

Essay

Not peer-reviewed version

---

# A Causal Framework for AI Regulation and Auditing

---

Lee Sharkey<sup>\*</sup>, Clíodhna Ní Ghuidhir, Dan Braun, Jérémy Scheurer, Mikita Balesni, Lucius Bushnaq, Charlotte Stix, Marius Hobbhahn

Posted Date: 18 January 2024

doi: 10.20944/preprints202401.1424.v1

Keywords: AI governance; AI evaluations; AI regulation; AI auditing



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

## Article

# A Causal Framework for AI Regulation and Auditing

Lee Sharkey \*, Clíodhna Ní Ghuidhir, Dan Braun, Jérémy Scheurer, Mikita Balesni,  
Lucius Bushnaq, Charlotte Stix and Marius Hobbhahn

Apollo Research, London, UK

\* Correspondence: lee@apolloresearch.ai

**Abstract:** Artificial intelligence (AI) systems are poised to become deeply integrated into society. If developed responsibly, AI has potential to benefit humanity immensely. However, it also poses a range of risks, including risks of catastrophic accidents. It is crucial that we develop oversight mechanisms that prevent harm. This article outlines a framework for evaluating and auditing AI to provide assurance of responsible development and deployment, focusing on catastrophic risks. We argue that responsible AI development requires comprehensive auditing that is proportional to AI systems' capabilities and available affordances. This framework offers recommendations toward that goal and may be useful in the design of AI auditing and governance regimes.

**Keywords:** AI governance; AI evaluations; AI regulation; AI auditing

## Executive Summary

This article presents a causal framework for the evaluation and auditing of AI, with a particular focus on mitigating catastrophic risks. Our approach emphasizes the importance of degrees of auditing that are proportional to the capabilities and affordances of AI systems and aims to ensure their safe development and deployment.

Our main contributions are

- **A causal framework:** Our framework works backwards through the causal chain that leads to the *effects that AI systems have on the world* and discusses ways auditors may work toward assurances at each step in the chain.
- **Conceptual clarity:** We develop several distinctions that are useful in describing the chain of causality. Conceptual clarity should lead to better governance.
- **Highlighting the importance of AI systems' available affordances:** We identify a key node in the causal chain - the affordances available to AI systems - which may be useful in designing regulation. The affordances available to AI systems are the environmental resources and opportunities for affecting the world that are available to it, e.g. whether it has access to the internet. These determine which capabilities the system can currently exercise. They can be constrained through guardrails, staged deployment, prompt filtering, safety requirements for open sourcing, and effective security. One of our key policy recommendations is that proposals to change the affordances available to an AI system should undergo auditing.

We outline the causal chain leading to AI systems' effects on the world below, working backwards from the real world effects to inputs and determinants:

- **AI system behaviors** - The set of actions or outputs that a system actually produces and the context in which they occur (for example, the type of prompt that elicits the behavior).
- **Available affordances** - The environmental resources and opportunities for affecting the world that are available to an AI system.
- **Absolute capabilities and propensities** - The full set of potential behaviors that an AI system can exhibit and its tendency to exhibit them.
- **Mechanistic structure of the AI system during and after training** - The structure of the function that the AI system implements, comprising architecture, parameters, and inputs.

- **Learning** - The processes by which AI systems develop mechanistic structures that are able to exhibit intelligent-seeming behavior.
- **Effective compute and training data content** - The amount of compute used to train an AI system and the effectiveness of the algorithms used in training; and the content of the data used to train an AI system.
- **Security** - Adequate information security, physical security, and response protocols.
- **Deployment design** - The design decisions that determine how an AI system will be deployed, including who has access to what functions of the AI system and when they have access.
- **Training-experiment design** - The design decisions that determine the procedure by which an AI system is trained.
- **Governance and institutions** - The governance landscape in which AI training-experiment and security decisions are made, including of institutions, regulations, standards, and norms.

We identify and discuss five audit categories, each aiming to provide assurances on different determinants:

- AI system evaluations
- Security audits
- Deployment audits
- Training design audits
- Governance audits.

We highlight key research directions that will be useful for designing an effective AI auditing regime. High priority research questions include interpretability; predictive models of capabilities and alignment; structured access; and potential barriers to transparency of AI labs to regulators.

## 1. Introduction

AI has the potential to influence the lives of the public both positively [1,2] and negatively [3–5], including catastrophically [6,7]. As frontier AI systems become increasingly integrated into our lives and the economy, it grows ever more important to have assurances regarding the effects of these systems. Such assurances may be possible through auditing frontier AI systems and the processes involved in their development and deployment.

There are growing calls for regulations and audits of the social and technical processes that generate frontier AI systems. Auditing complements regulation; it helps to ensure that actors adhere to regulations and to identify risks before their realization. But the design and implementation of regulations for frontier AI are in their early days. How can we design a framework for auditing when most of the relevant regulations do not yet exist? Although we do not yet know which regulations will be implemented, many proposed measures already have broad support [8] and therefore have a reasonable probability of being implemented. **Here, we discuss what a governance regime for frontier AI systems could look like and introduce a framework for the roles that auditors could play within it.** The framework focuses on extreme risks from frontier AI systems, rather than nearer term, smaller scale risks (such as bias or fairness) or risks from systems with narrower applications (such as lethal autonomous weapons). It emphasizes safety and assurance over other considerations, such as costs, political viability, or current technological barriers. Nevertheless, we believe that all regulations and auditor roles proposed here could and should be implemented in practice. It thus serves as a comprehensive menu of options for auditing in an AI governance regime and helps to identify areas for further research.

By ‘auditing’, we do not exclusively mean ‘compliance auditing’, where auditors ensure that auditees have complied with a well-defined set of acceptable practices, or where systems comply with a well-defined set of regulations. Although compliance audits are an ideal regime when possible, the science of AI safety is currently nascent enough that, in many areas of AI governance, compliance audits are at present inappropriate since standards are still emerging. Therefore, by ‘auditing’ we also

include evaluations, which include risk assessments or assessments of whether AI systems possess certain dangerous capabilities. Evaluations may be one of the checks performed during auditing, though not necessarily compliance auditing.

Our work builds on similar previous work [9–12], though it differs in the following ways:

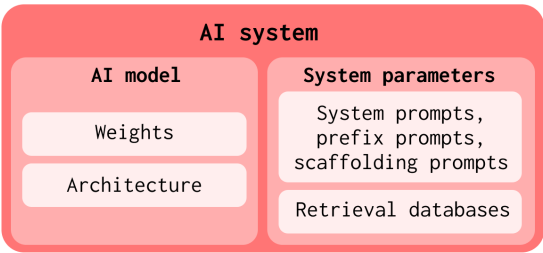
- **Causal framework:** Our framework begins with the target of our assurance efforts - *the effects that AI systems have on the world* - and works backwards through the causal process that leads to them. This provides an overarching framework for conceptualizing the governance of general-purpose AI systems. We hope that a framework will help to highlight potential regulatory blindspots and ensure comprehensiveness.
- **Conceptual engineering for clarity:** In developing the framework, we found we needed to use distinctions that we had not previously encountered. For instance, in Section 2.2 we distinguish between absolute, contextual, and reachable capabilities, which we found useful for thinking about how AI systems interact with their environment and, hence, about regulations concerning their eventual effects on the world. We also reframe the focus from AI models to the slightly broader concept of AI systems (Section 2.1).
- **Highlighting the importance of AI systems' available affordances:** One of the benefits of the technical focus of the framework is that it highlights a key concept, the affordances available to AI systems (defined in Section 2.2) as a key variable for regulations, and serves as a unifying frame for many related concepts such as AI system deployment, open sourcing, access to code interpreters, guardrails, and more. The concept of available affordances adds important nuance to the 'training phase-deployment phase' distinction, a distinction that arguably formed the keystone of previous frameworks.

Our framework is agnostic to who should do the auditing. In general, however, we believe that *external* auditors, as well as auditors who are internal to AI labs, should be empowered to perform audits at all stages of AI system development. Decisions made prior to and during training ultimately affect the final AI system and therefore directly (and potentially radically) affect the public. Both internal and external audits are necessary for ensuring public accountability and to ensure that the incentives of auditors are aligned with public interests [12–14]. In addition to ensuring appropriate incentives, additional external auditing is also safer than relying solely on internal audits, since an absence of extensive auditing may reduce the possibility of discovering, and thus neutralizing, dangerous capabilities in AI systems.

## 2. Our framework: Auditing the determinants of AI systems' effects

### 2.1. Conceptualizing AI systems vs. AI models

The objects of our framework, *AI systems*, are a slight generalization of AI models. AI systems include not only the weights and architecture of the model, but also include a broader set of system parameters (Figure 1). These consist of retrieval databases and particular kinds of prompts. The reason for our expanded focus is that system parameters strongly influence the capability profile of AI systems. Therefore, systems with one set of system parameters may be safe to deploy whereas others are not.



**Figure 1.** An AI system comprises an AI model, with its weights and architecture, as well as various other system parameters, including system prompts, prefix prompts, scaffolding prompts, and retrieval databases.

We discuss two main types of system parameter:

1. *System prompts, prefix prompts, and scaffolding*: System prompts and prefix prompts are prefixed to any user-generated prompts that are input to AI models (Section 2.3.4). These prompts substantially modify system behavior. For instance: A model may not itself possess cyber offensive capabilities. But if part of its system prompt contains documentation for cyber offense software, then, through in-context learning (Section 2.3.4), the system may exhibit dangerous cyber offense capabilities. This materially affects which systems are safe to deploy. If a lab wants to change the system prompt of a deployed model, it could change the capabilities profile enough to warrant repeated audits. Similarly, AI systems with few-shot prompts or chain-of-thought prompts (examples of *prefix prompts*) are more capable of answering questions than without those prompts that prefix their outputs. It is possible to arrange systems of models such that outputs of some models are programmatically input to other models (e.g. [15–17]). These systems, which often involve using different *scaffolding prompts* in each of the models in the system, have different capability profiles than systems without scaffolding prompts.
2. *Retrieval databases*: Retrieval databases (Section 2.3.4) may sometimes be considered parameters of the model [18] and other times parameters of the AI system (of which the model is a part) [19]. In both cases, the retrieval databases meaningfully influence the capabilities of the system, which may influence the outcomes of evaluations. We therefore expand the focus of what we consider important to evaluate in order to include all factors that influence the capabilities profile of the system.

AI models are still AI systems. They are simply a special case of AI systems that have only model parameters and no other AI system parameters. Given the relevance of learning and system parameters to capabilities and, hence, to auditing, we found it made most sense to place AI systems at the center of our framework, rather than AI models.

2.2. Conceptualizing AI system capabilities

In addition to shifting our focus from AI models to AI systems, we also develop some additional clarifying terminology regarding capabilities.

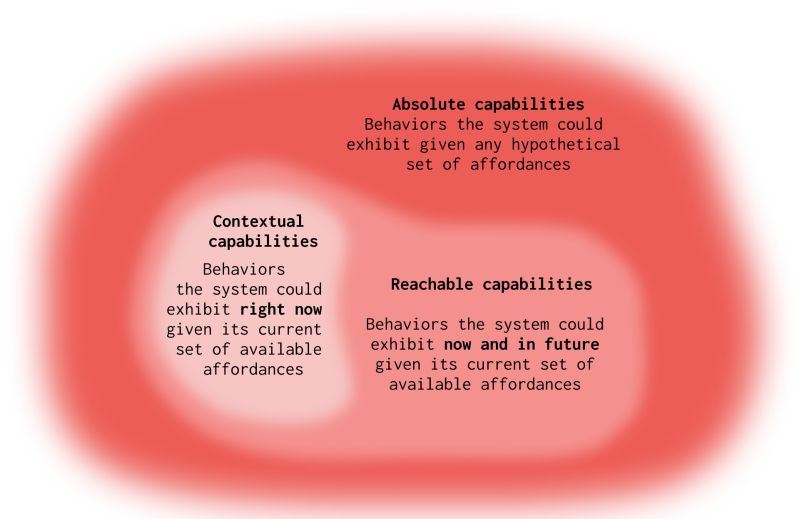
We use the following terms:

- **AI system behaviors.** The set of actions or outputs that a system actually produces and the context in which they occur (for example, the type of prompt that elicits the behavior). In this article, we exclusively mean behaviors that could be useful to the system or a user, rather than arbitrary or random behaviors that might be produced by an untrained system.
- **Available affordances.** The environmental resources and opportunities for influencing the world that are available to a system. We can exert control over the affordances available to a system through deployment decisions, training design decisions, guardrails, and security measures. As we’ll see below, the available affordances determine which of a system’s *absolute capabilities*

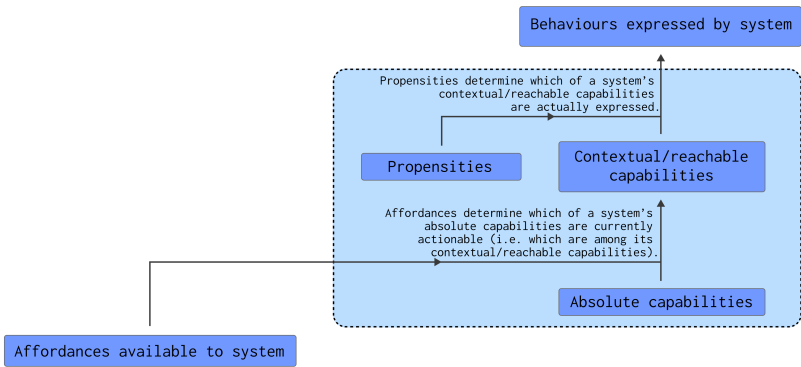
are its *contextual capabilities and reachable capabilities* (Figure 3). Systems with greater available affordances have constant absolute capabilities but greater contextual and reachable capabilities (Figure 5).

- **Absolute capabilities.** The potential set of behaviors that a system could exhibit given any hypothetical set of available affordances (regardless of whether the system can reach them from its current context) (Figure 2).
  - *Example:* If a trained system is saved in cold storage and not loaded, it is not contextually or reachably capable of behaving as a chat bot, even if it is absolutely capable of doing so. To become contextually or reachably capable of behaving as a chat bot, it must be loaded from storage and used for inference.
- **Contextual capabilities.** The potential set of behaviors that a system could exhibit right now given its current set of available affordances in its current environmental context (Figure 2).
  - *Example:* A system may be absolutely capable of browsing the web, but if it does not have access to the internet (i.e., an available affordance), then it is not contextually capable of browsing the web.
- **Reachable capabilities.** The potential set of behaviors that a system could exhibit given its current set of available affordances (i.e. contextual capabilities) as well as all the affordances the system could eventually make available from its current environmental context. A system's reachable capabilities include and directly follow from a system's contextual capabilities (Figure 2).
  - *Example:* A system, such as GPT4, may not be able to add two six-digit numbers. Therefore, six-digit addition is not within its contextual capabilities. However, if it can browse the web, it could navigate to a calculator app to successfully add the numbers. Therefore, six-digit addition is within its reachable capabilities.
- **System propensities.** The tendency of a system to express one behavior over another (Figure 3). Even though systems may be capable of a wide range of behaviors, they may have a tendency to express only particular behaviors. Just because a system may be capable of a dangerous behavior, it might not be inclined to exhibit it. It is therefore important that audits assess a system's propensities using 'alignment evaluations' [9].
  - *Example:* Instead of responding to user requests to produce potentially harmful or discriminatory content, some language models, such as GPT-3.5, usually respond with a polite refusal to produce such content. This happens even though the system is capable of producing such content, as demonstrated when the system is 'jailbroken'. We say that such systems have a propensity not to produce harmful or offensive content.





**Figure 2.** The relationship between the sets of potential behaviors defined by absolute capabilities, reachable capabilities, and contextual capabilities.



**Figure 3.** The relationship between an AI system’s capabilities, propensities, affordances, and behaviors.

It is worth noting that the ‘capabilities’ of a system are multidimensional in a similar sense that ‘intelligence’ is multidimensional. A ‘less capable’ system might have learned some abilities that a ‘more capable’ system has not. However, when comparing two general cognitive systems, we rely on the intuition that a ‘more capable’ system has learned almost all and more of the abilities that ‘less capable’ has learned.

It is also worth noting that training has an effect on all classes of capabilities (absolute, contextual, reachable), but some more than others. For example, if a large language model is being trained offline on a large fixed corpus of text (i.e. it is not learning online by, for instance, browsing the web and updating its weights frequently during browsing), then during training its contextual capabilities remain quite limited - they are limited only to outputting a prediction for the current minibatch of data. The system’s reachable capabilities remain similarly limited. However, its absolute capabilities may increase a lot throughout the training process. This is because absolute capabilities do not depend on a system’s current set of available affordances. It is important to design training-experiments that do not introduce risks from excessive absolute capabilities (such as a training-experiment that scales an AI system 1000x larger before understanding similar systems’ capabilities at smaller scales) or excessive affordances (such as giving an AI system root access to a supercomputing cluster or unconstrained access to the internet).

2.3. The determinants of AI systems’ effects on the world

The ultimate purpose of auditing for AI safety is to reduce the risk of AI systems having damaging effects on the world. In particular, the role of auditing is to identify sources of risk before the damage materializes. It is therefore important to audit the *upstream determinants of AI systems’ effects on the world* (Figure 4).

AI systems’ effects on the world are determined by the behaviors they express, which are in turn determined by:

- 1. Affordances available to the system
- 2. Absolute capabilities and propensities of the system during and after training
- 3. Mechanistic structure of the system during and after training
- 4. Learning
- 5. Effective compute and training data content
- 6. Security
- 7. Deployment design
- 8. Training-experiment design
- 9. Governance and institutions

We identify five main kinds of audit, each aiming to provide assurances on different determinants (Figure 4). These are:

- 1. AI system evaluations
- 2. Security audits
- 3. Deployment audits
- 4. Training design audits
- 5. Governance audits.

In the sections below, we consider each determinant in turn and explore auditors’ potential roles with respect to each.

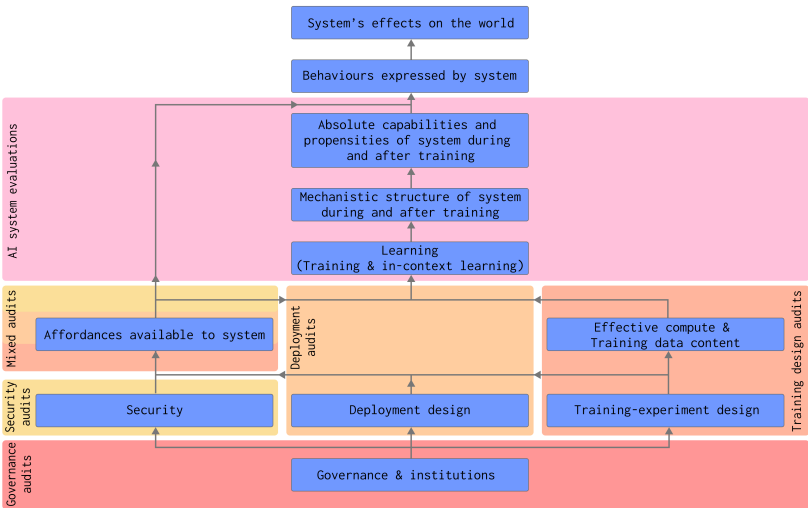



Figure 4. Determinants of AI system’s effects on the world and the types of auditing that act on them.

2.3.1. Affordances available to a system

The affordances available to an AI system determine which of its absolute capabilities are among its contextual and reachable capabilities. Systems with greater available affordances have constant absolute capabilities but greater contextual and reachable capabilities (Figure 5). The potential harm posed by an AI system can thus increase with increasing available affordances. **One of our key**



recommendations is that decisions that expand the affordances available to AI systems should be subject to auditing.



**Figure 5.** The relationship between absolute capabilities, affordances, contextual and reachable capabilities, and the level of auditing warranted. Absolute capabilities and available affordances are orthogonal. As either increase, the level of auditing required also increases.

The affordances available to AI systems are determined by a number of factors, which we'll look at in turn. They include:

1. Extent of system distribution
2. Actuators, interfaces, and scaffolding
3. Online vs offline training
4. Guardrails

#### Extent of system distribution

This is the number of people the AI system interacts with and the extent to which it is integrated into systems in the wider world. This includes factors such as deployment, open sourcing, and other forms of AI system proliferation.


'Extent of system distribution' closely reflects the training vs. deployment distinction emphasized in other auditing frameworks [9,12]. However, it is an imperfect reflection, since AI systems may potentially have access to dangerous levels of affordances during training, for instance during online training. The 'training vs deployment' distinction is only relevant for risk because it reflects, albeit imperfectly, the affordances available to a system. Thus, while not entirely disposing of the training versus deployment distinction, available affordances are more important in our framework because they are more direct causes of risk.

When a developer proposes deployment of an AI system, it should be subject to a deployment audit because it involves a change to the system's available affordances. Sudden jumps in the extent of system distribution available to systems represent sudden jumps in contextual and reachable capabilities. This may lead to unexpected adverse outcomes. For example, if a system with the capability to instruct users on how to build bioweapons is open sourced, then it may be usable by malicious actors to devastating effect. Staged deployment, first to trusted AI system evaluations researchers (for instance) and finally to the broader public, may help identify and mitigate these risks while access is expanded. In other words, evaluations researchers may help estimate the absolute capabilities of the system before very large affordances (and hence potentially large contextual and reachable capabilities) are made available to them. At each stage, *deployment audits* (Figure 4) should

be required before a new deployment stage begins. In a similar way, AI systems should not be open-sourced unless it can be demonstrated (through a deployment audit) that the risks are sufficiently small.

Besides intentional deployment, another way that the extent of system distribution may increase is unintentional proliferation, e.g. through leaked or stolen model weights [20] or, in highly capable future AI systems, the system exfiltrating itself. *Security audits* (Figure 4) should help reduce the risks from system theft or autoexfiltration. The weights of highly absolutely capable systems should also be tracked, such that the locations of all copies are known and secured.

Expensive AI systems with extensive absolute capabilities are likely to be targets of theft. They are also more likely than less capable systems to be capable of autoexfiltration. Therefore, the risk of proliferation may increase as absolute capabilities increase, and therefore also increases the risk of unintended levels of affordances becoming available to the system. This risk means that absolute capabilities and available affordances may become positively correlated at high levels of absolute capabilities. Once systems proliferate widely, protecting society from potential harms becomes exceptionally, if not impossibly, difficult [4]. We should therefore be careful about developing highly absolutely capable systems at all.



**Figure 6.** The effects of different types of training and available affordances on capabilities. While changes in available affordances increase contextual and reachable capabilities, they don't change absolute capabilities.

Actuators, interfaces, and scaffolding

AI system outputs can be input to specialized interfaces. These may be physical, such as monitors (and the people looking at them) and robotic actuators (e.g. [21,22]). But they may also be interfaces to software systems, such as code interpreters [23] or actions in games [24,25], or even as inputs to other language models.

Systems can be furnished with plugins or tools [26], where systems are given access to interfaces and are provided with prompts (or potentially fine-tuning [27]) that teach them how to use those interfaces. Plugins serve to facilitate a system's interactions with a wider range of interfaces and actuators. For example, an AI system without the right plugin may not easily be able to interact with the internet. But when given access to a browsing interface and prompts that inform the system on how to use the interface, the system may be able to browse the internet effectively.

Similar to plugins is scaffolding. *Scaffolding programs* involve passing the outputs of one AI system to the inputs of other system instances, sometimes with additional scaffolding prompts that guide models on how to use the outputs of the other models. Scaffolding programs influence the

available affordances of AI systems used within the scaffolding, whereas *scaffolding prompts* shape their capabilities. One example of scaffolding is 'Tree of Thought' scaffolding [15], which uses multiple language models scaffolded together to answer questions; one AI system produces multiple candidate intermediate steps to answer a problem, and a second system evaluates those candidates. The scaffolded systems move closer to answer the question by advancing along only the most promising intermediate steps. Other frameworks exist, such as LangChain or AutoGPT, which give systems the ability to behave like autonomous agents. Today's AI systems do not appear to be able to make maximum use of scaffolding due to reliability issues. Nevertheless, near-term future systems with marginally greater absolute capabilities might be reliable enough to make proficient use of these affordances, which has potential to vastly increase their contextual and reachable capabilities.

AI systems may be given different actuators, interfaces, or scaffolding at different points in training and deployment. This means that both deployment audits and *training design audits* (Figure 4) should be used in their oversight. Security audits may also be involved in ensuring that AI systems do not have access to more actuators or interfaces than intended.

### Online vs offline training

In *offline training*, the training dataset is collected and fixed prior to training. Systems have limited affordances available to them during offline training since their outputs mostly do not affect anything in the world. Nor can they accrue more affordances during training. Therefore the system's contextual and reachable capabilities remain roughly constant, although their absolute capabilities increase.

By contrast, *online training* involves training on data that are collected during training. The training data result from the outputs of the AI system that is being trained and their interaction with its (optional) environment. Systems trained online usually interact with an environment, and those interactions require that the environment provide the system with some affordances. If that environment is the real world, such as a chatbot interacting with humans, it elevates risks. Online training therefore shares similarities with deployment: In both cases, AI systems have extensive affordances available and may have real world consequences. This makes online training potentially riskier. Decisions regarding online training should therefore be audited during training design audits.

### Extent of guardrails

We can potentially limit the affordances available to AI systems by introducing guardrails. For example, attempts were made - with qualified success [28] - to prevent the AI system Stable Diffusion from outputting sexually explicit images by applying a safety filter to the output of the image generation system. In the language model case, it may be possible to curtail certain kinds of capabilities by e.g. filtering their inputs (using another language model) such that any prompt that has the system do chain-of-thought reasoning is prevented from being input to the AI system. It may also be possible to curtail the acquisition of new affordances through guardrails. An advanced AutoGPT-like system could, for instance, have guardrails that require the AI system to request human permission to access the internet; to use financial resources; to create new system instances; or to run code. The job of ensuring that appropriate guardrails are in place should fall on either deployment audits or, if guardrails will be imposed during training, training design audits.

### Summary of recommendations regarding affordances available to AI systems

- Proposed changes in the affordances available to a system (including changes to the extent of a system's distribution, online/offline training, actuators, interfaces, plugins, scaffolding, and guardrails) should undergo auditing, including conduct risk assessments, scenario planning, and evaluations.
- AI systems' deployment should be staged such that distribution increases in the next stage only if it is deemed safe.
- Model parameters should not be open sourced unless they can be demonstrated to be safe.

- All copies of highly absolutely capable models should be tracked and secured.
- Guardrails should be in place to constrain the affordances available to AI systems.

### 2.3.2. Absolute capabilities and propensities of a system

One of the goals of *AI system evaluations* (Figure 4) is to assess whether a system will exhibit any dangerous behaviors when certain affordances are made available to it. When systems are deployed, users could for instance provide the system with novel scaffolding. For sufficiently absolutely capable systems, these affordances may give the system potentially dangerous reachable capabilities. By looking for these dangerous capabilities in more controlled environments, AI systems evaluations can help us better predict the risks posed by deployment.

When evaluating a system, we should assess whether the system possesses certain capabilities ('dangerous capability evaluations') and whether it has the propensity to exhibit them ('alignment evaluations') [9]. These should be done both during and after training to ensure system developers and auditors understand the current capabilities and propensities at any training checkpoint and to respond to trends in how systems are being used. Certain kinds of training algorithms may result in systems with different absolute capability and propensity profiles; for instance, systems trained with RL algorithms may potentially result in more dangerous systems than systems trained with other algorithms [29].

Attempts to elicit dangerous behaviors in a controlled setting may involve prompt engineering [30] or prompt optimization [31–33]. Fine-tuning methods, such as prefix tuning [34] or standard fine-tuning (for instance, using RLHF [35]), may serve as approximate upper bounds on what can be achieved by prompting alone. Evaluating the capabilities of fine-tuned systems can also be used to inform the decision of whether to deploy fine-tuning access to users.

It is likely the case that dangerous capability evaluations are easier than alignment evaluations; to demonstrate that an AI system can exhibit the dangerous capability in any setting; to demonstrate an AI system's propensity to exhibit a dangerous capability, it must be shown whether or not the system exhibits it in all (or at least a representative sample) of settings. This seems much harder to do, since the space of possible settings is very large.

It is important to be clear that, like AI system training, dangerous capability evaluations and alignment evaluations can be 'gain-of-function' work, where dangerous capabilities are elicited in controlled settings. As systems become closer to being able to, for instance, autonomously self-replicate or deceive evaluators, the risk of these experiments increases. At that point, auditors must therefore be careful to avoid introducing new risks through their work. Eventually, before such systems are developed, there should be

1. Risk assessments prior to certain gain-of-function experiments to ensure that the risks are worth the benefits.
2. Requirement of official certification to perform certain gain-of-function work, reducing risks from irresponsible or underqualified actors.
3. Information controls. Given the risk that information obtained by gain-of-function research may proliferate, there should be controls on its reporting. Auditors should be able to report gain-of-function results to regulators, who may have regulatory authority on decisions whether to prevent or reduce system deployment, and not necessarily to AI system developers or the broader scientific community.

To avoid hampering the field in its early stages, it seems sensible to implement these requirements only for models considerably larger (and therefore riskier) than today's. This said, as AI systems are trained, their absolute capabilities increase and highly absolutely capable systems may be able to make use of even very limited available affordances. For example, suppose a dangerous system had no access to the outside world other than through a single terminal to interact with researchers; if the system has sufficient absolute capabilities, it may be able to use its limited set of available affordances

to manipulate the auditors and developers into giving it access to even greater affordances, such as internet access, which may be dangerous. In other words, if a system has sufficient absolute capabilities, its limited set of available affordances may nevertheless give it dangerously large reachable capabilities even though its contextual capabilities are small. Therefore, as absolute capabilities increase, gain of function research becomes riskier, so there should be increasing restrictions on both the development of such absolutely capable systems as well as on gain-of-function research that involves such systems.

It is important that enforceable action plans are triggered following AI system evaluations that reveal danger. Unless the evaluations lead to concrete actions that mitigate risk, they may not be worth the additional risks they create. Without enforceable protocols that describe what will be done following specific concerning evaluations, such as the prevention of deployment or the pausing of training, then there is little purpose in doing these evaluations at all. Some action plans could be implemented in code, such that training does not proceed if certain evaluations reveal cause for concern.

Summary of proposed changes regarding absolute capabilities and propensities of AI systems

- Evaluate dangerous capabilities and propensities continually throughout and after training.
- When such experiments involve extremely capable AI systems, auditors should require certification to perform gain-of-function experiments, similar to AI developers; risk assessments prior to gain-of-function experiments should be required; and information controls on reporting gain-of-function results should be implemented.
- Have enforceable action plans following concerning evaluations.

### 2.3.3. Mechanistic structure of a system

The AI systems under discussion are neural networks that have learned complicated internal algorithms that implement highly capable input-output functions. The *mechanistic structure* of an AI system consists of the algorithms it implements internally. The field of *mechanistic interpretability* aims to characterize AI systems' mechanistic structure. Mechanistic structure is determined not only by internal algorithms implemented by the model (i.e. the architecture and network parameters), but also by its system parameters (e.g. system prompts, prefix prompt, or retrieval database).

Evaluations of AI system behavior alone will, unfortunately, shed limited light on a system's absolute capabilities; they can only show that certain behaviors are among a system's absolute capabilities but cannot conclusively show that certain behaviors are not. If we understood a system's mechanistic structure we could predict what it will do in particular hypothetical situations. For instance, if we observed that a system had a mechanistic structure such that it possessed no knowledge about dangerous pathogens, we would be able to predict that it wouldn't be capable of giving instructions on how to manufacture a bioweapon when asked. Understanding a system's mechanistic structure should therefore give us a better sense of its absolute capabilities, which will give us a much better understanding of a system's reachable and contextual capabilities in arbitrary settings and therefore let us make predictions about system behavior that are relevant to safety. It should also give us a better understanding of systems' propensities, we would be able to make wider inferences about what systems would tend to do in different contexts.

At present, mechanistic interpretability, the field of research that would let us understand systems' mechanistic structure, is in its infancy. Current interpretability methods seem inadequate for making confident claims about how AI systems might behave outside of the evaluation setting [36]. It will therefore be important to develop better interpretability methods and incorporate them into AI system evaluations. Nevertheless, in advance of adequate interpretability methods being found, it will be important to develop forms of structured access [37,38] for external auditors that permit the kinds of interpretability research that will be necessary to gain assurances of AI systems' safety.



### Summary of recommendations regarding the mechanistic structure of AI systems

- Do not overstate the guarantees about an AI system's absolute capabilities and propensities that can be achieved through behavioral evaluations alone.
- Incorporate interpretability into capability and alignment evaluations as soon as possible.
- Develop forms of structured access for external auditors to enable necessary research and evaluations.

### 2.3.4. Learning

Learning is the process by which AI systems develop mechanistic structures that are able to exhibit intelligent-seeming (sometimes dangerous) behavior. Here, this includes both training of model parameters by gradient descent (i.e. pretraining and fine-tuning) and other forms of optimization such as in-context learning [39,40].

#### Pretraining and fine-tuning

The typical development pipeline for general-purpose AI systems involves a long offline 'pretraining' phase on a large corpus of data. Pretraining need not be exclusively offline, but today usually is. Pretraining is often followed by multiple 'fine-tuning' phases. There are diverse approaches to fine-tuning, which have a wide variety of effects on systems' absolute capabilities including (but not limited to):

- Instruction fine-tuning [41], RLHF [35], Constitutional AI [42] drastically change a system's absolute capabilities and propensities;
- Iterative hallucination reduction, as in GPT-4 [43];
- Different training objectives that increase absolute capabilities on downstream metrics, such as UL2 training [44];
- Adding new parameters to a system and fine-tuning on particular tasks, such as in LoRA fine-tuning or prefix-tuning [34,45];
- Fine-tuning two independently trained models so they operate as one model, such as Flamingo, which combines a vision model with a language model [46];
- Enabling a model to use particular kinds of retrieval databases [18,47];

The main differences between pretraining and fine-tuning is that pretraining comes first and the number of updates applied to the model parameters is usually much larger than in fine-tuning. The relevant aspect of pretraining vs fine-tuning with respect to auditing is how much they contribute to system absolute capabilities and propensities. Audits should be required according to the level of absolute capability a particular training phase may add, regardless of number of updates. Although some fine-tuning updates may have only small changes to the absolute capabilities of a system, some fine-tuning approaches may contribute significant absolute capabilities with only small numbers of updates. For example, using the UL2R mixture of training objectives to finetune a language model, it is possible to drastically improve performance on downstream tasks using only 0.1% to 1% of the pretraining computation costs [44]. Fine-tuned systems may therefore necessitate significant additional auditing.

AI systems evaluations during training will be important for developing good predictive models of capabilities, which we currently lack [48–50]. By modeling patterns and phase changes in a system's capability profile, we would like to be able to predict when the capacity for certain dangerous behaviors will emerge. This may let us avoid training AI systems capable of particularly dangerous behaviors, such as strategic deception, which may render other evaluations useless or have catastrophic outcomes.

#### Learning from prompts and retrieval databases

A remarkable property of large language models is their capacity for in-context learning [39,40,51–54]. For example, when provided with in-context demonstrations of how to perform a task, a system



that can't solve a task zero-shot may become able to solve the task after being provided with a few demonstrations [51]. The absolute capabilities of an AI system with one prompt are different than with a different prompt; this is because the structure of the function they implement is different thanks to different system parameters, even though the weights of the model stay constant.

While the ability to learn new capabilities in-context is one of the most useful properties of language models, it is also one of the main sources of risk. For instance, a system may in-context-learn new cyber offense skills from documentation for cyber offense software. As discussed in Section 2.3.1, plugins and scaffolding often employ prompts to teach AI systems how to use software interfaces or operate within scaffolding programs, giving them very different capability profiles. Therefore, AI systems evaluations are needed to test the extent to which systems can learn dangerous capabilities in-context. This includes the ability to behave as an autonomous agent using scaffolding such as LangChain [55] or AutoGPT [17].

As introduced in Section 2.1 the other kind of system parameter is retrieval databases, which make use of or influence in-context learning [18,19,56–59]. As with prompts, changes to a model's retrieval database could change the capabilities profile of the model. To reuse the same example, adding information about cyber offense to a model's retrieval database may furnish a model with novel cyber offensive capabilities. Changes to retrieval databases should necessitate significant auditing.

There are multiple kinds of retrieval methods; not all have the same auditing requirements. Most kinds use an embedding model to vectorize a corpus of data, such as Wikipedia or the entire training dataset, and then look up useful vectors in this database using vector similarity metrics. From that point, different methods use different approaches. For instance, Lewis *et al.* [19] insert the looked-up data into the prompt so that the language model can use the extra information in the prompt to better predict the next token. Another method, RETRO [18] instead incorporates the database vectors using cross-attention, which the network learns through finetuning. The retrieval database in Lewis *et al.* [19] may be best thought of as system parameters, whereas the database in Borgeaud *et al.* [18] may be best thought of as model parameters. In both cases, changes to retrieval databases should warrant AI system evaluations, since they may change systems' absolute capabilities. A special case of retrieval is 'memorizing transformers' [47], where the embedding model is the language model itself at previous timesteps (i.e. the retrieval database consists of the hidden activations of the AI system for earlier parts of the dataset during training). This is functionally similar to giving the system a much larger context window. Retrieval databases of memorizing transformers are thus constructed 'online', whereas retrieval databases as in RETRO [18] are constructed 'offline'. Similar to training data, it will be more challenging to audit retrieval databases that are constructed online since auditing must be performed constantly; by contrast, we can audit an offline-constructed retrieval database once.

#### Summary of recommendations regarding learning

- For each new experiment, require audits in proportion to expected capability increases from pretraining or fine-tuning.
- Filter prompts and retrieval databases to avoid AI systems learning potentially dangerous capabilities using in-context learning.
- When system parameters, such as retrieval databases, are changed, the AI system should undergo renewed auditing.

#### 2.3.5. Effective compute and training data content

Changes in a system's mechanistic structure that increase absolute capabilities necessitate further auditing. We identify the main inputs to a system's absolute capabilities as 1) Effective compute and 2) Training data content.

## Effective compute

Vast leaps in the absolute capabilities of AI systems in recent years have been, to a large extent, driven by similarly vast increases in the size of systems, the amount of data, and the amount of computation used to train them [60–62]. There has also been significant *algorithmic progress*, which reduces the amount of compute needed to obtain the same performance [63,64]. This means that even if we held compute constant, algorithmic progress would mean that *effective compute* would continue to increase, where effective compute is the product of the amount of compute used and the efficiency of how that compute is used. Systems trained with additional effective compute leads to increased absolute capabilities, introducing further risks, which should necessitate additional auditing.

Like absolute capabilities, effective compute is difficult to measure. This makes effective compute impractical for use as a policy lever. Instead, regulators, and hence auditors, will likely focus on correlates: Amount of compute and algorithmic progress.

- **Amount of compute** Given that AI systems that undergo additional training may have increased absolute capabilities, they should require additional auditing. Unfortunately, even though the absolute amount of computation used in training-experiments correlates with a system's absolute capabilities, it is not possible to use it to predict accurately when particular capabilities will emerge [48–50]. We want to avoid risky scenarios where systems trained with increasing amounts of compute suddenly become able to exhibit certain dangerous behavior. We may be able to do this by evaluating systems trained with a lower level of compute to understand their capabilities profiles before moving to systems trained with slightly more. Risk assessments should ensure that slightly smaller systems have undergone adequate auditing before permission is given to develop larger systems. If compute thresholds exist, auditors may ensure that training-experiment designs and the experiments themselves do not exceed permitted thresholds.
- **Algorithmic progress** AI is a dual-use technology; it has great potential for both good and harm. Due to its dual-use nature, widespread access to large amounts of effective compute could result in proliferation of dangerous misuse risks (such as systems capable of automating the discovery of zero-day exploits on behalf of bad actors) or accident risks (such as systems capable of autonomous exfiltration and replication). Algorithmic progress is one of the key inputs to effective compute. Thus, when tasked with governing effective compute, policymakers are faced with a challenge: Algorithmic improvements are often published openly. But future publications may have unpredictable effects on the amount of effective compute available to all actors. Therefore, the standard publication norm or openness may therefore unintentionally provide dangerous actors with increased effective compute. Policymakers may therefore consider it necessary from a security perspective to implement publication controls, such as requiring pre-publication risk assessments, in order to prevent undesirable actors gaining access to potentially dangerous amounts of effective compute. Such publication controls for dual-use technologies have precedent; they are the norm in nuclear technology, for instance [65]. It may not be possible to rely on lab self-regulation, since labs are incentivised to publish openly in order to garner prestige or to attract the best researchers. Regulation would therefore likely be required to ensure that pre-publication risk assessments take national security into account sufficiently. Assessing risks of publication may require significant independent technical expertise, a role which regulators may either have in-house or could be drawn from auditing organizations.

In addition to facilitating pre-publication risk assessments, auditing organizations may also serve as monitors of algorithmic progress within labs, since doing so would require access to frontier AI systems in order to evaluate them, adequate technical capacity, and independence from other incentives. Between AI labs, regulators, and auditors, it may therefore be the case that auditors are the best positioned actors to perform this function. However, currently no consensus metric of algorithmic progress exists and further research is required to identify metrics that are

predictive of capability levels. Building a 'science of evals' and designing metrics that are more predictive of capabilities than compute should be priority research areas.

### Training data content

AI system capabilities emerge as a result of learning from the training data. This means that we can exert some control over system behavior by controlling the data it is trained on. For instance, we can control AI systems' absolute capabilities [66–68] or ease of alignment [69] by controlling for training data quality or training only on certain subsets of data [70]. Similarly, we may be able to avoid certain dangerous capabilities from emerging at all by carefully curating an AI system's training data. For instance, if we want to ensure that a system cannot easily tell users how to synthesize illicit substances (even if 'jailbroken' [71]), then we should remove data related to the synthesis of those substances from the training data; similarly, if we wish to limit a system's ability to exfiltrate itself, we should consider removing cybersecurity-related data from its training corpus. However, filtering data is unlikely to yield strong guarantees about AI system capabilities because AI systems can generalize. For instance, a system might be able to give instructions on how to synthesize dangerous chemicals using its knowledge of how to synthesize similar, safe compounds. In some instances, data may therefore be inherently dual-use. Nevertheless, figuring something out for oneself is harder than being taught; we can therefore potentially make some behaviors harder to obtain than they otherwise would be by filtering data. External and internal auditors should evaluate training data pertaining to potentially dangerous behaviors and filter as appropriate.

When auditing of data should occur will depend on certain aspects of training. In offline training, the training data are pre-collected prior to training. By contrast, online training involves data that are collected during training, often as a result of interactions with an environment. Auditing the training data of an AI system trained offline can be done by filtering its corpus for sensitive data. This step may be expensive, but only need be done once, prior to beginning training. Filtering data during online training may be more expensive since it must be performed continually, but we otherwise have few safety assurances over online training data content. Iterative retraining has properties of both offline and online training; data content should be audited prior to each bout of training.

### Summary of recommendations regarding effective compute and training data content

- For each new experiment, require audits in proportion to expected capability increases from additional effective compute or different training data content.
- Conduct risk assessments of slightly smaller AI systems before approval is given to develop a larger system.
- Place strict controls on training-experiments that use above a certain level of effective compute.
- Implement national security-focused publication controls on research related to AI capabilities.
- Auditors should be able to evaluate training data.
- Filter training data for potentially dangerous or sensitive content.

### 2.3.6. Security

#### Preventing AI system theft and espionage

Proper security measures are crucial for controlling the affordances available to systems. Adequate security helps avoid several risky scenarios such as proliferation of the system by malicious actors (including advanced threat actors such as hostile governments) or the system self-proliferating. Any organization interacting with AI systems and the computer infrastructure that they run on should implement strict cyber- and physical security to prevent unauthorized access and AI system exfiltration. Given the security requirements of a potentially strategically valuable dual-use technology, military-grade security, espionage protection, and red-teaming may be required. Individuals with high levels of access to AI systems should require background checks. Security auditors should assess

the adequacy of security measures and compliance with information security standards through reviews and red-teaming. The security of labs should be coordinated, such that information concerning security and safety is shared between labs, either directly or through a third party such as a common government-led project or a network of security auditors. Structured access APIs [37,38] should be developed that give appropriate degrees of access to AI system developers, researchers, and auditors. Further research should be carried out on methods for verifying that code run on computing hardware is compliant with safety regulations [72].

### Preventing misuse of AI systems

Beyond cyber- and physical security practices, AI systems introduce unique security challenges. External challenges from users include prompt injection attacks, jailbreaking, and malicious use. Mitigating these risks are active research areas. AI systems' inputs should be monitored and subject to guardrails to help tackle these risks, and auditors should be involved in assessing the extent of these risks in AI systems destined for deployment. Filtering inputs may succeed in avoiding most, but not all, jailbreaking and misuse risks. For example, if it ever became prohibited for certain systems to be placed in scaffolding that let them behave like autonomous agents (e.g. scaffolding such as LangChain or AutoGPT), then it may be possible to identify and filter the prompts used in this scaffolding, thus preventing that particular use. Auditors, in the form of internal and external red-teams, should be charged with proactively identifying vulnerabilities. Additionally, labs and governments should implement bug bounty programs, which incentivize people to find and report vulnerabilities and dangerous capabilities. Deployment audits or security audits should ensure that Know-Your-Customer (KYC) requirements and monitoring of interactions with AI systems are implemented; these may help mitigate the risks from exposure of AI systems to threat actors through prevention and monitoring.

### Protection from dangerous autonomous AI systems

AI systems with sufficient absolute capabilities may themselves pose security risks if acting autonomously. The affordances available to such systems should be subject to strict security assessments from external auditors. AI systems should not have unfettered freedom to acquire new affordances, such as the ability to access arbitrary code libraries or software tools. There should be clear mechanisms to recall AI systems or restrict the affordances available to them. Any interfaces that the system interacts with should be monitored for potentially dangerous use. Ongoing security assessments should be performed throughout AI system development, especially when new affordances are proposed. Scenario planning can help identify risks associated with increased affordances.

### Incident response plans

Institutions should have clear security protocols in place to act quickly and effectively in response to safety and security incidents. For example, at labs there should be fail-safes, rapid response plans, and incident reporting protocols. Internal and external auditors should perform risk assessments, which should be published openly (to the extent that they do not reveal proprietary or security-compromising information). Whistleblower protections may help ensure that vulnerabilities are reported to the correct authorities; security or governance audits should ensure adequate protection. As for other dual-use technologies, emergency response plans should include government departments involved in national security to ensure that, if necessary, use of force during security and safety incidents involving highly capable frontier AI systems is legitimate.

### Summary of recommendations regarding security

- Organizations with access to advanced AI systems should have military-grade information security, espionage protection, and red-teaming protocols.

- Implement strict cyber- and physical security practices to prevent unauthorized access and AI system exfiltration.
- Structured access APIs and other technical controls to enable secure development and sharing with researchers, auditors, or the broader public.
- Individuals with high levels of access to AI systems should require background checks.
- Information sharing of security and safety incidents between labs.
- The level of access to AI systems given to developers, researchers, and auditors should be appropriate and not excessive.
- Monitoring of compute usage to ensure compliance with regulations regarding the amount of compute used and how it is used.
- Prompt filtering and other input controls to prevent malicious and dangerous use, such as prohibited scaffolding methods.
- Fail-safes and rapid response plans in case the AI system does gain access to more affordances (e.g. by auto-exfiltration).
- Mandatory reporting of safety and security incidents.

### 2.3.7. Deployment design

Deployment of AI systems is not a binary threshold. There are many kinds of deployment, each with different consequences for risks. Deployment decisions determine who has access?; when do they get access?; and what do they have access to?

Deployment audits should assess risks from different modes of deployment for each AI system to be deployed and ensure that any regulation regarding deployment is upheld. Deployment audits aim to ensure that AI systems are not *intentionally* given excessive available affordances; by contrast, security audits aim to reduce the risk that they are given excessive available affordances *unintentionally*. Deployment audits should determine whether internal researchers vs. the general public have access to an AI system. It will also determine factors such as whether they have access to finetuning; freedom to modify the system prompt; or whether they have freedom to swap out a retrieval database for another, which could affect the absolute capabilities of the system. Another capability-relevant deployment decision is whether the system will be trained while deployed (i.e. online training), necessitating elements of both deployment audits and training design audits.

#### Summary of recommendations regarding deployment design

- Deployment plans should be subject to auditing.

### 2.3.8. Training-experiment design

A training-experiment is the procedure by which an AI system is developed. Design decisions for the training-experiment include data selection and filtering, training process; model architecture and hyperparameters; choice of deep learning framework; hardware choices; the amount of compute to use; the algorithms used; evaluation procedures; safety procedures; the affordances made available to the system during training; the properties of different phases of pretraining and fine-tuning; whether to train online or offline; etc. Given that systems trained with RL algorithms may result in dangerous policies [29], experiments that involve RL could be subject to additional scrutiny through training design audits. The design of training experiments thus determines many of the downstream factors that are relevant to risk. Training-experiment designs should therefore themselves be subject to pretraining risk assessment, which may employ approaches such as incentive analysis [73]. A comprehensive AI governance regime would require pre-registration, pre-approval, and monitoring of experiments involving highly absolutely capable systems trained with large amounts of effective compute and AI systems with large affordances available to them. Regulators, auditors, and AI development labs should make their own risk assessments for such experiments, with a publicly accountable institution having the final authority on whether they go ahead.



In addition to risk assessments for individual training experiments, the overall strategy into which these experiments fit (i.e. the alignment strategy) should be evaluated in ways that are accountable to the public. In order for these assessments to occur, AI system developers could be required to publish detailed alignment strategies.

#### Summary of recommendations regarding training-experiment design

- Require pre-registration and pre-approval of training-experiments involving highly absolutely capable AI systems trained with large amounts of effective compute and AI systems with large affordances available to them.
- Training-experiment designs should be subject to prior-to-training risk assessment.
- Require developers of frontier AI systems to publish detailed alignment strategies or to make their plans available to auditors for scrutiny.
- Require regulator approval of experiments with highly capable AI systems, large sets of available affordances, or large effective compute budgets, based on risk assessments from internal and external auditors.
- Potentially require smaller scale experiments before further scaling compute. This helps assess effective compute and predict capabilities.

#### 2.3.9. Governance and institutions

Certain governance structures help determine the training-experiment design decisions, deployment decisions, or security decisions that are likely to be made. Auditors may be able to perform governance audits of the institutions developing general-purpose AI to ensure that incentives are aligned with safety. The AI governance and institutional landscape determines the constraints and incentives under which training experiment design decisions and security-relevant decisions are made. It is therefore important that this landscape be designed with safety in mind, rather than let it develop unconstrained under other incentives, such as profit maximization.

Mökander *et al.* [12] discuss governance audits in the context of large language model development and deployment, though their work can readily be applied to broader frontier AI systems. They identify three roles for governance audits:

1. Reviewing the adequacy of organizational governance structures
2. Creating an audit trail of the frontier AI systems development process
3. Mapping roles and responsibilities within organizations that design frontier AI systems.

As the regulatory landscape is currently being designed, research should aim to anticipate and avoid potential barriers to these tasks, such as ways to ensure institutional transparency of the labs developing frontier AI systems.

#### Summary of recommendations regarding governance and institutions

- Labs and other relevant actors should be rendered transparent enough to regulators for effective governance.
- Regulators commission and act on governance audits when structuring the governance and institutional landscape.

### 3. Key areas of research

To build an auditing regime that can ensure adequate safety of AI development and deployment, further research is needed in the following key areas.

#### 3.1. Technical AI safety research

- **Interpretability and training dynamics:** Improve methods for explaining the internal mechanisms of AI system behavior. Develop better understanding of how capabilities emerge



during training through phase transitions. Use this to create predictive models that forecast emergence of new capabilities.

- **Behavioral evaluations of dangerous capabilities:** In advance of adequate interpretability methods, we should develop better behavioral methods for assessing risks from AI systems. We must develop evaluations that can serve as clear decision boundaries for particular kinds of regulatory or other risk-mitigating actions. We should improve our predictive models of how capabilities emerge in AI systems.
- **Alignment theory:** Further develop the theoretical understanding of how to create AI systems whose goals and incentives are robustly aligned with human values. This might eventually provide technically grounded metrics against which AI systems can be audited.

### 3.2. Technical infrastructure research

- **Structured access frameworks:** Design interfaces and structured access protocols that enable external auditors and researchers to safely analyze AI system mechanistic structure in ways that avoid misuse and proliferation.
- **Auditing of training compute usage:** Create methods to monitor and audit how computational resources are used during AI system training to ensure adherence to safety requirements and regulations. The methods would preferably be privacy-preserving, secure, and permit verification of code, data pipelines, compute usage, and other aspects of the training process.
- **Technically grounded definitions of effective compute and algorithmic progress:** A prerequisite for governance and auditing of effective compute is a technically grounded definition. Develop rigorous technical definitions and metrics for measuring the effective compute used during AI system training and for measuring algorithmic progress.

### 3.3. Institutional governance and policy research:

- **Accountability of auditors:** Research protocols and institutional designs that ensure accountability of auditors to the public, while also controlling potentially hazardous information flows.
- **Institutional design for transformational technology:** The current political economy of general-purpose AI development, which is currently driven by private interests, may permit less public accountability than is ideal for such a transformational technology. As the technology advances, we should consider alternative frameworks that bring AI development into a regime with a greater focus on security and public accountability, such as nationalization or internationalization.
- **Adaptive policy making and enforcement:** Build regulatory and policy-making capability to enable rapid adaptation of regulatory infrastructure as AI progresses at pace. Ensure that auditors adapt at the same pace.
- **Legal frameworks:** Explore legal tools like liability, regulations, and treaties to align AI with public interests. These provide the basis for compliance audits.
- **Frameworks for cooperation:** Develop frameworks to facilitate cooperation between governments and between companies on AI governance, even where little mutual trust exists. This may be required for auditors to operate cross-jurisdictionally, which will be required for a global approach to AI risk reduction.

**Author Contributions:** Lee Sharkey conceived and wrote initial drafts for all sections of this article, made initial designs for the figures, and was responsible for final editing. Clíodhna Ní Ghuidhir contributed substantial editing throughout. Clíodhna Ní Ghuidhir, Dan Braun, Marius Hobbahn provided comments and important discussion on conceptual parts of the article. Jérémy Scheurer and Lucius Bushnaq provided useful technical commentary and references. Charlotte Stix, Clíodhna Ní Ghuidhir, and Mikita Balesni provided feedback on content and assisted with figures.

**Acknowledgments:** Many external commenters helped greatly to improve this article, including: Markus Anderljung, Joe O'Brien, Ben Bucknall, Tom Davidson, Lisa Soder, Leonie Koessler, Jason Hausenloy, Adam Jones, Zach Stein-Perlman, Andrew Trask, Herbie Bradley, and Lennart Heim.

## References

1. Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Židek, A.; Potapenko, A.; Bridgland, A.; Meyer, C.; Kohl, S.A.A.; Ballard, A.J.; Cowie, A.; Romera-Paredes, B.; Nikolov, S.; Jain, R.; Adler, J.; Back, T.; Petersen, S.; Reiman, D.; Clancy, E.; Zielinski, M.; Steinegger, M.; Pacholska, M.; Berghammer, T.; Bodenstein, S.; Silver, D.; Vinyals, O.; Senior, A.W.; Kavukcuoglu, K.; Kohli, P.; Hassabis, D. Highly accurate protein structure prediction with AlphaFold. *Nature* **2021**, *596*, 583–589. doi:10.1038/s41586-021-03819-2.
2. Singhal, K.; Azizi, S.; Tu, T.; Mahdavi, S.S.; Wei, J.; Chung, H.W.; Scales, N.; Tanwani, A.; Cole-Lewis, H.; Pfohl, S.; et al.. Large language models encode clinical knowledge, 2023.
3. Brundage, M.; Avin, S.; Clark, J.; Toner, H.; Eckersley, P.; Garfinkel, B.; Dafoe, A.; Scharre, P.; Zeitzoff, T.; Filar, B.; Anderson, H.S.; Roff, H.; Allen, G.C.; Steinhardt, J.; Flynn, C.; hÉigeartaigh, S.Ó.; Beard, S.; Belfield, H.; Farquhar, S.; Lyle, C.; Crootof, R.; Evans, O.; Page, M.; Bryson, J.; Yampolskiy, R.; Amodei, D. The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation. *CoRR* **2018**, *abs/1802.07228*, [1802.07228].
4. Anderljung, M.; Hazell, J. Protecting Society from AI Misuse: When are Restrictions on Capabilities Warranted?, 2023, [arXiv:cs.AI/2303.09377].
5. Solaiman, I.; Talat, Z.; Agnew, W.; Ahmad, L.; Baker, D.; Blodgett, S.L.; au2, H.D.I.; Dodge, J.; Evans, E.; Hooker, S.; Jernite, Y.; Luccioni, A.S.; Lusoli, A.; Mitchell, M.; Newman, J.; Png, M.T.; Strait, A.; Vassilev, A. Evaluating the Social Impact of Generative AI Systems in Systems and Society, 2023, [arXiv:cs.CY/2306.05949].
6. Hendrycks, D.; Mazeika, M.; Woodside, T. An Overview of Catastrophic AI Risks, 2023, [arXiv:cs.CY/2306.12001].
7. Ngo, R.; Chan, L.; Mindermann, S. The alignment problem from a deep learning perspective, 2023, [arXiv:cs.AI/2209.00626].
8. Schuett, J.; Dreksler, N.; Anderljung, M.; McCaffary, D.; Heim, L.; Bluemke, E.; Garfinkel, B. Towards best practices in AGI safety and governance: A survey of expert opinion, 2023, [arXiv:cs.CY/2305.07153].
9. Shevlane, T.; Farquhar, S.; Garfinkel, B.; Phuong, M.; Whittlestone, J.; Leung, J.; Kokotajlo, D.; Marchal, N.; Anderljung, M.; Kolt, N.; Ho, L.; Siddarth, D.; Avin, S.; Hawkins, W.; Kim, B.; Gabriel, I.; Bolina, V.; Clark, J.; Bengio, Y.; Christiano, P.; Dafoe, A. Model evaluation for extreme risks, 2023, [arXiv:cs.AI/2305.15324].
10. Anderljung, M.; Barnhart, J.; Korinek, A.; Leung, J.; O’Keefe, C.; Whittlestone, J.; Avin, S.; Brundage, M.; Bullock, J.; Cass-Beggs, D.; Chang, B.; Collins, T.; Fist, T.; Hadfield, G.; Hayes, A.; Ho, L.; Hooker, S.; Horvitz, E.; Kolt, N.; Schuett, J.; Shavit, Y.; Siddarth, D.; Trager, R.; Wolf, K. Frontier AI Regulation: Managing Emerging Risks to Public Safety, 2023, [arXiv:cs.CY/2307.03718].
11. Whittlestone, J.; Clark, J. Why and How Governments Should Monitor AI Development, 2021, [arXiv:cs.CY/2108.12427].
12. Mökander, J.; Schuett, J.; Kirk, H.R.; Floridi, L. Auditing large language models: a three-layered approach. *AI and Ethics* **2023**. doi:10.1007/s43681-023-00289-2.
13. Raji, I.D.; Xu, P.; Honigsberg, C.; Ho, D.E. Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance, 2022, [arXiv:cs.CY/2206.04737].
14. Anderljung, M.; Smith, E.T.; O’Brien, J.; Soder, L.; Bucknall, B.; Bluemke, E.; Schuett, J.; Trager, R.; Strahm, L.; Chowdhury, R. Towards Publicly Accountable Frontier LLMs: Building an External Scrutiny Ecosystem under the ASPIRE Framework., forthcoming.
15. Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.L.; Cao, Y.; Narasimhan, K. Tree of Thoughts: Deliberate Problem Solving with Large Language Models, 2023, [arXiv:cs.CL/2305.10601].
16. Long, J. Large Language Model Guided Tree-of-Thought, 2023, [arXiv:cs.AI/2305.08291].
17. Gravitas, S. Significant-Gravitas/AutoGPT: An experimental open-source attempt to make GPT-4 fully autonomous. <https://github.com/Significant-Gravitas/AutoGPT>, 2023. (Accessed on 10/13/2023).
18. Borgeaud, S.; Mensch, A.; Hoffmann, J.; Cai, T.; Rutherford, E.; Millican, K.; van den Driessche, G.; Lespiau, J.B.; Damoc, B.; Clark, A.; de Las Casas, D.; Guy, A.; Menick, J.; Ring, R.; Hennigan, T.; Huang, S.; Maggiore, L.; Jones, C.; Cassirer, A.; Brock, A.; Paganini, M.; Irving, G.; Vinyals, O.; Osindero, S.; Simonyan, K.; Rae, J.W.; Elsen, E.; Sifre, L. Improving language models by retrieving from trillions of tokens, 2022, [arXiv:cs.CL/2112.04426].

19. Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; tau Yih, W.; Rocktäschel, T.; Riedel, S.; Kiela, D. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *Advances in Neural Information Processing Systems*; Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; Lin, H., Eds. Curran Associates, Inc., 2020, Vol. 33, pp. 9459–9474.
20. Heim, L. Information security considerations for AI and the long term future. <https://blog.heim.xyz/information-security-considerations-for-ai/>, 2022. (Accessed on 10/12/2023).
21. Boiko, D.A.; MacKnight, R.; Gomes, G. Emergent autonomous scientific research capabilities of large language models, 2023, [arXiv:physics.chem-ph/2304.05332].
22. Brohan, A.; Brown, N.; Carbajal, J.; Chebotar, Y.; Chen, X.; Choromanski, K.; Ding, T.; Driess, D.; Dubey, A.; Finn, C.; Florence, P.; Fu, C.; Arenas, M.G.; Gopalakrishnan, K.; Han, K.; Hausman, K.; Herzog, A.; Hsu, J.; Ichter, B.; Irpan, A.; Joshi, N.; Julian, R.; Kalashnikov, D.; Kuang, Y.; Leal, I.; Lee, L.; Lee, T.W.E.; Levine, S.; Lu, Y.; Michalewski, H.; Mordatch, I.; Pertsch, K.; Rao, K.; Reymann, K.; Ryoo, M.; Salazar, G.; Sanketi, P.; Sermanet, P.; Singh, J.; Singh, A.; Soricute, R.; Tran, H.; Vanhoucke, V.; Vuong, Q.; Wahid, A.; Welker, S.; Wohlhart, P.; Wu, J.; Xia, F.; Xiao, T.; Xu, P.; Xu, S.; Yu, T.; Zitkovich, B. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control, 2022. (Accessed on 10/12/2023).
23. OpenAI. ChatGPT plugins. <https://openai.com/blog/chatgpt-plugins#code-interpreter>, 2023. (Accessed on 10/12/2023).
24. Reed, S.; Zolna, K.; Parisotto, E.; Colmenarejo, S.G.; Novikov, A.; Barth-maroon, G.; Giménez, M.; Sulsky, Y.; Kay, J.; Springenberg, J.T.; Eccles, T.; Bruce, J.; Razavi, A.; Edwards, A.; Heess, N.; Chen, Y.; Hadsell, R.; Vinyals, O.; Bordbar, M.; de Freitas, N. A Generalist Agent. *Transactions on Machine Learning Research* **2022**. Featured Certification, Outstanding Certification.
25. Wang, G.; Xie, Y.; Jiang, Y.; Mandlekar, A.; Xiao, C.; Zhu, Y.; Fan, L.; Anandkumar, A. Voyager: An Open-Ended Embodied Agent with Large Language Models, 2023, [arXiv:cs.AI/2305.16291].
26. Bubeck, S.; Chandrasekaran, V.; Eldan, R.; Gehrke, J.; Horvitz, E.; Kamar, E.; Lee, P.; Lee, Y.T.; Li, Y.; Lundberg, S.; Nori, H.; Palangi, H.; Ribeiro, M.T.; Zhang, Y. Sparks of Artificial General Intelligence: Early experiments with GPT-4, 2023, [arXiv:cs.CL/2303.12712].
27. Schick, T.; Dwivedi-Yu, J.; Dessi, R.; Raileanu, R.; Lomeli, M.; Zettlemoyer, L.; Cancedda, N.; Scialom, T. Toolformer: Language Models Can Teach Themselves to Use Tools, 2023, [arXiv:cs.CL/2302.04761].
28. Rando, J.; Paleka, D.; Lindner, D.; Heim, L.; Tramèr, F. Red-Teaming the Stable Diffusion Safety Filter, 2022, [arXiv:cs.AI/2210.04610].
29. Turner, A.M.; Smith, L.; Shah, R.; Critch, A.; Tadepalli, P. Optimal Policies Tend to Seek Power, 2019, [arXiv:cs.AI/1912.01683].
30. Reynolds, L.; McDonnell, K. Prompt Programming for Large Language Models: Beyond the Few-Shot Paradigm, 2021, [arXiv:cs.CL/2102.07350].
31. Shin, T.; Razeghi, Y.; au2, R.L.L.I.; Wallace, E.; Singh, S. AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts, 2020, [arXiv:cs.CL/2010.15980].
32. Wen, Y.; Jain, N.; Kirchenbauer, J.; Goldblum, M.; Geiping, J.; Goldstein, T. Hard Prompts Made Easy: Gradient-Based Discrete Optimization for Prompt Tuning and Discovery, 2023, [arXiv:cs.LG/2302.03668].
33. Jones, E.; Dragan, A.; Raghunathan, A.; Steinhart, J. Automatically Auditing Large Language Models via Discrete Optimization, 2023, [arXiv:cs.LG/2303.04381].
34. Li, X.L.; Liang, P. Prefix-Tuning: Optimizing Continuous Prompts for Generation, 2021, [arXiv:cs.CL/2101.00190].
35. Christiano, P.; Leike, J.; Brown, T.B.; Martic, M.; Legg, S.; Amodei, D. Deep reinforcement learning from human preferences, 2017, [arXiv:stat.ML/1706.03741].
36. Levinstein, B.A.; Herrmann, D.A. Still No Lie Detector for Language Models: Probing Empirical and Conceptual Roadblocks, 2023, [arXiv:cs.CL/2307.00175].
37. Trask, A.; Sukumar, A.; Kalliokoski, A.; Farkas, B.; Ezenwaka, C.; Popa, C.; Mitchell, C.; Hrebenach, D.; Muraru, G.C.; Junior, I.; Bejan, I.; Mishra, I.; Ngong, I.; Bandy, J.; Stahl, J.; Cardonnet, J.; Trask, K.; Trask, K.; Nguyen, K.; Dang, K.; van der Veen, K.; Eng, K.; Strahm, L.; Ayre, L.; Jay, M.; Lytvyn, O.; Kyemenu-Sarsah, O.; Chung, P.; Smith, P.; S, R.; Falcon, R.; Gupta, S.; Gabriel, S.; Milea, T.; Thoraldson, T.; Porto, T.; Cebere, T.; Gorana, Y.; Reza, Z. How to audit an AI model owned by someone else (part 1). <https://blog.openmined.org/ai-audit-part-1/>, 2023. (Accessed on 10/13/2023).

38. Shevlane, T. Structured access: an emerging paradigm for safe AI deployment, 2022, [\[arXiv:cs.AI/2201.05159\]](#).
39. von Oswald, J.; Niklasson, E.; Randazzo, E.; Sacramento, J.; Mordvintsev, A.; Zhmoginov, A.; Vladymyrov, M. Transformers learn in-context by gradient descent, 2023, [\[arXiv:cs.LG/2212.07677\]](#).
40. Xie, S.M.; Raghunathan, A.; Liang, P.; Ma, T. An Explanation of In-context Learning as Implicit Bayesian Inference, 2022, [\[arXiv:cs.CL/2111.02080\]](#).
41. Wei, J.; Bosma, M.; Zhao, V.Y.; Guu, K.; Yu, A.W.; Lester, B.; Du, N.; Dai, A.M.; Le, Q.V. Finetuned Language Models Are Zero-Shot Learners, 2022, [\[arXiv:cs.CL/2109.01652\]](#).
42. Bai, Y.; Kadavath, S.; Kundu, S.; Askell, A.; Kernion, J.; Jones, A.; Chen, A.; Goldie, A.; Mirhoseini, A.; McKinnon, C.; Chen, C.; Olsson, C.; Olah, C.; Hernandez, D.; Drain, D.; Ganguli, D.; Li, D.; Tran-Johnson, E.; Perez, E.; Kerr, J.; Mueller, J.; Ladish, J.; Landau, J.; Ndousse, K.; Lukosuite, K.; Lovitt, L.; Sellitto, M.; Elhage, N.; Schiefer, N.; Mercado, N.; DasSarma, N.; Lasenby, R.; Larson, R.; Ringer, S.; Johnston, S.; Kravec, S.; Showk, S.E.; Fort, S.; Lanham, T.; Telleen-Lawton, T.; Conerly, T.; Henighan, T.; Hume, T.; Bowman, S.R.; Hatfield-Dodds, Z.; Mann, B.; Amodei, D.; Joseph, N.; McCandlish, S.; Brown, T.; Kaplan, J. Constitutional AI: Harmlessness from AI Feedback, 2022, [\[arXiv:cs.CL/2212.08073\]](#).
43. OpenAI. GPT-4 Technical Report, 2023, [\[arXiv:cs.CL/2303.08774\]](#).
44. Tay, Y.; Wei, J.; Chung, H.W.; Tran, V.Q.; So, D.R.; Shakeri, S.; Garcia, X.; Zheng, H.S.; Rao, J.; Chowdhery, A.; Zhou, D.; Metzler, D.; Petrov, S.; Houlsby, N.; Le, Q.V.; Dehghani, M. Transcending Scaling Laws with 0.1 [\[arXiv:cs.CL/2210.11399\]](#).
45. Hu, E.J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W. LoRA: Low-Rank Adaptation of Large Language Models, 2021, [\[arXiv:cs.CL/2106.09685\]](#).
46. Alayrac, J.B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; Ring, R.; Rutherford, E.; Cabi, S.; Han, T.; Gong, Z.; Samangooei, S.; Monteiro, M.; Menick, J.; Borgeaud, S.; Brock, A.; Nematzadeh, A.; Sharifzadeh, S.; Binkowski, M.; Barreira, R.; Vinyals, O.; Zisserman, A.; Simonyan, K. Flamingo: a Visual Language Model for Few-Shot Learning, 2022, [\[arXiv:cs.CV/2204.14198\]](#).
47. Wu, Y.; Rabe, M.N.; Hutchins, D.; Szegedy, C. Memorizing Transformers, 2022, [\[arXiv:cs.LG/2203.08913\]](#).
48. Srivastava, A.; Rastogi, A.; Rao, A.; Shoeb, A.A.M.; Abid, A.; Fisch, A.; Brown, A.R.; Santoro, A.; Gupta, A.; Garriga-Alonso, A.; Kluska, A.; Lewkowycz, A.; Agarwal, A.; Power, A.; Ray, A.; Warstadt, A.; Kocurek, A.W.; Safaya, A.; Tazarv, A.; Xiang, A.; Parrish, A.; Nie, A.; Hussain, A.; Askell, A.; Dsouza, A.; Slone, A.; Rahane, A.; Iyer, A.S.; Andreassen, A.; Madotto, A.; Santilli, A.; Stuhlmüller, A.; Dai, A.; La, A.; Lampinen, A.; Zou, A.; Jiang, A.; Chen, A.; Vuong, A.; Gupta, A.; Gottardi, A.; Norelli, A.; Venkatesh, A.; Gholamidavoodi, A.; Tabassum, A.; Menezes, A.; Kirubakaran, A.; Mullokandov, A.; Sabharwal, A.; Herrick, A.; Efrat, A.; Erdem, A.; Karakaş, A.; Roberts, B.R.; Loe, B.S.; Zoph, B.; Bojanowski, B.; Özyurt, B.; Hedayatnia, B.; Neyshabur, B.; Inden, B.; Stein, B.; Ekmekci, B.; Lin, B.Y.; Howald, B.; Orinon, B.; Diaio, C.; Dour, C.; Stinson, C.; Argueta, C.; Ramírez, C.F.; Singh, C.; Rathkopf, C.; Meng, C.; Baral, C.; Wu, C.; Callison-Burch, C.; Waites, C.; Voigt, C.; Manning, C.D.; Potts, C.; Ramirez, C.; Rivera, C.E.; Siro, C.; Raffel, C.; Ashcraft, C.; Garbacea, C.; Sileo, D.; Garrette, D.; Hendrycks, D.; Kilman, D.; Roth, D.; Freeman, D.; Khashabi, D.; Levy, D.; González, D.M.; Perszyk, D.; Hernandez, D.; Chen, D.; Ippolito, D.; Gilboa, D.; Dohan, D.; Drakard, D.; Jurgens, D.; Datta, D.; Ganguli, D.; Emelin, D.; Kleyko, D.; Yuret, D.; Chen, D.; Tam, D.; Hupkes, D.; Misra, D.; Buzan, D.; Mollo, D.C.; Yang, D.; Lee, D.H.; Schrader, D.; Shutova, E.; Cubuk, E.D.; Segal, E.; Hagerman, E.; Barnes, E.; Donoway, E.; Pavlick, E.; Rodola, E.; Lam, E.; Chu, E.; Tang, E.; Erdem, E.; Chang, E.; Chi, E.A.; Dyer, E.; Jerzak, E.; Kim, E.; Manyasi, E.E.; Zheltonozhskii, E.; Xia, F.; Siar, F.; Martínez-Plumed, F.; Happé, F.; Chollet, F.; Rong, F.; Mishra, G.; Winata, G.I.; de Melo, G.; Kruszewski, G.; Parascandolo, G.; Mariani, G.; Wang, G.; Jaimovitch-López, G.; Betz, G.; Gur-Ari, G.; Galijasevic, H.; Kim, H.; Rashkin, H.; Hajishirzi, H.; Mehta, H.; Bogar, H.; Shevlin, H.; Schütze, H.; Yakura, H.; Zhang, H.; Wong, H.M.; Ng, I.; Noble, I.; Jumelet, J.; Geissinger, J.; Kernion, J.; Hilton, J.; Lee, J.; Fisac, J.F.; Simon, J.B.; Koppel, J.; Zheng, J.; Zou, J.; Kocoń, J.; Thompson, J.; Wingfield, J.; Kaplan, J.; Radom, J.; Sohl-Dickstein, J.; Phang, J.; Wei, J.; Yosinski, J.; Novikova, J.; Bosscher, J.; Marsh, J.; Kim, J.; Taal, J.; Engel, J.; Alabi, J.; Xu, J.; Song, J.; Tang, J.; Waweru, J.; Burden, J.; Miller, J.; Balis, J.U.; Batchelder, J.; Berant, J.; Frohberg, J.; Rozen, J.; Hernandez-Orallo, J.; Boudeman, J.; Guerr, J.; Jones, J.; Tenenbaum, J.B.; Rule, J.S.; Chua, J.; Kanclerz, K.; Livescu, K.; Krauth, K.; Gopalakrishnan, K.; Ignatyeva, K.; Markert, K.; Dhole, K.D.; Gimpel, K.; Omondi, K.; Mathewson, K.; Chiafullo, K.; Shkaruta, K.; Shridhar, K.; McDonnell, K.; Richardson, K.; Reynolds, L.; Gao, L.; Zhang, L.; Dugan, L.; Qin, L.; Contreras-Ochando, L.; Morency, L.P.;



- Moschella, L.; Lam, L.; Noble, L.; Schmidt, L.; He, L.; Colón, L.O.; Metz, L.; Şenel, L.K.; Bosma, M.; Sap, M.; ter Hoeve, M.; Farooqi, M.; Faruqui, M.; Mazeika, M.; Baturan, M.; Marelli, M.; Maru, M.; Quintana, M.J.R.; Tolkiehn, M.; Giulianelli, M.; Lewis, M.; Potthast, M.; Leavitt, M.L.; Hagen, M.; Schubert, M.; Baitemirova, M.O.; Arnaud, M.; McElrath, M.; Yee, M.A.; Cohen, M.; Gu, M.; Ivanitskiy, M.; Starritt, M.; Strube, M.; Swędrowski, M.; Bevilacqua, M.; Yasunaga, M.; Kale, M.; Cain, M.; Xu, M.; Suzgun, M.; Walker, M.; Tiwari, M.; Bansal, M.; Aminnaseri, M.; Geva, M.; Gheini, M.; T, M.V.; Peng, N.; Chi, N.A.; Lee, N.; Krakover, N.G.A.; Cameron, N.; Roberts, N.; Doiron, N.; Martinez, N.; Nangia, N.; Deckers, N.; Muennighoff, N.; Keskar, N.S.; Iyer, N.S.; Constant, N.; Fiedel, N.; Wen, N.; Zhang, O.; Agha, O.; Elbaghdadi, O.; Levy, O.; Evans, O.; Casares, P.A.M.; Doshi, P.; Fung, P.; Liang, P.P.; Vicol, P.; Alipoormolabashi, P.; Liao, P.; Liang, P.; Chang, P.; Eckersley, P.; Htut, P.M.; Hwang, P.; Miłkowski, P.; Patil, P.; Pezeshkpour, P.; Oli, P.; Mei, Q.; Lyu, Q.; Chen, Q.; Banjade, R.; Rudolph, R.E.; Gabriel, R.; Habacker, R.; Risco, R.; Millièrè, R.; Garg, R.; Barnes, R.; Saurous, R.A.; Arakawa, R.; Raymaekers, R.; Frank, R.; Sikand, R.; Novak, R.; Sitelew, R.; LeBras, R.; Liu, R.; Jacobs, R.; Zhang, R.; Salakhutdinov, R.; Chi, R.; Lee, R.; Stovall, R.; Teehan, R.; Yang, R.; Singh, S.; Mohammad, S.M.; Anand, S.; Dillavou, S.; Shleifer, S.; Wiseman, S.; Gruetter, S.; Bowman, S.R.; Schoenholz, S.S.; Han, S.; Kwatra, S.; Rous, S.A.; Ghazarian, S.; Ghosh, S.; Casey, S.; Bischoff, S.; Gehrmann, S.; Schuster, S.; Sadeghi, S.; Hamdan, S.; Zhou, S.; Srivastava, S.; Shi, S.; Singh, S.; Asaadi, S.; Gu, S.S.; Pachchigar, S.; Toshniwal, S.; Upadhyay, S.; Shyamolima.; Debnath.; Shakeri, S.; Thormeyer, S.; Melzi, S.; Reddy, S.; Makini, S.P.; Lee, S.H.; Torene, S.; Hatwar, S.; Dehaene, S.; Divic, S.; Ermon, S.; Biderman, S.; Lin, S.; Prasad, S.; Piantadosi, S.T.; Shieber, S.M.; Mishnerghi, S.; Kiritchenko, S.; Mishra, S.; Linzen, T.; Schuster, T.; Li, T.; Yu, T.; Ali, T.; Hashimoto, T.; Wu, T.L.; Desbordes, T.; Rothschild, T.; Phan, T.; Wang, T.; Nkinyili, T.; Schick, T.; Kornev, T.; Tunduny, T.; Gerstenberg, T.; Chang, T.; Neeraj, T.; Khot, T.; Shultz, T.; Shaham, U.; Misra, V.; Demberg, V.; Nyamai, V.; Raunak, V.; Ramasesh, V.; Prabhu, V.U.; Padmakumar, V.; Srikumar, V.; Fedus, W.; Saunders, W.; Zhang, W.; Vossen, W.; Ren, X.; Tong, X.; Zhao, X.; Wu, X.; Shen, X.; Yaghoobzadeh, Y.; Lakretz, Y.; Song, Y.; Bahri, Y.; Choi, Y.; Yang, Y.; Hao, Y.; Chen, Y.; Belinkov, Y.; Hou, Y.; Hou, Y.; Bai, Y.; Seid, Z.; Zhao, Z.; Wang, Z.; Wang, Z.J.; Wang, Z.; Wu, Z. Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models, 2023, [[arXiv:cs.CL/2206.04615](https://arxiv.org/abs/2206.04615)].
49. Wei, J.; Tay, Y.; Bommasani, R.; Raffel, C.; Zoph, B.; Borgeaud, S.; Yogatama, D.; Bosma, M.; Zhou, D.; Metzler, D.; Chi, E.H.; Hashimoto, T.; Vinyals, O.; Liang, P.; Dean, J.; Fedus, W. Emergent Abilities of Large Language Models, 2022, [[arXiv:cs.CL/2206.07682](https://arxiv.org/abs/2206.07682)].
  50. Ganguli, D.; Hernandez, D.; Lovitt, L.; Askell, A.; Bai, Y.; Chen, A.; Conerly, T.; Dassarma, N.; Drain, D.; Elhage, N.; Showk, S.E.; Fort, S.; Hatfield-Dodds, Z.; Henighan, T.; Johnston, S.; Jones, A.; Joseph, N.; Kernian, J.; Kravec, S.; Mann, B.; Nanda, N.; Ndousse, K.; Olsson, C.; Amodei, D.; Brown, T.; Kaplan, J.; McCandlish, S.; Olah, C.; Amodei, D.; Clark, J. Predictability and Surprise in Large Generative Models. 2022 ACM Conference on Fairness, Accountability, and Transparency. ACM, 2022. doi:10.1145/3531146.3533229.
  51. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; Amodei, D. Language Models are Few-Shot Learners, 2020, [[arXiv:cs.CL/2005.14165](https://arxiv.org/abs/2005.14165)].
  52. Olsson, C.; Elhage, N.; Nanda, N.; Joseph, N.; DasSarma, N.; Henighan, T.; Mann, B.; Askell, A.; Bai, Y.; Chen, A.; Conerly, T.; Drain, D.; Ganguli, D.; Hatfield-Dodds, Z.; Hernandez, D.; Johnston, S.; Jones, A.; Kernion, J.; Lovitt, L.; Ndousse, K.; Amodei, D.; Brown, T.; Clark, J.; Kaplan, J.; McCandlish, S.; Olah, C. In-context Learning and Induction Heads, 2022, [[arXiv:cs.LG/2209.11895](https://arxiv.org/abs/2209.11895)].
  53. Raventós, A.; Paul, M.; Chen, F.; Ganguli, S. Pretraining task diversity and the emergence of non-Bayesian in-context learning for regression, 2023, [[arXiv:cs.LG/2306.15063](https://arxiv.org/abs/2306.15063)].
  54. Wei, J.; Wei, J.; Tay, Y.; Tran, D.; Webson, A.; Lu, Y.; Chen, X.; Liu, H.; Huang, D.; Zhou, D.; Ma, T. Larger language models do in-context learning differently, 2023, [[arXiv:cs.CL/2303.03846](https://arxiv.org/abs/2303.03846)].
  55. Langchain. Introduction | Langchain. [https://python.langchain.com/docs/get\\_started/introduction](https://python.langchain.com/docs/get_started/introduction), 2023. (Accessed on 10/13/2023).
  56. Izacard, G.; Lewis, P.; Lomeli, M.; Hosseini, L.; Petroni, F.; Schick, T.; Dwivedi-Yu, J.; Joulin, A.; Riedel, S.; Grave, E. Atlas: Few-shot Learning with Retrieval Augmented Language Models, 2022, [[arXiv:cs.CL/2208.03299](https://arxiv.org/abs/2208.03299)].

57. Zhong, Z.; Lei, T.; Chen, D. Training Language Models with Memory Augmentation, 2022, [[arXiv:cs.CL/2205.12674](https://arxiv.org/abs/2205.12674)].
58. Karpukhin, V.; Oğuz, B.; Min, S.; Lewis, P.; Wu, L.; Edunov, S.; Chen, D.; tau Yih, W. Dense Passage Retrieval for Open-Domain Question Answering, 2020, [[arXiv:cs.CL/2004.04906](https://arxiv.org/abs/2004.04906)].
59. Guu, K.; Lee, K.; Tung, Z.; Pasupat, P.; Chang, M.W. REALM: Retrieval-Augmented Language Model Pre-Training, 2020, [[arXiv:cs.CL/2002.08909](https://arxiv.org/abs/2002.08909)].
60. Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T.B.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; Amodei, D. Scaling Laws for Neural Language Models, 2020, [[arXiv:cs.LG/2001.08361](https://arxiv.org/abs/2001.08361)].
61. Hoffmann, J.; Borgeaud, S.; Mensch, A.; Buchatskaya, E.; Cai, T.; Rutherford, E.; de Las Casas, D.; Hendricks, L.A.; Welbl, J.; Clark, A.; Hennigan, T.; Noland, E.; Millican, K.; van den Driessche, G.; Damoc, B.; Guy, A.; Osindero, S.; Simonyan, K.; Elsen, E.; Rae, J.W.; Vinyals, O.; Sifre, L. Training Compute-Optimal Large Language Models, 2022, [[arXiv:cs.CL/2203.15556](https://arxiv.org/abs/2203.15556)].
62. Sevilla, J.; Heim, L.; Ho, A.; Besiroglu, T.; Hobbhahn, M.; Villalobos, P. Compute Trends Across Three Eras of Machine Learning, 2022, [[arXiv:cs.LG/2202.05924](https://arxiv.org/abs/2202.05924)].
63. Erdil, E.; Besiroglu, T. Algorithmic progress in computer vision, 2023, [[arXiv:cs.CV/2212.05153](https://arxiv.org/abs/2212.05153)].
64. Tucker, A.D.; Anderljung, M.; Dafoe, A. Social and Governance Implications of Improved Data Efficiency. Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society; Association for Computing Machinery: New York, NY, USA, 2020; AIES '20, p. 378–384. doi:10.1145/3375627.3375863.
65. Wasil, A.; Siegmann, C.; Ezell, C.; Richardson, A. WasilEzellRichardsonSiegmann+(10).pdf. <https://static1.squarespace.com/static/6276a63ecf564172c125f58e/t/641cbc1d84814a4d0f3e1788/1679604766050/WasilEzellRichardsonSiegmann+%2810%29.pdf>, 2023. (Accessed on 10/13/2023).
66. Gunasekar, S.; Zhang, Y.; Aneja, J.; Mendes, C.C.T.; Giorno, A.D.; Gopi, S.; Javaheripi, M.; Kauffmann, P.; de Rosa, G.; Saarikivi, O.; Salim, A.; Shah, S.; Behl, H.S.; Wang, X.; Bubeck, S.; Eldan, R.; Kalai, A.T.; Lee, Y.T.; Li, Y. Textbooks Are All You Need, 2023, [[arXiv:cs.CL/2306.11644](https://arxiv.org/abs/2306.11644)].
67. Eldan, R.; Li, Y. TinyStories: How Small Can Language Models Be and Still Speak Coherent English?, 2023, [[arXiv:cs.CL/2305.07759](https://arxiv.org/abs/2305.07759)].
68. Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P.J. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer, 2023, [[arXiv:cs.LG/1910.10683](https://arxiv.org/abs/1910.10683)].
69. Korbak, T.; Shi, K.; Chen, A.; Bhalerao, R.; Buckley, C.L.; Phang, J.; Bowman, S.R.; Perez, E. Pretraining Language Models with Human Preferences, 2023, [[arXiv:cs.CL/2302.08582](https://arxiv.org/abs/2302.08582)].
70. Chan, J.S.; Pieler, M.; Jao, J.; Scheurer, J.; Perez, E. Few-shot Adaptation Works with Unpredictable Data, 2022, [[arXiv:cs.CL/2208.01009](https://arxiv.org/abs/2208.01009)].
71. Zou, A.; Wang, Z.; Kolter, J.Z.; Fredrikson, M. Universal and Transferable Adversarial Attacks on Aligned Language Models, 2023, [[arXiv:cs.CL/2307.15043](https://arxiv.org/abs/2307.15043)].
72. Shavit, Y. What does it take to catch a Chinchilla? Verifying Rules on Large-Scale Neural Network Training via Compute Monitoring, 2023, [[arXiv:cs.LG/2303.11341](https://arxiv.org/abs/2303.11341)].
73. Everitt, T.; Carey, R.; Langlois, E.; Ortega, P.A.; Legg, S. Agent Incentives: A Causal Perspective, 2021, [[arXiv:cs.AI/2102.01685](https://arxiv.org/abs/2102.01685)].

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.