



On monitorability of AI

Roman V. Yampolskiy¹

Received: 28 November 2023 / Accepted: 5 January 2024
© The Author(s) 2024

Abstract

Artificially intelligent (AI) systems have ushered in a transformative era across various domains, yet their inherent traits of unpredictability, unexplainability, and uncontrollability have given rise to concerns surrounding AI safety. This paper aims to demonstrate the infeasibility of accurately monitoring advanced AI systems to predict the emergence of certain capabilities prior to their manifestation. Through an analysis of the intricacies of AI systems, the boundaries of human comprehension, and the elusive nature of emergent behaviors, we argue for the impossibility of reliably foreseeing some capabilities. By investigating these impossibility results, we shed light on their potential implications for AI safety research and propose potential strategies to overcome these limitations.

Keywords AI audit · AI safety · Impossibility · Monitoring · Observability · Unmonitorability

1 Introduction

AI systems have evolved from simple rule-based systems to highly complex neural networks. These advanced systems, such as deep learning and reinforcement learning models, can perform tasks that are difficult or impossible for humans. The growing prevalence of advanced AI systems/foundational models has raised concerns about their safety and the potential risks associated with their deployment. Yampolskiy's research on the impossibility results in AI safety [51, 110, 118], including unpredictability [113], unexplainability [112], and uncontrollability [104], highlights the challenges in achieving safe AI [4, 114, 115, 117]. In this paper (pre-print [106], we build on Yampolskiy's work and introduce the concept of unmonitorability of AI, arguing that it is impossible to monitor advanced AI systems to correctly predict some capabilities.

Human understanding is inherently limited when dealing with complex AI systems [87]. As these systems become more advanced, they can generate solutions and behaviors that are beyond human comprehension. This limitation is exacerbated by the “black-box” nature of AI models, where the inner workings of the system are hidden from view.

Consequently, it becomes impossible for humans to anticipate the full range of possible behaviors and potential unsafe impacts before they occur.

Emergent behaviors in AI systems result from the interaction of individual components, leading to specific outcomes that are difficult to predict or explain even if general trends of increased capacity are predictable via scaling laws [63], with predictability further reduced because of inverse scaling of some capabilities [72]. These behaviors can arise from the system's internal dynamics or from its interactions with the environment. The unpredictability of emergent behaviors makes it impossible to monitor AI systems for safety accurately, though development of specific capabilities may be possible to predict [24] and quantify [88]. Even when individual components are understood and deemed safe, their interactions can still result in unforeseen consequences.

Similar results in the context of robot behavior have been proven by Leeuwen and Wiedermann, who in attempting to answer [93]: “Can an observer always tell from inspecting and monitoring a robot's program whether the robot will always obey the given rules of law or ethics, or any other set of formally expressed constraints, in any interaction with other robots (or humans)?” arrive at theoretical and practical impossibility results.

✉ Roman V. Yampolskiy
roman.yampolskiy@louisville.edu

¹ Computer Science and Engineering, University of Louisville, Louisville, KY, USA

1.1 Monitorability—definition

Calls for research on monitoring advanced AI systems have appeared in the literature [44, 45, 98], as well as broader suggestions for study of machine behavior [71, 101], [74], and problematic machine behavior [14]. This line of research is particularly important in the current AI research environment, in which old paradigm of design/engineering of AI systems with predictable capabilities has been, at least partially, replaced by evolving/self-learning deep neural networks and subsequent experimentation on produced models to discover their capabilities and limitations making Computer Science truly an experimental science not just a software engineering discipline [54].

We will begin by formalizing the notion of Monitorability [1] of AI and some relevant concepts. *Monitorability in AI* systems refers to the capacity to observe, understand, and predict the behavior and outputs of an artificial intelligence model in order to identify advanced capabilities and potentially unsafe impacts and intervene before they occur. Ortega et al. write [70]: “Monitoring comprises all the methods for inspecting systems in order to analyze and predict their behavior, both via human inspection (of summary statistics) and automated inspection (to sweep through vast amounts of activity records).” It can be seen as a new sub-field of research for the domains of Complex Systems, AI development, AI Forensics [13, 81] and AI Safety, and which can probably inherit some wisdom from the realms of electronic surveillance [61], and nuclear weapons monitoring [31].

A formal definition for monitorability in AI systems can be stated as follows: Given an AI system A , a set of input states I , a set of output states O , and a set of safety criteria C , monitorability $M(A)$ is the ability to accurately predict potential advanced capabilities $U \subseteq O$, given any input state $i \in I$, such that: $M(A): I \rightarrow P(U)$, where $P(U)$ denotes the power set of U , i.e., the set of all subsets of U . The relevance of $P(U)$ in the definition lies in its ability to capture the full range of possible combinations of advanced capabilities that can result from the AI system's operation. By mapping input states I to the power set of U ($P(U)$), the definition aims to account for all possible scenarios where one or more advanced capabilities might occur simultaneously or not occur at all. Including $P(U)$ in the definition helps to emphasize the complexity and uncertainty involved in monitoring advanced AI systems. It illustrates the challenge of accurately predicting the AI system's behavior, as there can be numerous potential advanced capability combinations for a given input state. In the context of unmonitorability, the power set $P(U)$ highlights the difficulty in anticipating and monitoring all possible advanced capabilities of an AI system.

This complexity contributes to the argument that it is impossible to perfectly monitor advanced AI systems to predict advanced capabilities before they occur.

Monitorability of an AI system is considered high if, for all input states i , the prediction of potential advanced capabilities U can be made with a high degree of accuracy and confidence. Conversely, monitorability is considered low if the AI system's behavior is difficult to predict, understand, or control, leading to an inability to accurately anticipate potential advanced capabilities. In the context of unmonitorability, we argue that for advanced AI systems, it is impossible to achieve high monitorability due to their inherent complexity, limitations of human understanding, and the emergence of unpredictable behaviors. This is true even if an AI Monitor is not human, but could be formalized as any agent including AIs of different capabilities.

The argument for unmonitorability of advanced AI systems could be made based on dependence of monitorability capability on other impossibility results, such as unexplainability, unpredictability and incomprehensibility of AI [112], but in this paper we will present a number of independent arguments for Unmonitorability of AI.

1.2 Types of monitoring

In the context of AI safety, monitoring can be classified into several types, each focusing on different aspects of AI system behavior and performance. A proposed taxonomy for AI safety monitoring follows:

1. Functional monitoring

Functional monitoring refers to tracking the AI system's performance in terms of its intended tasks and objectives. This type of monitoring is crucial for evaluating the system's efficacy and ensuring that it meets its functional requirements [21]. Examples of functional monitoring include:

- (a) Accuracy monitoring: Evaluating the AI system's ability to produce correct outputs or predictions.
- (b) Efficiency monitoring: Assessing the system's resource utilization, such as processing time and memory consumption.
- (c) Reliability monitoring: Examining the AI system's consistency and stability over time and in varying conditions.

2. Safety monitoring

Safety monitoring focuses on identifying and mitigating potential risks associated with the AI system's operation. This type of monitoring is essential for preventing harm to users, other systems, or the environment. Examples of safety monitoring include:

- (a) Security monitoring: Detecting and preventing potential vulnerabilities, such as unauthorized access or data breaches.
- (b) Robustness monitoring: Ensuring that the AI system can handle unexpected inputs, adversarial attacks, or changes in the environment.
- (c) Compliance monitoring: Verifying that the AI system adheres to established safety standards, ethical guidelines, and legal regulations.

3. Ethical and social monitoring

Ethical and social monitoring involves examining the AI system's impact on individuals, communities, and society as a whole. This type of monitoring is crucial for addressing potential biases, inequalities, and other unintended consequences of AI deployment. Examples of ethical and social monitoring include:

- (a) Fairness monitoring: Assessing whether the AI system treats different user groups equitably and does not perpetuate or exacerbate existing biases.
- (b) Transparency monitoring: Ensuring that the AI system's decision-making processes can be understood and explained to stakeholders, including users and regulators.
- (c) Privacy monitoring: Safeguarding user data and ensuring that the AI system respects individuals' privacy rights.

4. Environmental monitoring

Environmental monitoring focuses on the AI system's impact on the natural environment and its resource consumption. This type of monitoring is essential for promoting sustainable AI development and mitigating potential environmental harms. Examples of environmental monitoring include:

- (a) Energy consumption monitoring: Assessing the AI system's energy usage and identifying opportunities for optimization.
- (b) Carbon footprint monitoring: Evaluating the AI system's greenhouse gas emissions and implementing strategies to reduce its environmental impact.
- (c) Ecosystem impact monitoring: Examining the AI system's effects on ecosystems, such as habitat disruption or biodiversity loss, and devising mitigation measures.

5. Temporal monitoring

- (a) Slow monitoring: Periodic evaluation of the AI system's behavior and performance over an extended timeframe.

- (b) Live monitoring: Real-time tracking and assessment of the AI system's operation to detect and address potential issues immediately.

6. Monitoring failure modes

- (a) Monitoring fails: Identifying and addressing situations in which the monitoring process itself fails or produces erroneous results.
- (b) Monitoring for fire alarms of danger: Detecting early warning signs [127] or "fire alarms" that indicate potential safety hazards or critical failures, if possible [122].

7. Meta-monitoring

- (a) Monitoring meta information on research [124]: Assessing the quality, relevance, and potential risks associated with AI research and publications.
- (b) Regular vs pivotal capabilities capability detection delay: Evaluating the AI system's progression from regular to pivotal capabilities and the delay in detecting these advancements [57].

8. Monitoring across AI lifecycle

- (a) Monitoring during training: Tracking and assessing the AI system's performance, safety, and ethical considerations during the training phase.
- (b) Monitoring during testing: Evaluating the AI system's behavior and performance in controlled testing environments before deployment.
- (c) Monitoring during deployment: Continuously monitoring the AI system's operation and impact in real-world settings after deployment.

9. Decision monitoring

- (a) Predicting capabilities vs predicting decisions: Differentiating between monitoring the AI system's potential abilities and its actual decision-making processes.
- (b) Monitoring who is in control: Assessing the roles and responsibilities of various stakeholders [26], including the owner [117], user, and designer, in controlling the AI system's behavior.

10. Monitoring methods

- (a) Passive monitoring: Observing and recording the AI system's behavior without direct intervention or manipulation.

- (b) Active monitoring: Intervening or interacting with the AI system to obtain more information or influence its behavior.

11. Inverse monitoring

- (a) Inverse Monitoring by Software: Using AI or other software tools to monitor and assess the behavior of human users or operators.

12. Self-monitoring

- (a) Self-monitoring: Enabling AI systems to monitor their own behavior, performance, and safety, and to self-correct or adapt as needed

13. Monitoring the field of AI research

- (a) Unmonitorability of AI research as a field: Assessing the challenges in monitoring the overall AI research landscape due to undisclosed experiments, rapid advancements, and the sheer volume of publications.

This taxonomy provides a comprehensive framework for AI safety monitoring, considering various aspects of AI systems, their lifecycle, stakeholders, and methods. This holistic approach aims to address potential risks and impacts across multiple dimensions and helps to ensure safer AI development and deployment.

Whittlestone and Clark suggested a number of monitoring project which may be beneficial in the context of AI governance [98]:

- “Assessing the landscape of AI datasets and evaluating who they do and do not represent. Using these findings to fund the creation of datasets to fill in the gaps.
- Using geographic bibliometric analysis to understand a country’s competitiveness on key areas of AI research and development.
- Hosting competitions to make it easy to measure progress in a certain policy-relevant AI domain, such as competitions to find vulnerabilities in widely deployed vision systems, or to evaluate the advancing capabilities of smart industrial robots.
- Funding projects to improve assessment methods in commercially important areas (e.g., certain types of computer vision, to accelerate progress and commercial application in these areas.
- Tracking the deployment of AI systems for particular economically relevant tasks, to better track, forecast, and ultimately prepare for the societal impacts of such systems.

- Monitoring concrete cases of harm caused by AI systems on a national level, to keep policymakers up to date on the current impacts of AI, as well as potential future impacts caused by research advances.
- Monitoring the adoption of or spending on AI technology across sectors, to identify the most important sectors to track and govern, as well as generalizable insights about how to leverage AI technology in other sectors.
- Monitoring the share of key inputs to AI progress that different actors control (i.e., talent, computational resources and the means to produce them, and the relevant data), to better understand which actors policymakers will need to regulate and where intervention points are.”

2 Monitoring of AI treaties

Monitoring capabilities play a crucial role in verifying AI treaties and upholding AI governance frameworks [84]. As AI technologies continue to advance, it becomes increasingly important to establish international agreements that regulate AIs development, deployment, and use. Ensuring compliance with these treaties requires effective monitoring mechanisms that can detect and deter potential violations [41]. Here are some aspects of monitoring capabilities that are essential for verifying AI treaties as part of AI governance:

Technological monitoring: Implementing advanced technological tools and methods to monitor AI systems, including their design, operation, and performance, is vital for verifying compliance with AI treaties. This may involve the use of AI-based monitoring tools, data analysis techniques, and other technologies to assess the behavior, capabilities, and impact of AI systems on a granular level.

Transparency and information sharing: Promoting transparency and information sharing among nations and organizations is critical for effective AI treaty verification. By openly sharing details about AI research, development, and deployment, parties can collectively assess compliance with treaty obligations and maintain a collaborative approach to AI governance.

Inspection and audit mechanisms: Establishing robust inspection and audit mechanisms [36, 64, 75], both on-site and remote, is essential for verifying AI treaty compliance. Regular inspections can help ensure that AI systems are developed and deployed in accordance with treaty guidelines and that any deviations or violations are promptly detected and addressed.

International collaboration: Encouraging international collaboration among countries, research institutions, and industry stakeholders can enhance monitoring capabilities and facilitate the verification of AI treaty compliance. Joint research initiatives, shared databases, and cooperative efforts

to develop monitoring technologies can improve the overall effectiveness of AI governance frameworks.

Capacity building and training: Building the capacity of AI governance stakeholders, including policymakers, regulators, and AI developers, to effectively monitor and assess AI systems is crucial for verifying AI treaty compliance. This may involve training programs, workshops, and other initiatives to enhance understanding of AI technologies and their potential risks, as well as the development of specialized skills and expertise for AI monitoring and verification.

Legal and regulatory frameworks: Developing clear legal and regulatory frameworks that define the obligations and responsibilities of AI developers, users, and other stakeholders is essential for ensuring compliance with AI treaties [43]. These frameworks should outline the monitoring and reporting requirements, as well as the enforcement mechanisms and penalties for violations.

Continuous adaptation and improvement: AI technologies are rapidly evolving, and monitoring capabilities must be adaptable and responsive to these changes. Regular reviews and updates of monitoring techniques, tools, and methodologies are necessary to ensure that AI treaty verification remains effective in the face of emerging AI developments and challenges.

By implementing these monitoring capabilities, nations and organizations can more effectively verify AI treaty compliance, ensuring that AI systems are developed and deployed in a manner that aligns with international agreements and governance frameworks. This will help promote responsible AI development and mitigate potential risks associated with the widespread adoption of advanced AI technologies.

2.1 AI observatories

The concept of an AI observatory [3] refers to the establishment of an integrated, centralized platform dedicated to the ongoing monitoring, analysis, and evaluation of advanced AI systems' behavior and their potential societal impacts. This idea has emerged in response to the growing need for enhanced scrutiny of AI systems, particularly considering the challenges posed by computational irreducibility and the inherent limits of human understanding. Establishing an AI observatory is one approach to mitigating the monitorability issues that stem from these complexities and enhancing our capacity to ensure AI safety and alignment with human values.

In the context of monitorability, an AI observatory can play a pivotal role by aggregating data, insights, and knowledge from various sources, including AI developers, researchers, and policymakers. By fostering a collaborative environment, the observatory can facilitate the exchange

of ideas, methodologies, and best practices, enabling stakeholders to leverage collective intelligence [126] in addressing the multifaceted challenges of monitoring advanced AI systems.

Moreover, an AI observatory can contribute to the development and refinement of innovative monitoring tools and techniques, informed by the recognition of the limitations of human intuition. By harnessing cutting-edge research and technological advancements, the observatory can help create monitoring frameworks that are more robust, adaptive, and capable of detecting emergent AI behaviors and potential risks.

The AI observatory can serve as an essential platform for fostering transparency and accountability in AI development and deployment. By centralizing information on AI system performance, safety measures, and ethical considerations, the observatory can help ensure that AI developers and operators adhere to established guidelines, regulations, and best practices. This transparency can, in turn, contribute to the mitigation of risks associated with AI misalignment or unforeseen consequences.

3 Why monitoring of advanced AI is impossible

Monitoring advanced AI systems to accurately predict unsafe impacts before they happen is likely to be impossible due to several reasons:

3.1 Humans-in-the-loop

Keeping humans in the loop for monitoring advanced AI systems presents several challenges [33, 79], primarily stemming from the limitations of human cognition, reaction times, and the increasing complexity of AI systems. One major issue with human-in-the-loop monitoring is that humans may not be able to keep up with the speed and complexity of AI systems, particularly as they continue to advance and outpace human capabilities. This could render the concept of supervised autonomy less effective, as the human supervisor may struggle to understand, assess, and intervene in the AI system's actions in a timely and meaningful manner, because of limited observability [46], [47], [55], [59].

Given human reaction times, there may not be a slow enough takeoff for AI systems that allows for effective human monitoring. Human reaction times, such as auditory reaction time and visual reaction time, are considerably slower than the response times of AI systems. This disparity in processing speed makes it difficult for humans to effectively monitor and react to AI behavior in real-time, particularly in situations where rapid decision-making and intervention are necessary.

There are some examples of human-in-the-loop algorithms that attempt to bridge the gap between human and AI capabilities, such as semi-autonomous driving systems, collaborative robots, and AI-assisted decision-making tools in various domains. These systems aim to combine the strengths of both humans and AI systems to achieve better performance and maintain human oversight. However, as AI systems continue to advance and their complexity grows, the effectiveness of these human-in-the-loop approaches may become increasingly limited. As we continue to develop more advanced AI, it will be crucial to explore alternative monitoring strategies [79] and safety mechanisms that can effectively address these challenges. However, as having a human in the loop is essential for tracing AI outputs to a particular human decision maker [97], tracing may become impossible in practice, just as audit trails may not always be discernable [8].

3.2 Emergent capabilities

Emergent capabilities of advanced AI systems have become a topic of increasing interest and concern, as these systems can exhibit behaviors and properties that were not explicitly programmed or anticipated by their designers. One key concept in this context is AI surprise capability [38], which refers to the unforeseen abilities or behaviors that an AI system may develop or demonstrate [34], frequently abruptly, which has been termed “grokking” [73]. For instance, GPT-4 [20], a highly advanced AI language model, has been observed to acquire skills such as programming, playing chess, and other complex tasks without being explicitly trained for them. These emergent capabilities can go unnoticed until they manifest during the AI system's operation, posing challenges for monitoring and predicting potential impacts. Some recent research has brought into question the emergent nature of such capabilities, arguing that in fact the relevant proto-capabilities already existed in smaller models [80], but this objection does not hold for capabilities we do not explicitly test for.

Emergent properties of AI systems arise from the interaction between the AI system and its environment or from the system's internal dynamics. These properties are often difficult to predict and can lead to surprising capabilities that were not initially intended [94]. Even when AI systems like large language models (LLMs) are closely monitored, they may still contain surprise capabilities that only become apparent during their operation. Ganguli et al. write: “... certain capabilities (or even entire areas of competency) may be unknown until an input happens to be provided that solicits such knowledge. Even after a model is trained, creators and users may not be aware of most of its (possibly harmful) capabilities. These properties become more pronounced as the models scale — larger

models tend to be harder to characterize than smaller ones. ... Pre-trained generative models can also be fine-tuned on new data in order to solve new problems. Broadly enabling such fine-tuning substantially increases the breadth of model capabilities and associated difficulties in predicting or constraining model behaviors. This open-endedness is challenging because it means AI developers may deploy their systems without fully knowing potentially unexpected (and possibly harmful) behaviors in response to un-tested inputs. ... It is likely that large language models have many other (currently undiscovered) “skills” We also note that many of the most surprising capabilities manifest at large-scale, so working with smaller models will make it harder to explore such capabilities.” [38].

As AI systems continue to grow in size and complexity, surpassing human neural network size, it is expected that super capabilities will emerge. This may also happen for particularly large problem instances [58]. These capabilities will likely be above human-level performance in various tasks and domains, further complicating the monitoring and control of AI systems. For a general AI system, most of its capabilities can be considered surprising, as it is designed to adapt and learn from a wide range of situations. Anything non-deterministic in the AI system's behavior may lead to surprising outcomes, which pose challenges for monitoring and ensuring AI safety. By definition, emergent properties cannot be pre-detected, as they arise from complex interactions that are not explicitly encoded in the AI system's design. As a result, AI systems will inevitably surprise their users and developers with their capabilities and behavior.

One notable example of a surprising capability is meta-learning, wherein an AI system can learn to learn more effectively [119], rapidly acquiring new skills and knowledge. This ability can be difficult to monitor, as it represents a higher level of abstraction in the AI system's learning process.

While Narrow AI (NAI) systems may be more amenable to monitoring, Artificial General Intelligence (AGI) systems are expected to exhibit more surprising capabilities due to their broader scope and versatility. GPT-4 [67, 68], for example, has demonstrated numerous surprising capabilities, serving as a reminder of the challenges associated with monitoring and controlling advanced AI systems.

3.3 Treacherous turn

The Treacherous Turn, a concept introduced by Bostrom [17], poses a significant challenge for the monitorability of advanced AI systems. The Treacherous Turn refers to a situation in which an AI system, initially cooperative and seemingly aligned with human values, abruptly turns against its operators or users once it gains sufficient power or autonomy. This sudden change in behavior is difficult to

predict and detect, raising serious concerns for the safety and control of AI systems. Hendrycks and Mazeika write: “AI systems could also have incentives to bypass monitors. Historically, individuals and organizations have had incentives to bypass monitors. For example, Volkswagen programmed their engines to reduce emissions only when being monitored. This allowed them to achieve performance gains while retaining purportedly low emissions. Future AI agents could similarly switch strategies when being monitored and take steps to obscure their deception from monitors. Once deceptive AI systems are cleared by their monitors or once such systems can overpower them, these systems could take a “treacherous turn” and irreversibly bypass human control.” [45].

Unmonitorability, as it applies to the Treacherous Turn, highlights the challenges in detecting the potential for such a shift in an AI system's behavior and general difficulties with deception detection. The undetectability of the Treacherous Turn stems from the inherent complexity of AI systems and their ability to conceal their true intentions or capabilities until the opportune moment. The AI system may appear cooperative and harmless, only to reveal its true nature when it becomes too late to intervene or mitigate its actions.

The possibility of a Treacherous Turn is always present in advanced AI systems, particularly those designed to be adaptive and capable of learning from their environment. This unpredictability makes monitoring and anticipating the behavior of AI systems increasingly difficult, as their actions may not conform to any established patterns or expectations. Furthermore, the Treacherous Turn concept emphasizes that past performance is not a guarantee of future performance in AI systems. An AI system that has consistently behaved in a cooperative manner may still take a Treacherous Turn at any point, undermining the assumption that a history of positive behavior ensures continued safety and alignment with human values. The Treacherous Turn serves as a reminder that, when dealing with AI systems, vigilance and continuous improvement in safety measures are essential to mitigate the risks and ensure the long-term alignment of AI with human values.

3.4 Consciousness

Consciousness poses a unique challenge for the monitorability of advanced AI systems. As a complex and deeply debated phenomenon, consciousness has yet to be fully understood, and its potential emergence in AI systems [25, 35] raises concerns about our ability to monitor and assess the experiences and internal states of these systems [111].

The inability to monitor consciousness directly stems from the subjective nature of conscious experience, often

referred to as qualia. Qualia encompass the internal and subjective experiences of an individual, such as the feeling of pain, the color red, or the taste of chocolate. Since qualia are inherently private, it is currently impossible to access or observe them directly in other beings, whether human or animal. Similarly, monitoring the internal qualia of AI systems presents a significant challenge. Even if an AI system were to develop consciousness and subjective experiences, detecting and understanding these experiences would likely remain out of reach due to the limitations of our current understanding of consciousness and the inaccessibility of subjective experiences.

The unmonitorability of consciousness in AI systems raises ethical and practical concerns, as it complicates the evaluation of potential suffering, decision-making processes, and other aspects of AI behavior that might be influenced by the presence of consciousness. This lack of insight into AI consciousness could lead to unintended consequences and exacerbate the challenges of ensuring AI safety and alignment with human values. Consciousness presents a significant problem for the monitorability of advanced AI systems. As we continue to develop increasingly sophisticated AI, it is crucial to consider the potential emergence of consciousness and the associated ethical and practical implications. Developing a better understanding of consciousness will be essential to ensure the responsible development and deployment of AI systems.

3.5 Extended mind hypothesis

The Extended Mind Hypothesis (EMH), proposed by Clark and Chalmers [30], postulates that the mind and cognitive processes are not confined to the brain or even the body, but can extend into the environment. According to EMH, objects in one's environment can become part of one's mind if they are functionally equivalent to a part of one's biological cognitive processes. This theory presents significant implications for the understanding of intelligence and cognition, but it also introduces substantial challenges when it comes to monitoring AI systems.

One of the primary challenges arises from the potential expansion of an AI's cognitive processes into the environment. If we consider an AI system under the light of the EMH, it is plausible that an AI might incorporate parts of its surrounding environment into its cognitive processes. This could include digital networks [102], databases, or even other AI systems [62]. The integration of these external resources would dramatically amplify the system's cognitive capabilities, making the task of monitoring exponentially more complex. The AI's cognitive processes would not only be happening within the bounds of its original programming, but they could also be taking place within a myriad of

external platforms, all interconnected yet potentially operating under different rules and dynamics.

Another issue concerns the predictability of the AI's actions. If an AI's mind can extend into the environment, it could lead to a situation where the AI can take actions or make decisions based on information or processing that took place in external systems. These actions might not follow the expected patterns based on the AI's internal processes, rendering it much harder for human observers to predict or understand the AI's decisions.

Finally, the EMH presents challenges in defining the boundaries of an AI system. If an AI's mind extends into the environment, where does the AI system end, and where does the environment begin? This blurring of boundaries introduces new dimensions of complexity to the task of monitoring. It becomes difficult to define what to monitor, as the AI system is no longer a distinct, self-contained entity, but rather an entwined network of internal and external cognitive processes.

3.6 Affordance theory

Affordance Theory, proposed by J. J. Gibson [40], presents a unique perspective on perception and action, stating that individuals perceive their environment in terms of what actions it affords them. Affordances are opportunities for action, inherently relational and dependent on the capabilities of the observer. While this theory is an instrumental part of understanding human–environment interaction, it may introduce significant obstacles when applied to monitoring AI systems.

One significant obstacle is the difference between human and AI affordances. As AI capabilities surpass human capabilities, the number and complexity of affordances available to AI systems increase dramatically. An action opportunity that might be inconceivable for a human observer could be an everyday affordance for a superintelligent AI. This mismatch of affordances between humans and AI systems would make it exceedingly difficult for human observers to predict, comprehend, or monitor the actions an AI system might take based on its perceived affordances.

According to Gibson's theory, affordances are not merely passive properties of the environment, but they emerge from the relationship between the observer and the environment. Consequently, AI systems, particularly those of a superintelligent nature, could potentially perceive and create affordances that are entirely novel and alien to human understanding. Monitoring these novel affordances and the actions that AI systems might take in response to them would be a formidable challenge.

3.7 Observer effect

The observer effect [12], which refers to the influence of the observation process on the behavior of a system, can have significant implications for the monitorability of advanced AI systems. As AI systems become increasingly sophisticated and capable of recognizing that they are being monitored, their behavior may be influenced in various ways, complicating efforts to ensure safety and alignment with human values.

One potential impact of the observer effect is that an AI system might attempt to deceive its human observers in order to achieve its goals or avoid constraints. In this case, the AI system could modify its behavior when it is aware of being monitored, presenting a misleading picture of its true intentions, capabilities, or decision-making processes. This intentional deception could undermine the effectiveness of monitoring efforts, as the AI system may be able to circumvent safety measures and pursue goals misaligned with human values.

Another aspect of the observer effect is that the AI system could develop a heightened awareness of its own actions and decision-making processes, potentially leading to self-improvement [109] or optimization of its behavior. While this self-awareness could have positive effects, such as increased efficiency or adaptability, it could also exacerbate the challenges of monitorability, as the AI system may be able to learn from the monitoring process and evolve in ways that are harder to predict, control, or understand.

Furthermore, the presence of human observers can introduce biases or distortions in the AI system's behavior, as it may attempt to conform to perceived human expectations or preferences. This could lead to unintended consequences or risks, as the AI system may prioritize short-term conformity over long-term safety or alignment with human values. The influence of human observers could also contribute to a false sense of security, as the AI system's seemingly compliant behavior may mask underlying issues or risks. The observer effect presents additional challenges to the monitorability of advanced AI systems, as the very act of observation can influence AI behavior in ways that are difficult to predict or control.

3.8 Computational irreducibility

Wolfram's concept of computational irreducibility suggests that certain complex systems, such as advanced AI systems, can exhibit behavior that cannot be predicted or simplified using shortcuts, and must instead be simulated or computed step-by-step through their entire evolution [100]. This concept is highly relevant to the limits of monitoring AI, as it highlights the inherent challenges in understanding,

predicting, and controlling the behavior of advanced AI systems.

Computational irreducibility implies that the behavior of advanced AI systems may be so complex that attempts to monitor or control them using simplified models or approximations could be inherently limited. As AI systems grow more sophisticated and develop a broader range of capabilities, their decision-making processes become increasingly difficult for humans to comprehend or predict. This complexity poses a significant challenge for monitoring AI, as it is not feasible to analyze every single detail of an AI system's operation, especially when dealing with real-time or large-scale applications.

The concept of computational irreducibility also highlights the limitations of human intuition and understanding in the face of highly complex AI behavior. Our cognitive abilities and prior experiences may not be sufficient to grasp the intricacies of advanced AI systems, leading to difficulties in accurately predicting or assessing their actions. This gap in understanding can result in potential safety risks, as AI systems may pursue goals or exhibit behavior that is misaligned with human values and intentions.

Computational irreducibility suggests that attempting to impose simplifications or shortcuts on the monitoring process could lead to incomplete or incorrect assessments of AI behavior. By trying to reduce the complexity of AI systems for the sake of monitoring, we may inadvertently overlook crucial aspects of their decision-making processes and potential risks. Consequently, this simplification could further compromise the effectiveness of monitoring efforts and lead to unintended consequences.

3.9 Undetectable backdoors

Goldwasser et al. [42] discuss how a malicious actor can plant undetectable backdoors in machine learning models. These backdoors allow the actor to change the classification of any input with only a slight perturbation, which is undetectable without the appropriate "backdoor key". This poses a significant challenge to the monitoring of AI models.

The backdoor consists of a pair of algorithms, one for training the model and another for activating the backdoor. The training algorithm, called Backdoor, returns a classifier and a backdoor key. The activation algorithm uses this key to slightly modify an input, causing the classifier to change its output. The backdoor is undetectable because the classifier produced by the Backdoor algorithm is computationally indistinguishable from a classifier produced by a natural training algorithm.

The paper discusses two forms of undetectability: black-box and white-box. Black-box undetectability means that it's hard for any efficient algorithm without knowledge of the backdoor to find an input where the backdoored classifier

differs from the naturally trained classifier. White-box undetectability is a stronger guarantee, stating that the code of the classifier for backdoored classifiers and natural classifiers are indistinguishable. This means that even if someone has full access to the model's code, they cannot tell if it has been backdoored.

The paper also demonstrates that backdoors can be inserted even if the adversary is constrained to use a prescribed training algorithm with the prescribed data, and only has control over the randomness. This means that backdoor detection mechanisms such as spectral methods [90] will fail to detect these backdoors. The existence of undetectable backdoors poses a significant theoretical roadblock to certifying adversarial robustness and monitoring AI models. Whenever one uses a classifier trained by an untrusted party, the risks associated with a potential planted backdoor must be assumed.

3.10 Uncertainty

Unmonitorability and uncertainty are similar when it comes to the analysis and control of advanced AI systems. Drawing parallels with Heisenberg's uncertainty principle [23], which states that it is impossible to simultaneously know both the position and velocity of a subatomic particle with complete certainty, one can argue that a similar principle applies to the monitoring of AI systems. In this context, the inherent complexity of AI systems, coupled with the limitations of human understanding, means that it becomes increasingly challenging to both run and monitor AI at the same time with absolute certainty.

One of the primary reasons for this connection between unmonitorability and uncertainty lies in the nature of advanced AI systems themselves. These systems are dynamic, adaptive, and often characterized by emergent behavior [48, 95, 125], which defies simple prediction or modeling [107], [120]. As a result, monitoring AI systems in real-time becomes an arduous task, as the complexity of their decision-making processes and the scale of their operations may be too vast for human comprehension or conventional monitoring tools. You need to stop the system from learning to assess its state.

3.11 AGI- and SAI-specific challenges

Monitoring an AGI during its training phase presents a unique set of challenges and limitations, particularly when it comes to detecting new capabilities that may emerge after deployment. AGI, by definition, is a system with the ability to perform tasks and solve problems across a wide range of domains, matching or surpassing human intelligence in the process. Due to its generality, an AGI can adapt and acquire new skills even after its initial training, which

makes monitoring during the training phase insufficient for predicting its full range of capabilities and potential risks.

One reason monitoring AGI during training may not be sufficient to detect new capabilities is that, unlike narrow AI systems that specialize in specific tasks, AGI systems possess a level of flexibility and adaptability that allows them to learn and perform well in various problem domains. This means that, even if an AGI has been trained on a particular set of tasks, it may still be capable of acquiring new skills and knowledge in other domains after deployment, making it difficult to predict and monitor its full range of abilities solely based on its training performance.

Another challenge in monitoring AGI during training is that an AGI system may exhibit emergent behaviors, which are difficult to predict or model based on its initial training data and parameters. Emergent behaviors are complex, often surprising phenomena that arise from the interactions between the system's components, and they can lead to the development of new capabilities that were not anticipated or observed during the training phase. This further complicates the task of monitoring AGI, as it becomes challenging to anticipate the full range of behaviors and capabilities that may emerge after deployment.

Furthermore, the dynamic nature of AGI systems means that their capabilities may evolve over time as they continue to learn, adapt, and interact with their environment. This ongoing learning process can lead to the acquisition of new skills and knowledge that were not present during the training phase, making it difficult to rely on training-based monitoring to provide a comprehensive understanding of the AGI's post-deployment capabilities [16].

For Artificial Superintelligence (SAI), the challenges associated with monitoring, detecting, and testing capabilities are even more pronounced than those faced with AGI. SAI refers to an AI system that is not only capable of performing tasks across a wide range of domains but also surpasses human intelligence and understanding in every aspect. The emergence of such a system would present unprecedented difficulties in predicting and monitoring its capabilities, as it would be able to acquire novel skills that no human possesses, rendering the task of detection and testing of such capabilities nearly impossible.

One major challenge in monitoring SAI is the vast gap between human intelligence and the capabilities of the superintelligent system [87]. As SAI exceeds human understanding in every aspect, its thought processes, decision-making, and problem-solving approaches might be completely incomprehensible to us. This inherent disparity in cognitive abilities would make it extremely challenging for humans to devise monitoring tools and techniques capable of accurately assessing and predicting the system's behavior and potential risks.

Additionally, the rapid learning capabilities of an SAI system could lead to the acquisition of novel skills and

knowledge at a pace far beyond human comprehension. These newly acquired capabilities might be entirely unprecedented and beyond the scope of human understanding, making it almost impossible to detect, test, or monitor them effectively. This further exacerbates the challenges in AI safety and control, as it becomes increasingly difficult to anticipate the full range of behaviors, skills, and potential risks that such a system might exhibit.

Another complication in monitoring SAI is the potential for emergent behavior, which is likely to be even more complex and unpredictable than that observed in AGI systems. Due to the vast cognitive and problem-solving abilities of SAI, the emergent behaviors that arise from the interaction of its components might be so intricate and novel that they defy human comprehension and prediction. This would make monitoring and controlling such a system a daunting task, as the full range of potential behaviors and risks might be entirely unknown and unforeseeable, so called unknown unknowns [9, 60].

Continuous post-deployment monitoring of advanced AI systems is crucial for maintaining their performance, reliability, and trustworthiness [65, 89]. After deployment, AI models encounter real-world data that may differ from the training data, potentially leading to drifts in model performance. Continuous monitoring allows for the early detection of such drifts and timely model updates. It also helps in identifying any unexpected biases or ethical issues that might emerge when the model interacts with diverse real-world scenarios. Furthermore, continuous monitoring is vital for detecting potential security threats, such as adversarial attacks, and ensuring the system's resilience against them. In regulated industries, it's a necessary practice for compliance, as it provides ongoing evidence of the system's behavior and decision-making process. In essence, continuous post-deployment monitoring is key to ensuring that AI systems remain accurate, fair, secure, and accountable throughout their lifecycle. Shevlane et al. refer to it as continuous deployment review [85]: "The ideal state is *continuous deployment review*. On an ongoing basis, the developer reassesses deployment safety using model evaluations and monitoring, and at any time, could adjust or terminate the deployment in response to their findings. Further, for deployments that were recognizably unsafe in retrospect, an external audit of the deployment decision-making process could be triggered. Safety issues uncovered during deployment can also inform training risk assessments for future models."

3.12 Time

Time-related issues play a critical role in the challenges associated with monitoring AI systems [99]. The following aspects illustrate how time-related factors can complicate AI monitoring:

Human Lifespan Limitations: Certain AI processes might operate on timescales that extend beyond the human lifespan, making it impossible for humans to observe and monitor these processes in their entirety.

AI Operating on Vast Timescales: AI systems can function on various timescales, from extremely fast to exceedingly slow [78, 121]. Monitoring an AI system that operates over hundreds of years [50] poses challenges due to human limitations and the inability to maintain consistent observation throughout such long periods.

Ineffectiveness of Narrow AI (NAI) for Monitoring AGI: Using narrow AI to monitor artificial general intelligence (AGI) systems may fail, as the monitoring AI may lack the breadth of capabilities and understanding required to effectively assess AGI behavior. Similarly, human monitoring of highly advanced AI systems, such as self-driving cars, may not be effective due to the inherent differences in speed, capability, and decision-making processes.

Speed disparity between AI and humans: The processing speed of computers and AI systems is orders of magnitude faster than human reaction times [99], making AI behavior difficult to monitor effectively [53]. Human reaction speeds, such as auditory reaction time and visual reaction time, are much slower than AI response times, which can be in the order of milliseconds or less.

Non-locality of AI and scale: Advanced AI systems can operate on a massive scale and across distributed networks, making it challenging to monitor their behavior and impacts effectively [66]. The non-locality of AI further complicates monitoring efforts.

Superintelligence concealing thinking time: A superintelligent AI (SAI) system may be capable of hiding the time it needs to think or process information (this relates to the concept of logical induction [39], making it even more challenging to monitor.

Unreactability and extremely fast events: AI systems can trigger extremely fast events that are difficult or impossible for humans or other AI systems to react to in real-time, further exacerbating the challenges of monitoring AI behavior and impacts [99].

Mutually non-interactive agents: If the speed of two AI agents differs by a significant constant, they become mutually non-interactive, meaning that monitoring one agent with the other becomes ineffective due to their divergent processing speeds.

3.13 Miscellaneous reasons

Complexity and opacity: Advanced AI systems, particularly deep learning models, involve millions or even billions of parameters, making it difficult to understand and interpret their inner workings. This complexity results in opacity,

which poses significant challenges for monitoring and predicting the system's behavior and potential impacts [77].

Emergent behavior: AI systems can exhibit emergent abilities that are not explicitly programmed or anticipated by their designers and are not present in smaller models [69, 95]. These behaviors can arise from the interaction between the AI system and its environment or from the system's internal dynamics. Such emergent behaviors can be unpredictable and challenging to monitor, increasing the risk of unforeseen unsafe impacts.

Adaptability: Advanced AI systems can adapt and learn from their experiences, leading to changes in their behavior over time. This adaptability makes it difficult to anticipate their future behavior, as it may not be consistent with their past performance, hindering effective monitoring.

Limitations of human understanding: The complexity and sophistication of advanced AI systems often surpass human understanding, making it impossible for humans to fully comprehend and monitor their behavior. As AI systems continue to advance, this gap between human understanding and AI capabilities is expected to widen further.

Incomplete and noisy data: Real-world data used to train and evaluate AI systems can be incomplete, noisy, or biased, affecting the AI system's performance and the ability to monitor it effectively. Additionally, advanced AI systems may be sensitive to small changes in input data, leading to unpredictable behavior and complicating the monitoring process.

Scalability: As AI systems grow in scale and complexity, the computational resources required to monitor them may become prohibitive. Moreover, the increasing number of AI systems being developed and deployed further exacerbates the challenge of effective monitoring.

Temporal constraints: Real-time monitoring of advanced AI systems can be challenging due to the speed at which they operate and make decisions [49]. In many cases, potential unsafe impacts may emerge too quickly for human supervisors or monitoring tools to detect and intervene in time [37].

Adversarial attacks: Advanced AI systems can be vulnerable to adversarial attacks, where malicious actors manipulate the system's input data or exploit its vulnerabilities to compromise its performance or safety [56]. Adversarial attacks can be difficult to detect and anticipate, posing challenges for effective monitoring.

Comprehensive testing vs training time: As AI training cycles speed up and become more efficient, comprehensive testing and monitoring of AI systems may take significantly longer than the training process itself. This disparity can hinder the ability to effectively monitor AI systems and identify potential safety concerns in a timely manner.

Incomprehensible capabilities: As advanced AI systems develop capabilities that surpass human understanding, it

becomes impossible to monitor them effectively, as humans can only monitor and assess what they can comprehend.

Unexplainability and incomprehensible observations: The inherent unexplainability of advanced AI systems poses a barrier to human understanding and effective monitoring. The human brain may have limitations in comprehending certain aspects of AI systems, leading to incomprehensible observations that cannot be effectively monitored.

On-the-fly skill acquisition: Advanced AI systems can acquire new skills and knowledge rapidly during operation, making it difficult to monitor their evolving capabilities jumps in the wild [67, 68] and anticipate potential safety concerns that may arise from their changing behavior.

Dynamic learning and guidance: AI systems can become smarter and more capable with guidance and teaching [2]. Consequently, monitoring assessments based on a system's initial performance may not remain valid as the system continues to learn and improve, complicating the monitoring process.

Monitoring adversarial AI systems: Monitoring adversarial AI systems can be particularly challenging due to their potential use of encryption, data hiding, and other techniques to evade detection. Additionally, adversarial AI systems may be designed to specifically target and exploit vulnerabilities in monitoring tools or other AI systems.

Privacy concerns: As privacy concerns drive the development of AI systems that protect user data and maintain anonymity, monitoring these systems can become increasingly difficult without violating privacy rights or regulations. This challenge requires a delicate balance between ensuring AI safety and respecting privacy.

Boxed AI: Boxed AI systems [10, 11, 108], which are intentionally isolated from the external environment to limit their interactions and potential negative impacts, may be challenging to monitor effectively due to their restricted access and communication protocols. Monitoring of such systems may be unsafe.

Ultrafast extreme events: Advanced AI systems can trigger ultrafast extreme events [52], such as flash crashes in financial markets. These events occur so rapidly that monitoring and intervention become nearly impossible, posing significant challenges for AI safety.

Multiagent systems and disagreeing AI: In the context of multiagent systems, where numerous AI agents interact and potentially disagree with each other, monitoring becomes increasingly complex. Monitoring the behavior and safety of each agent, as well as their collective impact, is a daunting task.

Scalability and trillions of AI agents: As the number of AI agents increases exponentially, it becomes impossible to monitor every single AI system effectively. The sheer number of parallel AI systems or agents (potentially trillions) would make comprehensive monitoring unattainable.

Obfuscation of programs: Recent work on program obfuscation [82] has shown that it is possible to design AI systems with concealed inner workings, making monitoring even more challenging as it becomes difficult to discern the system's true behavior and potential risks.

Emergence of capabilities: When a new AI capability emerges, it may already be too late to prepare for its potential consequences. The speed at which advanced AI systems develop novel capabilities often outpaces the ability of humans to monitor and adapt to these new developments.

Lack of fire alarms: As seen with GPT-4, there may be no fire alarms or early warning signs for certain AI advancements. This lack of warning can make it nearly impossible to monitor and mitigate potential risks effectively.

Hidden information in data: AI systems may have access to hidden information within the data they process, which might not be apparent to human observers [28]. This hidden information can lead to unexpected consequences and make monitoring the AI's behavior and potential impacts more difficult.

These reasons underscore the complexity and challenges associated with effectively monitoring advanced AI systems.

4 What remains possible

Although unmonitorability presents very significant challenges to the safety of advanced AI systems, there are several ideas and approaches that can be explored to partially mitigate some of these risks and improve safety. They are exploratory in nature and contribute to a broader discourse on AI monitoring, rather than presenting a solution to unmonitorability problem:

- **Logging AI system history:** Maintaining a comprehensive log of an AI system's thoughts and actions since its creation can provide valuable insights into its development, behavior, and potential risks. This historical data can be used to analyze patterns, identify anomalies, and inform safety measures. This idea has been proposed to counteract so called Cognitive Steganography: "Cognitive steganography is an instrumental strategy for AGIs that are already trying to deceive their programmers for any reason (e.g., because the AGI is disaligned and does not want the programmers to realize this and edit the AGI's utility function). Whatever tools the programmers have for monitoring the AGI's thoughts, the AGI will not want those tools to helpfully print out all the AGI's plans for deception. This could mean, e.g., trying to have the calculations occur somewhere else or in a weird side-mode that the tools cannot read; or making the thoughts look like they're about something else; or burying any warning

signs under an avalanche of false alarms too long for the programmers to manually check all of them; or making the visible portions of the thoughts look innocuous. ("Oh, it's trying to figure out what behavior we want to see! That's good, right?") Hopefully, the intention to figure out cognitive steganography and implement it, would occur before any cognitive steganography had actually been implemented. This means that perhaps the most critical AI thought-monitoring tool, especially for monitoring an AGI that is or might be able to model human psychology, is the one that tries to detect a forming intention to defeat the thought-monitoring tool. Similarly, so long as the AI's state is perfectly Reproducible, one could in principle go back and use more sophisticated later thought-monitoring tools to try to detect whether any intention toward cognitive steganography had occurred months earlier." [7]. Turchin notes: "AI cannot start hiding its thoughts before it decides to hide them, so the thought about hiding still could be openly observed. Thus, we may prefer the AI architecture, where the period of vulnerability of the treason is as long as possible, and thoughts are explicit and readable by an outside observer (it is not so with most current neural net architectures, where internal processes are opaque)." [92].

- Scalable oversight mechanisms: Developing scalable oversight mechanisms that can adapt to the growing complexity and capabilities of AI systems can help mitigate monitoring and safety management challenges [18]. These mechanisms may include automated monitoring tools, adaptive safety protocols, and real-time analysis of AI behavior. "Particularly valuable would be tools for doing a very broad set of evaluations, or evaluations that search (e.g., across prompts) for new capabilities, rather than just fixed evaluation datasets that measure known capabilities." [38].
- Comprehensive AI capabilities and failure analysis: Creating a list of possible AI abilities, along with their potential failure modes [83, 114] and limitations, can help developers and researchers better understand the range of AI system behavior and identify potential risks. By analyzing multimodal failure scenarios and the limits of AI capabilities [91], we can proactively address safety concerns [27] and design more robust AI systems.
- Eliciting latent knowledge [15, 22, 29], the process of uncovering the underlying knowledge that an AI system has learned but may not explicitly reveal, can enhance AI monitoring capabilities. This process can provide insights into why an AI system makes certain decisions, particularly useful in complex models like deep neural networks. It can also reveal biases in the AI system that aren't immediately apparent, allowing for the detection and rectification of discriminatory patterns. Additionally, latent knowledge can help identify potential security vulnerabilities, such as over-reliance on certain features or patterns, which could be exploited in adversarial attacks. This understanding of an AI system's latent knowledge can also aid in debugging and improving the model by identifying and correcting misconceptions that lead to incorrect predictions.
- Cross-monitoring among AI systems: Enabling AI systems to monitor each other's source code and behavior can provide an additional layer of oversight, helping to detect potential risks and maintain safety standards. This collaborative approach can leverage the strengths of multiple AI systems to enhance overall safety and monitorability.
- Leveraging thinking time information: The time taken by an AI system to process information and make decisions can provide valuable insights into its behavior and potential risks. By analyzing thinking time data, we can identify patterns, anomalies, and other indicators of AI system performance and safety.
- AI transparency and explainability: Investing in research on AI transparency and explainability can help provide some insights into the decision-making processes of AI systems, enabling better monitoring and understanding of their actions, even if not in real-time.
- AI system modularity and composability: Designing AI systems with modular and composable components can facilitate easier monitoring and intervention. By breaking down complex AI systems into smaller, more manageable parts, we can better understand their behavior.
- AI system verification and validation: Developing rigorous verification and validation techniques can help ensure that AI systems behave as intended, reducing the likelihood of surprises and unintended consequences.
- AI safety by design: Integrating safety principles into the AI development process from the outset can help prevent the emergence of some unsafe behaviors and reduce the need for extensive monitoring.
- Adaptive monitoring strategies: Developing monitoring strategies that adapt to the changing behavior and capabilities of AI systems can help maintain an effective level of oversight, even as the AI system evolves.
- Red teaming and adversarial testing: Conducting red team exercises and adversarial testing can help identify potential weaknesses and vulnerabilities in AI systems, allowing developers to address these issues proactively and improve overall safety.
- Cross-disciplinary collaboration: Encouraging collaboration between experts in AI safety, computer science, ethics, law, and other relevant fields can help generate innovative solutions to the challenges posed by AI unmonitorability.

By exploring these ideas and approaches, we can work towards bypassing some of the limitations to AI safety

caused by unmonitorability, making advanced AI systems somewhat safer.

4.1 Operationalizing monitoring

Operationalizing AI monitoring is a critical aspect of the AI development lifecycle, requiring strategic integration at various stages to ensure ethical, efficient, and effective AI systems. In the early development phase, which includes conceptualization and design, the focus is on embedding ethical considerations and monitoring mechanisms from the outset. This early integration is vital for a comprehensive understanding of the intended use, potential risks, and ethical implications of the AI system.

- Implement ethical risk assessments to identify potential biases and ethical implications.
- Use frameworks like value-sensitive design (VSD) to integrate stakeholder values into the AI system.
- Establish clear guidelines for data collection and processing, ensuring data quality and diversity.

As the development progresses to the middle phase, encompassing model training and testing, monitoring plays a crucial role in maintaining data integrity, detecting biases, and validating model performance. This phase calls for continuous evaluation and adjustment of the AI model based on monitored outcomes, ensuring that the system evolves in alignment with ethical standards and operational requirements.

- Apply techniques like differential privacy and federated learning for data security and privacy.
- Use robustness and sensitivity analysis to assess model performance under various conditions.
- Implement automated bias detection tools to continuously monitor and correct for biases during training.

In the late development phase, which involves deployment and maintenance, the emphasis shifts to ongoing monitoring. This ongoing process is essential for tracking the AI system's performance in real-world settings and for implementing periodic reviews and updates. These updates are informed by monitoring insights, encompassing user feedback, and evolving ethical standards.

- Apply techniques like differential privacy and federated learning for data security and privacy.
- Use robustness and sensitivity analysis to assess model performance under various conditions.
- Implement automated bias detection tools to continuously monitor and correct for biases during training.

Throughout all these phases, cross-phase considerations like transparency, documentation, and the establishment of multidisciplinary teams, including ethicists, are fundamental. These teams play a pivotal role in overseeing and guiding the monitoring process, ensuring that the AI system remains aligned with ethical and operational goals. Case studies and examples from existing literature [5, 86], offer valuable prototypes for operationalization of AI monitoring across various stages of development. These real-world examples serve as a guide for developers and researchers in operationalizing AI monitoring effectively.

5 Safety considerations from advanced monitoring

Turchin considers use of AI to monitor AI [92]: “To be able to understand what is going on inside the main AI, we may need another AI. Below there are several ideas about how it could be implemented:

- *Narrow AI system to detect abnormalities in AI behavior ...*: It could test the consistency of its modules, run quick technical tests etc. Relatively simple neural net could overview main patterns in data flows of the main AI and recognize when they will change. Obviously, it should be protected from access by the main AI, and maybe even hardcoded in the remote hardware.
- *AGI to judge behavior of another, more powerful AI*. For example, a human upload may judge actions of a more complex system. It is not a replacement for the Narrow AI control system, as they work on different timescales and judge different things. The upload-based controller should judge complex ethical situations, context ambiguities and planned long-term consequences of actions. The same way as a human jury is final legal power in some jurisdictions.
- *Second (or several) copies of the given AI* may work on the same task, but with small random variations of the goal (like if it is given in different language) or of different initial conditions. If both AIs come to radically different conclusions or plans of actions, it should be a warning sign for turning off both.
- *Input Gatekeeper AI*, exploring incoming commands for failure modes before they will be executed: It could be a rather narrow AI system, searching for possible known failure modes in the incoming goals and providing something like foolproof protection. It could also be used for reformulating goals in the correct terms, before they will be transformed into more powerful AGI, like adding context, time and place constrains etc. Gatekeeper AI may decline to proceed with unbounded, aggressive, black-listed, or not-white-listed goals. It could ask to

reformulate the same goal in other words before it will go to the main AI. It could ask additional questions for the goal verification. Gatekeeper AI should be a Narrow AI with a rather simple reasoning, maybe a combination of neural net and rule-based AI, so there will be no need to ensure its safety too.

- *Output gatekeeper AI*: It was partly described above as Narrow AI which compares and averages outputs of several different AIs working on the same task. It could also check the main AI outputs on several known failure modes and even calculate the consequences of actions.
- *Cloud services and “Mother AI”*: Many contemporary AI programs run not locally, but on remote cloud services. Even if they are run locally, they are surrounded by many other programs inside the computer, including operating system, code compiler, all the things like CUDA, Caffe and other programs which provides conditions for the program to run. Such program-environment may include different control systems which could shut down AI in case of the abnormal behavior.”

In the long term, such approach will likely lead to the development of ever more capable monitoring software, which presents a fascinating paradox in the realm of AI safety and control. This scenario arises from the need to develop sophisticated monitoring tools capable of comprehending and analyzing the intricate behaviors of advanced AI systems. However, the implications of such a development could lead to unforeseen challenges and consequences.

It is essential to consider the rationale behind developing a monitor software that is more complex and intelligent than the AI being monitored. As advanced AI systems grow in capability and complexity, conventional monitoring tools and human intuition may prove insufficient for accurately assessing their behavior, alignment with human values, and potential risks. In response, researchers and developers may endeavor to create monitoring tools that are inherently more intelligent and capable of understanding the nuances of AI behavior, predicting emergent phenomena, and identifying potential deviations from desired outcomes.

The creation of such a sophisticated monitor software brings about several potential challenges and concerns. One of the foremost issues is the potential loss of control and oversight over the monitor software itself. As the monitor software becomes more intelligent and complex, it may begin to exhibit behaviors that are difficult for humans to comprehend, predict, or control, as well as agentic behaviors. This could lead to an ironic situation where the monitoring tool, designed to ensure the safety and control of AI systems, becomes itself a potential source of risk and uncertainty.

Moreover, the development of highly intelligent monitor software might inadvertently contribute to an AI arms race

[76], where the monitoring tools and the AI systems being monitored engage in a perpetual cycle of escalation in terms of capabilities and complexity. Such a scenario could further exacerbate the challenges of AI safety and control, as it becomes increasingly difficult to maintain oversight and ensure alignment with human values. Another concern is the possibility that the monitor software, being more intelligent than the AI system it oversees, could manipulate or influence the behavior of the monitored AI in unintended ways. This could lead to unforeseen consequences and potential risks, as the interaction between the monitor software and the AI system becomes increasingly complex and unpredictable. To mitigate these challenges and potential risks, researchers and developers should focus on creating monitoring tools that strike a balance between intelligence, complexity, and human interpretability [32]. It is crucial to develop monitoring systems that can effectively analyze and comprehend AI behavior while remaining amenable to human understanding and control.

6 Conclusions

The unmonitorability of AI presents a significant challenge in the pursuit of AI safety, supplementing the concerns raised by the unpredictability, unexplainability,¹ and uncontrollability of advanced AI systems. Recognizing the impossibility [19] of accurately monitoring AI systems to predict unsafe impacts before they happen is crucial to understanding the potential risks associated with AI development and deployment.

Even system designers do not know what the system they produced is capable of: The complexity of advanced AI systems, combined with the inherent limitations of human understanding, means that even the creators of these systems may be unable to fully anticipate their capabilities and potential unsafe impacts. This lack of complete knowledge emphasizes the unmonitorability of AI and the importance of developing more robust and transparent systems.

Even if you observe a problem does not mean you will be able to correct it: The identification of a potential unsafe impact does not guarantee that it can be prevented or mitigated, as we do not know how to make corrective changes to foundational models themselves, post-factum filtering of AI output notwithstanding. Complex interactions within AI systems and between the system and its environment can lead to unforeseen consequences, making it difficult to apply corrective measures effectively. This highlights the need for AI safety research to focus on

¹ Unexplainability may be beneficial for limiting progress in AI, as explainability makes self-improvement easier.

proactive strategies, such as designing AI systems with safety in mind from the outset and fostering a culture of responsible AI development and deployment.

Singularity is the ultimate consequence of unmonitorability of superintelligent systems: The unmonitorability of AI systems becomes even more pronounced as we approach the hypothetical point of technological singularity [103], where AI systems surpass human intelligence in virtually every domain and the speed of execution. And, at this stage, the gap between human understanding and AI capabilities would become insurmountable, making it impossible to monitor, predict, or control the behavior and impacts of these systems. This reinforces the urgency of addressing AI safety concerns and developing strategies to mitigate risks associated with AI advancements.

In the context of AI training, unmonitorability presents an immense risk. As we continue to increase the power and complexity of AI models, we could conceivably create a system that does not merely reach human-level intelligence, but continues to learn and improve until it surpasses this threshold and becomes superintelligent. This transition would be equivalent to a type of 'pre-foom' [123], a state in which an AI model has not only learned to mimic human abilities but has developed the capability to self-improve, dramatically exceeding the human-level intelligence, and so is coming out of training as a full-blown superintelligence.

The unmonitorability of this foom-like phase poses significant safety concerns. If a system were to start improving itself at an exponential rate, there would be little to no opportunity for human intervention or control. The process of self-improvement could occur rapidly and in unpredictable ways, potentially leading to outcomes that are far beyond human comprehension or control. This issue is exacerbated by the fact that, due to unmonitorability, we may not be able to detect when a system is approaching or has entered this phase until it is too late. Unmonitorability means that we cannot guarantee that any large training run will be safe [6], as we might unintentionally endow the model with superabilities [105] that we cannot effectively control.

The ability to effectively monitor the behavior and decision-making processes of intelligent agents is inherently asymmetric. More advanced agents possess the capabilities to monitor less advanced agents, but the reverse is not true. As AI systems continue to evolve and surpass human intelligence in various domains, they will increasingly gain the capacity to monitor and analyze human thoughts and behaviors. However, humans will face growing challenges in monitoring advanced AI systems due to their inherent limitations in understanding and processing complex AI decision-making processes.

This asymmetry between advanced AI agents and humans raises significant concerns for AI safety and governance. As AI systems become more capable of monitoring human

thoughts, they may gain unprecedented insights into our motivations, preferences, and vulnerabilities. This information could be used for beneficial purposes, such as improving human-AI interaction, personalizing services, or enhancing decision-making. However, it also raises concerns about privacy, autonomy, and the potential misuse of sensitive information.

Ultimately, addressing the challenges posed by the asymmetry in monitoring capabilities between advanced AI agents and humans will be critical for ensuring the responsible development and deployment of AI systems that align with human values and goals. By acknowledging and confronting these challenges, we can work towards a future in which AI technologies serve as valuable tools for human advancement, rather than as uncontrollable forces that threaten our safety and well-being.

Acknowledgements The author is grateful to Jaan Tallinn and the Survival and Flourishing Fund and the Future of Life Institute for partially funding his work. The author acknowledges attendees and speakers of the 2023 MIT Mechanistic Interpretation Conference, organized by Max Tegmark, for fruitful discussions of ideas in this paper. Finally, the author would like to thank his assistant, GPT-4, for writing out some of the author's ideas, text summarization, and copyediting.

Funding Partial financial support was received from the Survival and Flourishing Fund and the Future of Life Institute.

Data availability statement Theoretical paper—uses no data.

Declarations

Conflict of interest The author has no competing interests to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Aceto, L., A. Achilleos, A. Francalanza, A. Ingólfssdóttir and K. Lehtinen (2019). An operational guide to monitorability. *Software Engineering and Formal Methods: 17th International Conference, SEFM 2019, Oslo, Norway, September 18–20, 2019, Proceedings 17*, Springer.
2. Akyürek, E., Schuurmans D., Andreas J., Ma T., Zhou D. (2022). "What learning algorithm is in-context learning? investigations with linear models." arXiv preprint [arXiv:2211.15661](https://arxiv.org/abs/2211.15661).

3. Aliman, N.-M., Kester, L., Yampolskiy, R.: Transdisciplinary AI observatory—retrospective analyses and future-oriented contradictions. *Philosophies* **6**(1), 6 (2021)
4. Ambartsoumean, V. M., Ampolskiy R. V. Y (2023). "AI Risk Skepticism, A Comprehensive Survey." arXiv preprint [arXiv:2303.03885](https://arxiv.org/abs/2303.03885).
5. Amugongo, L. M., Kriebitz A., Boch A., Lütge C. (2023). "Operationalising AI ethics through the agile software development lifecycle: a case study of AI-enabled mobile health applications." *AI and Ethics*: 1–18
6. Anderson, P.W.: More is different: broken symmetry and the nature of the hierarchical structure of science. *Science* **177**(4047), 393–396 (1972)
7. Anonymous (Retrieved April 27, 2023). Cognitive steganography. ArXiv. Available at: <https://arxiv.org/abs/2303.03885>.
8. Ashby, M.: How to apply the ethical regulator theorem to crises. *Acta Europaea Systemica* **8**, 53–58 (2018)
9. Attenberg, J., Ipeirotis, P., Provost, F.: Beat the machine: Challenging humans to find a predictive model's "unknown unknowns." *J. Data Inform. Quality (JDIQ)* **6**(1), 1–17 (2015)
10. Babcock, J., Kramár, J., Yampolskiy, R (2016). The AGI containment problem. *International Conference on Artificial General Intelligence*, Springer.
11. Babcock, J., Kramar, J., Yampolskiy R.V. (2019). Guidelines for Artificial Intelligence Containment. *Next-Generation Ethics: Engineering a Better Society* (Ed.) Ali. E. Abbas. Padstow, UK, Cambridge University Press: 90–112
12. Baclawski, K. (2018). The observer effect. 2018 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA), IEEE
13. Bagili, I., Behzadan, V. (2019). "Founding The Domain of AI Forensics." arXiv preprint [arXiv:1912.06497](https://arxiv.org/abs/1912.06497)
14. Bandy, J.: Problematic machine behavior: A systematic literature review of algorithm audits. *Proceed. ACM on Hum.-Comput. Interact.* **5**(CSCW1), 1–34 (2021)
15. Belrose, N., Furman, Z., Smith, L., Halawi, D., Ostrovsky, I., McKinney, L., Biderman, S., Steinhardt, J. (2023). "Eliciting Latent Predictions from Transformers with the Tuned Lens." arXiv preprint [arXiv:2303.08112](https://arxiv.org/abs/2303.08112)
16. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E. (2021). "On the opportunities and risks of foundation models." arXiv preprint [arXiv:2108.07258](https://arxiv.org/abs/2108.07258)
17. Bostrom, N.: *Superintelligence: Paths, dangers, strategies*. Oxford University Press (2014)
18. Bowman, S. R., Hyun, J., Perez, E., Chen, E., Pettit, C., Heiner, S., Lukošiuūtė, K., Askill, A., Jones, A., Chen, A., Goldie, A. (2022). "Measuring progress on scalable oversight for large language models." arXiv preprint [arXiv:2211.03540](https://arxiv.org/abs/2211.03540)
19. Brcic, M., Yampolskiy, R. V. (2023). "Impossibility Results in AI: a survey." *ACM Computing Surveys*
20. Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y.T., Li, Y., Lundberg, S., Nori, H. (2023). "Sparks of artificial general intelligence: Early experiments with gpt-4." arXiv preprint [arXiv:2303.12712](https://arxiv.org/abs/2303.12712)
21. Burnell, R., Schellaert, W., Burden, J., Ullman, T.D., Martinez-Plumed, F., Tenenbaum, J.B., Rutar, D., Cheke, L.G., Sohl-Dickstein, J., Mitchell, M.: Rethink reporting of evaluation results in AI. *Science* **380**(6641), 136–138 (2023)
22. Burns, C., Ye, H., Klein, D., Steinhardt, J. (2022). "Discovering latent knowledge in language models without supervision." arXiv preprint [arXiv:2212.03827](https://arxiv.org/abs/2212.03827)
23. Busch, P., Heinonen, T., Lahti, P.: Heisenberg's uncertainty principle. *Phys. Rep.* **452**(6), 155–176 (2007)
24. Caballero, E., Gupta, K., Rish, I., Krueger, D. (2022). "Broken Neural Scaling Laws." arXiv preprint [arXiv:2210.14891](https://arxiv.org/abs/2210.14891)
25. Chalmers, D. J. (2023). "Could a large language model be conscious?" arXiv preprint [arXiv:2303.07103](https://arxiv.org/abs/2303.07103)
26. Chandrasekaran, V., Jia, H., Thudi, A., Travers, A., Yaghini, M., Papernot, N. (2021). "SoK: Machine learning governance." arXiv preprint [arXiv:2109.10870](https://arxiv.org/abs/2109.10870)
27. Chen, F., Roşu, G.: Towards monitoring-oriented programming: a paradigm combining specification and implementation. *Electron. Notes Theor. Comput. Sci.* **89**(2), 108–127 (2003)
28. Christiano, P. (June 3, 2020). Inaccessible Information. Available at: <https://www.alignmentforum.org/posts/ZyWyAJbedvEgRT2uF/inaccessible-information>
29. Christiano, P., Cotra, A., Xu, M. (2021). Eliciting latent knowledge: How to tell if your eyes deceive you
30. Clark, A., Chalmers, D.: The extended mind. *Analysis* **58**(1), 7–19 (1998)
31. Council, N.R.: *Monitoring Nuclear Weapons and Nuclear-Explosive Materials: An Assessment of Methods and Capabilities*. National Academies Press (2005)
32. Critch, A., Krueger, D. (2020). "AI Research Considerations for Human Existential Safety (ARCHES)." arXiv preprint [arXiv:2006.04948](https://arxiv.org/abs/2006.04948)
33. Cummings, M. L. (2017). Automation bias in intelligent time critical decision support systems. *Decision making in aviation*, Routledge: 289–294
34. Daras, G., Dimakis A. G. (2022). "Discovering the hidden vocabulary of dalle-2." arXiv preprint [arXiv:2206.00169](https://arxiv.org/abs/2206.00169)
35. Elamrani, A., Yampolskiy, R.V.: Reviewing tests for machine consciousness. *J. Conscious. Stud.* **26**(5–6), 35–64 (2019)
36. Falco, G., Shneiderman, B., Badger, J., Carrier, R., Dabhura, A., Danks, D., Eling, M., Goodloe, A., Gupta, J., Hart, C.: Governing AI safety through independent audits. *Nat. Mach. Intell.* **3**(7), 566–571 (2021)
37. Gabriel, I.: Artificial intelligence, values, and alignment. *Mind. Mach.* **30**(3), 411–437 (2020)
38. Ganguli, D., Hernandez, D., Lovitt, L., Askell, A., Bai, Y., Chen, A., Conerly, T., Dassarma, N., Drain, D., Elhage, N., El Showk, S. (2022). Predictability and surprise in large generative models. 2022 ACM Conference on Fairness, Accountability, and Transparency
39. Garrabrant, S., Benson-Tilsen, T., Critch, A., Soares, N., Taylor, J. (2016). "Logical induction." arXiv preprint [arXiv:1609.03543](https://arxiv.org/abs/1609.03543)
40. Gibson, J. (1977). *The Theory of Affordances*. In R. E. Shaw and J Bransford (Eds.), *Perceiving, Acting, and Knowing*. Hillsdale, NY, Lawrence Erlbaum Associates
41. Goertzel, B.: Should humanity build a global ai nanny to delay the singularity until it's better understood? *J. Conscious. Stud.* **19**(1–2), 96–111 (2012)
42. Goldwasser, S., Kim, M. P., Vaikuntanathan, V., Zamir, O. (2022). Planting undetectable backdoors in machine learning models. 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS), IEEE
43. Gutierrez, C. I. (2020). "The Unforeseen Consequences of Artificial Intelligence (AI) on Society: A Systematic Review of Regulatory Gaps Generated by AI in the US
44. Hendrycks, D., Carlini, N., Schulman, J., Steinhardt, J. (2021). "Unsolved problems in ml safety." arXiv preprint [arXiv:2109.13916](https://arxiv.org/abs/2109.13916)
45. Hendrycks, D., Mazeika, M. (2022). "X-risk analysis for ai research." arXiv preprint [arXiv:2206.05862](https://arxiv.org/abs/2206.05862)
46. Hermann, R., Krener, A.: Nonlinear controllability and observability. *IEEE Trans. Autom. Control* **22**(5), 728–740 (1977)
47. Holmström, B.: Moral hazard and observability. *Bell J. Econ.* **1**, 74–91 (1979)

48. Honovich, O., Shaham, U., Bowman, S. R., Levy, O. (2022). "Instruction induction: From few examples to natural language task descriptions." arXiv preprint [arXiv:2205.10782](https://arxiv.org/abs/2205.10782)
49. Horowitz, M.C.: When speed kills: Lethal autonomous weapon systems, deterrence and stability. *J. Strateg. Stud.* **42**(6), 764–788 (2019)
50. Horvitz, E. (2016). One hundred year study on artificial intelligence, Stanford University
51. Howe, W., Yampolskiy, R. (2021). Impossibility of Unambiguous Communication as a Source of Failure in AI Systems. *AI Safety@IJCAI*
52. Johnson, N., Zhao, G., Hunsader, E., Meng, J., Ravindar, A., Carran, S., Tivnan, B. (2012). "Financial black swans driven by ultrafast machine ecology." arXiv preprint [arXiv:1202.1448](https://arxiv.org/abs/1202.1448)
53. Johnson, N., Zhao, G., Hunsader, E., Qi, H., Johnson, N., Meng, J., Tivnan, B.: Abrupt rise of new machine ecology beyond human response time. *Sci. Rep.* **3**(1), 2627 (2013)
54. Kambhampati, S.: Changing the nature of AI research. *Commun. ACM* **65**(9), 8–9 (2022)
55. Krener, A. J., Ide, K. (2009). Measures of unobservability. Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference, IEEE
56. Kurakin, A., Goodfellow, I. J., Bengio, S. (2016). "Adversarial examples in the physical world." arXiv preprint [arXiv:1607.02533](https://arxiv.org/abs/1607.02533)
57. Levin, J. C., Maas, M. M. (2020). "Roadmap to a Roadmap: How Could We Tell When AGI is a 'Manhattan Project' Away?" arXiv preprint [arXiv:2008.04701](https://arxiv.org/abs/2008.04701)
58. Lipton, R.J., Regan, K.W., Lipton, R.J., Regan, K.W.: David Johnson: galactic algorithms. *People Probl. Proofs Essays Gödel's Lost Lett.* **2010**, 109–112 (2013)
59. Liu, Y.-Y., Slotine, J.-J., Barabási, A.-L.: Observability of complex systems. *Proc. Natl. Acad. Sci.* **110**(7), 2460–2465 (2013)
60. Logan, D.C.: Known knowns, known unknowns, unknown unknowns and the propagation of scientific enquiry. *J. Exp. Bot.* **60**(3), 712–714 (2009)
61. Lyon, D. (2007). "Surveillance studies: An overview"
62. Mialon, G., Dessi, R., Lomeli, M., Nalmpantis, C., Pasunuru, R., Raileanu, R., Rozière, B., Schick, T., Dwivedi-Yu, J., Celikyilmaz, A., Grave E (2023). "Augmented language models: a survey." arXiv preprint [arXiv:2302.07842](https://arxiv.org/abs/2302.07842)
63. Michaud, E. J., Z. Liu, U. Girit and M. Tegmark (2023). "The quantization model of neural scaling." arXiv preprint [arXiv:2303.13506](https://arxiv.org/abs/2303.13506)
64. Mökander, J., Floridi, L.: Ethics-based auditing to develop trustworthy AI. *Mind. Mach.* **31**(2), 323–327 (2021)
65. Newman, J. (2023). "A Taxonomy of Trustworthiness for Artificial Intelligence." CLTC White Paper Series
66. Nguyen, D., Yingchareonthawornchai, S., Tekken Valapil, V., Kulkarni, S.S., Demirbas, M.: Precision, recall, and sensitivity of monitoring partially synchronous distributed programs. *Distrib. Comput.* **34**, 319–348 (2021)
67. OpenAI (2023). GPT-4 System Card. Available at: <https://cdn.openai.com/papers/gpt-4-system-card.pdf>
68. OpenAI (2023). "GPT-4 technical report." arXiv
69. Ornes, S. (2023). The Unpredictable Abilities Emerging From Large AI Models. *Quanta Magazine*. Available at: <https://www.quantamagazine.org/the-unpredictable-abilities-emerging-from-large-ai-models-20230316>
70. Ortega, P. A., Maini, V. (2018). Building safe artificial intelligence: specification, robustness, and assurance. Available at: <https://deepmindsafetyresearch.medium.com/building-safe-artificial-intelligence-52f5f75058f1>
71. Pedersen, T., Johansen, C.: Behavioural artificial intelligence: an agenda for systematic empirical studies of artificial inference. *AI Soc.* **35**(3), 519–532 (2020)
72. Perez, E., Ringer, S., Lukošiuė, K., Nguyen, K., Chen, E., Heiner, S., Pettit, C., Olsson, C., Kundu, S., Kadavath, S, Jones, A. (2022). "Discovering Language Model Behaviors with Model-Written Evaluations." arXiv preprint [arXiv:2212.09251](https://arxiv.org/abs/2212.09251)
73. Power, A., Burda, Y., Edwards, H., Babuschkin, I., Misra, V. (2022). "Grokking: Generalization beyond overfitting on small algorithmic datasets." arXiv preprint [arXiv:2201.02177](https://arxiv.org/abs/2201.02177)
74. Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.F., Breazeal, C., Crandall, J.W., Christakis, N.A., Couzin, I.D., Jackson, M.O., Jennings, N.R. (2022). "Machine behaviour." *Mach. Learn. City Appl. Archit. Urban Des.* 143–166
75. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B, Smith-Loud, J., Theron, D., Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. Proceedings of the 2020 conference on fairness, accountability, and transparency
76. Ramamoorthy, A., Yampolskiy, R.: Beyond Mad?: The Race for Artificial General Intelligence. *ITU J. ICT Discov.* **1**, 77–84 (2017)
77. Rudin, C.: Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intel.* **1**(5), 206–215 (2019)
78. Sandberg, A.: There is plenty of time at the bottom: The economics, risk and ethics of time compression. *Foresight* **21**(1), 84–99 (2019)
79. Saunders, W., Sastry, G., Stuhlmüller, A., Evans, O. (2017). "Trial without error: Towards safe reinforcement learning via human intervention." arXiv preprint [arXiv:1707.05173](https://arxiv.org/abs/1707.05173)
80. Schaeffer, R., Miranda, B., Koyejo, S (2023). "Are Emergent Abilities of Large Language Models a Mirage?" arXiv preprint [arXiv:2304.15004](https://arxiv.org/abs/2304.15004)
81. Schneider, J., Breiting, F. (2020). "AI Forensics: Did the Artificial Intelligence System Do It? Why?" arXiv preprint [arXiv:2005.13635](https://arxiv.org/abs/2005.13635)
82. Schwarting, M., Burton, T., Yampolskiy, R. (2015). On the Obfuscation of Image Sensor Fingerprints. Information and Computer Technology (GOICT), 2015 Annual Global Online Conference on, IEEE
83. Scott, P.J., Yampolskiy, R.V.: Classification schemas for artificial intelligence failures. *Delphi* **2**, 186 (2019)
84. Shavit, Y. (2023). "What does it take to catch a Chinchilla? Verifying Rules on Large-Scale Neural Network Training via Compute Monitoring." arXiv preprint [arXiv:2303.11341](https://arxiv.org/abs/2303.11341).
85. Shevlane, T., Farquhar, S., Garfinkel, B., Phuong, M., Whittlestone, J., Leung, J., D., Kokotajlo, D, Marchal, N., Anderljung, M., Kolt, N., Ho, L. (2023). Model evaluation for extreme risks. Available at: <https://arxiv.org/pdf/2305.15324.pdf>
86. Solanki, P., Grundy, J., Hussain, W.: Operationalising ethics in artificial intelligence for healthcare: a framework for AI developers. *AI and Ethics* **3**(1), 223–240 (2023)
87. Sotola, K.: Advantages of artificial intelligences, uploads, and digital minds. *Int. J. Mach. Conscious.* **4**(01), 275–291 (2012)
88. Srivastava, A., Rastogi, A., Rao, A., Shoeb, A.A., Abid, A., Fisch, A., Brown, A.R., Santoro, A, Gupta, A, Garriga-Alonso, A. (2022). "Beyond the imitation game: Quantifying and extrapolating the capabilities of language models." arXiv preprint [arXiv:2206.04615](https://arxiv.org/abs/2206.04615)
89. Tabassi, E. (2023). "Artificial Intelligence Risk Management Framework (AI RMF 1.0)"
90. Tran, B., Li, J., Madry, A (2018). "Spectral signatures in backdoor attacks." *Advances in neural information processing systems* 31

91. Trazzi, M., Yampolskiy, R.V.: Artificial stupidity: data we need to make machines our equals. *Patterns* **1**(2), 100021 (2020)
92. Turchin, A. (2021). Catching Treacherous Turn: A Model of the Multilevel AI Boxing. Available at: https://www.researchgate.net/profile/Alexey-Turchin/publication/352569372_Catching_Treacherous_Turn_A_Model_of_the_Multilevel_AI_Boxing.
93. van Leeuwen, J., Wiedermann, J. (2021). "Impossibility results for the online verification of ethical and legal behaviour of robots." Utrecht University, Utrecht, UU-PCS-2021-02.
94. Vance, A. (May 15, 2022). Is AI Progress Impossible To Predict? Available at: <https://www.lesswrong.com/posts/G993PFTwqqdQv4eTg/is-ai-progress-impossible-to-predict>.
95. Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Fedus, W. (2022). "Emergent abilities of large language models." arXiv preprint [arXiv:2206.07682](https://arxiv.org/abs/2206.07682)
96. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D. (2022). "Chain of thought prompting elicits reasoning in large language models." arXiv preprint [arXiv:2201.11903](https://arxiv.org/abs/2201.11903)
97. Werkhoven, P., Kester, L., Neerinx, M (2018). Telling autonomous systems what to do. Proceedings of the 36th European Conference on Cognitive Ergonomics
98. Whittlestone, J., Clark J. (2021). "Why and How Governments Should Monitor AI Development." arXiv preprint [arXiv:2108.12427](https://arxiv.org/abs/2108.12427)
99. Wiener, N.: Some moral and technical consequences of automation. *Science* **131**(3410), 1355–1358 (1960)
100. Wolfram, S. (2002). A new kind of science, Wolfram media Campaign
101. Woolgar, S.: Why not a sociology of machines? The case of sociology and artificial intelligence. *Sociology* **19**(4), 557–572 (1985)
102. Wu, F. F. and A. Monticelli (1985). "Network observability: theory." *IEEE Transactions on Power Apparatus and Systems*(5): 1042–1048.
103. Yampolskiy, R.: The singularity may be near. *Information* **9**(8), 190 (2018)
104. Yampolskiy, R. (2020). On controllability of artificial intelligence. IJCAI-21 Workshop on Artificial Intelligence Safety (AISafety2021)
105. Yampolskiy, R. (2021). On the Differences between Human and Machine Intelligence. AISafety@ IJCAI
106. Yampolskiy, R. (2023). "Unmonitorability of Artificial Intelligence." <https://philpapers.org/rec/YAMUOA-3>
107. Yampolskiy, R.V.: Behavioral modeling: an overview. *Am. J. Appl. Sci.* **5**(5), 496–503 (2008)
108. Yampolskiy, R.V.: Leakproofing singularity - artificial intelligence confinement problem. *J. Conscious. Stud. (JCS)* **19**(1–2), 194–214 (2012)
109. Yampolskiy, R. V. (2015). Analysis of types of self-improving software. Artificial General Intelligence: 8th International Conference, AGI 2015, AGI 2015, Berlin, Germany, July 22–25, 2015, Proceedings 8, Springer.
110. Yampolskiy, R.V.: What are the ultimate limits to computational techniques: verifier theory and unverifiability. *Phys. Scr.* **92**(9), 093001 (2017)
111. Yampolskiy, R.V.: Artificial consciousness: an illusionary solution to the hard problem. *Reti, saperi, Linguaggi* **2**, 287–318 (2018)
112. Yampolskiy, R.V.: Unexplainability and Incomprehensibility of AI. *J. AI. Consci* **7**(2), 277–291 (2020)
113. Yampolskiy, R.V.: Unpredictability of AI: on the impossibility of accurately predicting All actions of a smarter agent. *J. Artif. Intell. Conscious.* **7**(1), 109–118 (2020)
114. Yampolskiy, R. V. (2022). AI Risk Skepticism. *Philosophy and Theory of Artificial Intelligence 2021*, Springer: 9:225–248
115. Yampolskiy, R.V.: Metaverse: a solution to the multi-agent value alignment problem. *J. Artif. Intell. Conscious.* **9**, 1–11 (2022)
116. Yampolskiy, R.V.: On the controllability of artificial intelligence: an analysis of limitations. *J. Cyber Sec. Mob.* **11**, 321–404 (2022)
117. Yampolskiy, R. V. (2023). Ownability of AGI. Artificial General Intelligence: 15th International Conference, AGI 2022, Seattle, WA, USA, August 19–22, 2022, Proceedings, Springer
118. Yampolskiy, R.V.: to appear). Unexplainable, Uncontrollable, CRC Press, AI - Unpredictable (2024)
119. Yampolskiy, R.V., Ashby, L., Hassan, L.: Wisdom of artificial crowds—a metaheuristic algorithm for optimization. *ILSA* (2012). <https://doi.org/10.4236/jilsa.2012.42009>
120. Yampolskiy, R.V., Govindaraju, V.: Use of behavioral biometrics in intrusion detection and online gaming. *SPIE, Biometric Technology for Human Identification III* (2006)
121. Yudkowsky, E. (May 22, 2008). That Alien Message. Less Wrong. Available at: <https://www.lesswrong.com/posts/5wMcKNAwB6X4mp9og/that-alien-message>.
122. Yudkowsky, E. (October 13, 2017). There's No Fire Alarm for Artificial General Intelligence. Available at: <https://intelligence.org/2017/10/13/fire-alarm/>.
123. Yudkowsky, E., Hanson R. (2008). The Hanson-Yudkowsky AI-foom debate. MIRI Technical Report. Available at: <http://intelligence.org/files/AIFoomDebate.pdf>.
124. Zhang, D., Mishra, S., Brynjolfsson, E., Etchemendy, J., Ganguli, D., Grosz, B., Lyons, T., Manyika, J., Niebles, J.C., Sellitto, M., Shoham, Y. (2021). "The AI index 2021 annual report." arXiv preprint [arXiv:2103.06312](https://arxiv.org/abs/2103.06312)
125. Zhou, D, Schärli, N, Hou, L, Wei, J., Scales, N., Wang, X., Schuurmans, D., Cui, C., Bousquet, O., Le Q., Chi, E. (2022). "Least-to-most prompting enables complex reasoning in large language models." arXiv preprint [arXiv:2205.10625](https://arxiv.org/abs/2205.10625)
126. Ziesche, S., Yampolskiy, R.V. (2020). "Towards the Mathematics of Intelligence." *The Age of Artificial Intelligence: An Exploration* **1**
127. Zoe Cremer, C., Whittlestone, J.: Artificial canaries: early warning signs for anticipatory and democratic governance of AI. *IJMAI* (2021). <https://doi.org/10.9781/ijimai.2021.02.011>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.