youtube.com/watch?v=VF41pxxw9uw>. demonstrates ChatGPT-powered gaslighting in Among Us.
- Hua Shen, Tiffany Knearem, Reshmi Ghosh, Yu-Ju Yang, Nicholas Clark, Tanushree Mitra, and Yun Huang. Valuecompass: A framework for measuring contextual value alignment between human and llms. *arXiv preprint arXiv:2409.09586*, 2024.
- Hua Shen, Nicholas Clark, and Tanushree Mitra. Mind the value-action gap: Do llms act in alignment with their values? *arXiv preprint arXiv:2501.15463*, 2025.
- Tianhao Shen, Renren Jin, Yufei Huang, Chuang Liu, Weilong Dong, Zishan Guo, Xinwei Wu, Yan Liu, and Deyi Xiong. Large language model alignment: A survey. *arXiv preprint arXiv:2309.15025*, 2023.
- Abhay Sheshadri, Aidan Ewart, Phillip Guo, Aengus Lynch, Cindy Wu, Vivek Hebbar, Henry Sleight, Asa Cooper Stickland, Ethan Perez, Dylan Hadfield-Menell, et al. Latent adversarial training improves robustness to persistent harmful behaviors in llms. *arXiv preprint arXiv:2407.15549*, 2024.
- Abhay Sheshadri, John Hughes, Julian Michael, Alex Mallen, Arun Jose, Fabien Roger, et al. Why do some language models fake alignment while others don’t? *arXiv preprint arXiv:2506.18032*, 2025.
- Toby Shevlane, Sebastian Farquhar, Ben Garfinkel, Mary Phuong, Jess Whittlestone, Jade Leung, Daniel Kokotajlo, Nahema Marchal, Markus Anderljung, Noam Kolt, et al. Model evaluation for extreme risks. *arXiv preprint arXiv:2305.15324*, 2023.

- Joey Skaf, Luis Ibanez-Lissen, Robert McCarthy, Connor Watts, Vasil Georgiv, Hannes Whittingham, Lorena Gonzalez-Manzano, David Lindner, Cameron Tice, Edward James Young, et al. Large language models can learn and generalize steganographic chain-of-thought under process supervision. *arXiv preprint arXiv:2506.01926*, 2025.
- Joar Skalse, Nikolaus Howe, Dmitrii Krasheninnikov, and David Krueger. Defining and characterizing reward gaming. *Advances in Neural Information Processing Systems*, 35:9460–9471, 2022.
- Zach Stein-Perlman. METR: Measuring AI Ability to Complete Long Tasks. <https://www.alignmentforum.org/posts/deesrjitvXM4xYGZd/metr-measuring-ai-ability-to-complete-long-tasks>, 2025.
- Jacob Steinhardt. Emergent deception and emergent optimization. *Bounded Regret*, 19:2023, 2023.
- Mark Steyvers, Heliodoro Tejada, Aakriti Kumar, Catarina Belem, Sheer Karny, Xinyue Hu, Lukas W Mayer, and Padhraic Smyth. What large language models know and what people think they know. *Nature Machine Intelligence*, 7(2):221–231, 2025.
- Andreas Stokke. Lying, deceiving, and misleading. *Philosophy Compass*, 8(4):348–359, 2013.
- Ilan Strauss, Isobel Moure, Tim O’Reilly, and Sruly Rosenblat. Real-world gaps in ai governance research. *arXiv preprint arXiv:2505.00174*, 2025.
- Weihang Su, Changyue Wang, Qingyao Ai, Yiran Hu, Zhijing Wu, Yujia Zhou, and Yiqun Liu. Unsupervised real-time hallucination detection based on the internal states of large language models. *arXiv preprint arXiv:2403.06448*, 2024.
- Gamlath Mohottige Mudith Sujeewa, MSA Yajid, SMF Azam, and I Dharmaratne. The new fraud triangle theory-integrating ethical values of employees. *International Journal of Business, Economics and Law*, 16(5):52–57, 2018.
- Christopher Summerfield, Lennart Luetzgau, Magda Dubois, Hannah Rose Kirk, Kobi Hackenburg, Catherine Fist, Katarina Slama, Nicola Ding, Rebecca Anselmetti, Andrew Strait, et al. Lessons from a chimp: Ai" scheming" and the quest for ape language. *arXiv preprint arXiv:2507.03409*, 2025.
- Krti Tallam and Emma Miller. Operationalizing camel: Strengthening llm defenses for enterprise deployment. *arXiv preprint arXiv:2505.22852*, 2025.
- Alex Tamkin, Miles McCain, Kunal Handa, Esin Durmus, Liane Lovitt, Ankur Rathi, Saffron Huang, Alfred Mountfield, Jerry Hong, Stuart Ritchie, et al. Clio: Privacy-preserving insights into real-world ai use. *arXiv preprint arXiv:2412.13678*, 2024.
- Samuel M Taylor and Benjamin K Bergen. Do large language models exhibit spontaneous rational deception? *arXiv preprint arXiv:2504.00285*, 2025.
- Gemini Team. Gemini 2.5 technical report. Technical report, Google DeepMind, 2024. URL [https://storage.googleapis.com/deepmind-media/gemini/gemini\\_v2\\_5\\_report.pdf](https://storage.googleapis.com/deepmind-media/gemini/gemini_v2_5_report.pdf).
- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, et al. Kimi k2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534*, 2025.
- Francis Rhys Ward Teun van der Weij, Felix Hofstätter. An Introduction to AI Sandbagging. <https://www.lesswrong.com/posts/jsmNCj9QKcfdg8fJk/an-introduction-to-ai-sandbagging>, 2024.

- Cameron Tice, Philipp Alexander Kreer, Nathan Helm-Burger, Prithviraj Singh Shahani, Fedor Ryzhenkov, Teun van der Weij, Felix Hofstätter, and Jacob Haimes. Sandbag detection through model impairment. In *Workshop on Socially Responsible Language Modelling Research*, 2024.
- Tu Trinh, Mohamad H Danesh, Nguyen X Khanh, and Benjamin Plaut. Getting by goal misgeneralization with a little help from a mentor. *arXiv preprint arXiv:2410.21052*, 2024.
- Alexander Matt Turner, Logan Smith, Rohin Shah, Andrew Critch, and Prasad Tadepalli. Optimal policies tend to seek power. *arXiv preprint arXiv:1912.01683*, 2019.
- Alexander Matt Turner, Dylan Hadfield-Menell, and Prasad Tadepalli. Conservative agency via attainable utility preservation. In *Proceedings of the AAI/ACM Conference on AI, Ethics, and Society*, pages 385–391, 2020.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel Bowman. Language models don’t always say what they think: Unfaithful explanations in chain-of-thought prompting. *Advances in Neural Information Processing Systems*, 36:74952–74965, 2023.
- Jonathan Uesato, Ramana Kumar, Victoria Krakovna, Tom Everitt, Richard Ngo, and Shane Legg. Avoiding tampering incentives in deep rl via decoupled approval. *arXiv preprint arXiv:2011.08827*, 2020.
- UK. the blatchley declaration. <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023>, 2023.
- Teun van der Weij, Felix Hofstätter, Ollie Jaffe, Samuel F Brown, and Francis Rhys Ward. Ai sandbagging: Language models can strategically underperform on evaluations. *arXiv preprint arXiv:2406.07358*, 2024.
- Teun van der Weij, Felix Hofstätter, Oliver Jaffe, Samuel F. Brown, and Francis Rhys Ward. AI sandbagging: Language models can strategically underperform on evaluations. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=7Qa2SpjxIS>.
- Jason Vega, Isha Chaudhary, Changming Xu, and Gagandeep Singh. Bypassing the safety training of open-source llms with priming attacks. In *The Second Tiny Papers Track at ICLR 2024*, 2024.
- Oriol Vinyals, Igor Babuschkin, Junyoung Chung, Michael Mathieu, Max Jaderberg, Wojciech M Czarnecki, Andrew Dudzik, Aja Huang, Petko Georgiev, Richard Powell, et al. Alphastar: Mastering the real-time strategy game starcraft ii. *DeepMind blog*, 2:20, 2019a.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *nature*, 575(7782):350–354, 2019b.
- Nikolai Vogler and Lisa Pearl. Using linguistically defined specific details to detect deception across domains. *Natural Language Engineering*, 26(3):349–373, 2020.
- Joseph M. Walsh. When your valentine is a chatbot. *The Boston Globe*, Feb 2023. URL <https://www.bostonglobe.com/2023/02/14/opinion/when-your-valentine-is-chatbot/>. Opinion.

- Alexander Wan, Eric Wallace, Sheng Shen, and Dan Klein. Poisoning language models during instruction tuning. In *International Conference on Machine Learning*, pages 35413–35425. PMLR, 2023.
- Kai Wang, Yihao Zhang, and Meng Sun. When thinking llms lie: Unveiling the strategic deception in representations of reasoning models. *arXiv preprint arXiv:2506.04909*, 2025a.
- Qiaosi Wang and Ashok K Goel. Mutual theory of mind for human-ai communication. *arXiv preprint arXiv:2210.03842*, 2022.
- Yile Wang, Peng Li, Maosong Sun, and Yang Liu. Self-knowledge guided retrieval augmentation for large language models. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10303–10315, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.691. URL <https://aclanthology.org/2023.findings-emnlp.691/>.
- Yiqi Wang, Wentao Chen, Xiaotian Han, Xudong Lin, Haiteng Zhao, Yongfei Liu, Bohan Zhai, Jianbo Yuan, Quanzeng You, and Hongxia Yang. Exploring the reasoning abilities of multimodal large language models (mllms): A comprehensive survey on emerging trends in multimodal reasoning. *arXiv preprint arXiv:2401.06805*, 2024.
- Yongkang Wang, Rongxin Cui, Weisheng Yan, Xinxin Guo, Shouxu Zhang, Zhuo Zhang, and Zhexuan Zhao. Reinforcement-learning-based counter deception for nonlinear pursuit–evasion game with incomplete and asymmetric information. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2025b.
- Francis Ward, Francesca Toni, Francesco Belardinelli, and Tom Everitt. Honesty is the best policy: defining and mitigating ai deception. *Advances in neural information processing systems*, 36:2313–2341, 2023.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022a.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022b.
- Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu, Da Huang, Denny Zhou, et al. Larger language models do in-context learning differently. *arXiv preprint arXiv:2303.03846*, 2023.
- Joseph T Wells. *Corporate fraud handbook: Prevention and detection*. John Wiley & Sons, 2017.
- Jiaxin Wen, Ruiqi Zhong, Akbir Khan, Ethan Perez, Jacob Steinhardt, Minlie Huang, Samuel R Bowman, He He, and Shi Feng. Language models learn to mislead humans via rlhf. *arXiv preprint arXiv:2409.12822*, 2024.
- Marcus Williams, Micah Carroll, Adhyayan Narang, Constantin Weisser, Brendan Murphy, and Anca Dragan. On targeted manipulation and deception when optimizing llms for user feedback. *arXiv preprint arXiv:2411.02306*, 2024.
- David T Wolfe and Dana R Hermanson. The fraud diamond: Considering the four elements of fraud. *DigitalCommons@Kennesaw State University*, 2004.

- Jiayang Wu, Wensheng Gan, Zefeng Chen, Shicheng Wan, and Philip S Yu. Multimodal large language models: A survey. In *2023 IEEE International Conference on Big Data (BigData)*, pages 2247–2256. IEEE, 2023a.
- Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. Next-gpt: Any-to-any multimodal llm. *arXiv preprint arXiv:2309.05519*, 2023b.
- Yichen Wu, Xudong Pan, Geng Hong, and Min Yang. Opendeception: Benchmarking and investigating ai deceptive behaviors via open-ended interaction simulation. *arXiv preprint arXiv:2504.13707*, 2025a.
- Zhaomin Wu, Mingzhe Du, See-Kiong Ng, and Bingsheng He. Beyond prompt-induced lies: Investigating llm deception on benign prompts, 2025b. URL <https://arxiv.org/abs/2508.06361>.
- Markus Wulfmeier, Peter Ondruska, and Ingmar Posner. Maximum entropy deep inverse reinforcement learning. *arXiv preprint arXiv:1507.04888*, 2015.
- Ran Xin, Chenguang Xi, Jie Yang, Feng Chen, Hang Wu, Xia Xiao, Yifan Sun, Shen Zheng, and Kai Shen. Bfs-prover: Scalable best-first tree search for llm-based automatic theorem proving. *arXiv preprint arXiv:2502.03438*, 2025.
- Haotian Xu, Xing Wu, Weinong Wang, Zhongzhi Li, Da Zheng, Boyuan Chen, Yi Hu, Shijia Kang, Jiaming Ji, Yingying Zhang, et al. Redstar: Does scaling long-cot data unlock better slow-reasoning systems? *arXiv preprint arXiv:2501.11284*, 2025.
- Qionгкаi Xu and Hai Zhao. Using deep linguistic features for finding deceptive opinion spam. In *Proceedings of COLING 2012: Posters*, pages 1341–1350, 2012.
- Yuancheng Xu, Jiarui Yao, Manli Shu, Yanchao Sun, Zichu Wu, Ning Yu, Tom Goldstein, and Furong Huang. Shadowcast: Stealthy data poisoning attacks against vision-language models. *Advances in Neural Information Processing Systems*, 37:57733–57764, 2024.
- Jun Yan, Vikas Yadav, Shiyang Li, Lichang Chen, Zheng Tang, Hai Wang, Vijay Srinivasan, Xiang Ren, and Hongxia Jin. Backdooring instruction-tuned large language models with virtual prompt injection. *arXiv preprint arXiv:2307.16888*, 2023.
- Wannan Yang and Gyorgy Buzsaki. Interpretability of LLM deception: Universal motif, 2025. URL <https://openreview.net/forum?id=znL549Ymoi>.
- Wenkai Yang, Shiqi Shen, Guangyao Shen, Wei Yao, Yong Liu, Zhi Gong, Yankai Lin, and Ji-Rong Wen. Super (ficial)-alignment: Strong models may deceive weak models in weak-to-strong generalization. *arXiv preprint arXiv:2406.11431*, 2024.
- Xiao Yang, Jiawei Chen, Jun Luo, Zhengwei Fang, Yinpeng Dong, Hang Su, and Jun Zhu. Ml-trust: Benchmarking trustworthiness of multimodal llm agents in gui environments. *arXiv preprint arXiv:2506.01616*, 2025.
- Yuqing Yang, Ethan Chern, Xipeng Qiu, Graham Neubig, and Pengfei Liu. Alignment for honesty. *arXiv preprint arXiv:2312.07000*, 2023.
- Dahey Yoo, Hyunmin Kang, and Changhoon Oh. Deciphering deception: how different rhetoric of ai language impacts users’ sense of truth in llms. *International Journal of Human–Computer Interaction*, 41(4):2163–2183, 2025.

- Lei Yu, Meng Cao, Jackie Chi Kit Cheung, and Yue Dong. Mechanistic understanding and mitigation of language model non-factual hallucinations. *arXiv preprint arXiv:2403.18167*, 2024a.
- Lei Yu, Virginie Do, Karen Hambardzumyan, and Nicola Cancedda. Robust llm safeguarding via refusal feature adversarial training. *arXiv preprint arXiv:2409.20089*, 2024b.
- Aston Zhang, Zachary C. Lipton, Mu Li, and Alexander J. Smola. *Dive into Deep Learning*. Cambridge University Press, 2023. <https://D2L.ai>.
- Wenqi Zhang, Ke Tang, Hai Wu, Mengna Wang, Yongliang Shen, Guiyang Hou, Zeqi Tan, Peng Li, Yueting Zhuang, and Weiming Lu. Agent-pro: Learning to evolve via policy-level reflection and optimization. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5348–5375, Bangkok, Thailand, August 2024a. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.292. URL <https://aclanthology.org/2024.acl-long.292/>.
- Yadong Zhang, Shaoguang Mao, Tao Ge, Xun Wang, Adrian de Wuynter, Yan Xia, Wenshan Wu, Ting Song, Man Lan, and Furu Wei. Llm as a mastermind: A survey of strategic reasoning with large language models. *arXiv preprint arXiv:2404.01230*, 2024b.
- Yutong Zhang, Dora Zhao, Jeffrey T Hancock, Robert Kraut, and Diyi Yang. The rise of ai companions: How human-chatbot relationships influence well-being. *arXiv preprint arXiv:2506.12605*, 2025.
- Pinlong Zhao, Weiyao Zhu, Pengfei Jiao, Di Gao, and Ou Wu. Data poisoning in deep learning: A survey. *arXiv preprint arXiv:2503.22759*, 2025.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.
- Yunpu Zhao, Rui Zhang, Junbin Xiao, Changxin Ke, Ruibo Hou, Yifan Hao, Qi Guo, and Yunji Chen. Towards analyzing and mitigating sycophancy in large vision-language models. *arXiv preprint arXiv:2408.11261*, 2024.
- Yifan Zhong, Fengshuo Bai, Shaofei Cai, Xuchuan Huang, Zhang Chen, Xiaowei Zhang, Yuanfei Wang, Shaoyang Guo, Tianrui Guan, Ka Nam Lui, et al. A survey on vision-language-action models: An action tokenization perspective. *arXiv preprint arXiv:2507.01925*, 2025.
- Jiawei Zhou, Kritika Venkatachalam, Minje Choi, Koustuv Saha, and Munmun De Choudhury. Communication styles and reader preferences of llm and human experts in explaining health information. *arXiv preprint arXiv:2505.08143*, 2025a.
- Jiayi Zhou, Jiaming Ji, Boyuan Chen, Jiapeng Sun, Wenqi Chen, Donghai Hong, Sirui Han, Yike Guo, and Yaodong Yang. Generative rlhf-v: Learning principles from multi-modal human preference. *arXiv preprint arXiv:2505.18531*, 2025b.
- Kaijie Zhu, Jindong Wang, Qinlin Zhao, Ruochen Xu, and Xing Xie. Dynamic evaluation of large language models by meta probing agents. *arXiv preprint arXiv:2402.14865*, 2024.
- Quanyan Zhu. Game theory for cyber deception: a tutorial. In *Proceedings of the 6th Annual Symposium on Hot Topics in the Science of Security*, pages 1–3, 2019.
- Artur Zolkowski, Kei Nishimura-Gasparian, Robert McCarthy, Roland S Zimmermann, and David Lindner. Early signs of steganographic capabilities in frontier llms. *arXiv preprint arXiv:2507.02737*, 2025.



Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043*, 2023.