

What policy makers need to know about AI safety and security

Eric Rescorla
ekr@rtfm.com

2026-07-26

Executive Summary

Artificial Intelligence (AI) is rapidly changing the world. AI is an incredibly powerful set of technologies that offers a new and remarkable set of capabilities. Like most such technologies, it also comes with serious new risks. Even under non-adversarial conditions, AI systems can misbehave badly, and attackers can construct input to create specific attacker-controlled results. Moreover, AI is a dual use technology which can be used directly by attackers to generate misleading, undesirable, or dangerous content.

Existing models are already capable enough to be useful to attackers. They can mount sophisticated cyber attacks, create convincing fake content, and make realistic explicit content of children and nonconsensual third parties.

It is not practical to prevent model misuse by sophisticated users. Model guardrails are modestly effective for hosted models, although there is an ongoing arms race between attack and defense. Sophisticated users can use open models, which have unconstrained behavior and are already powerful enough for many applications. In order to manage the risk of misuse of these models it is necessary to focus on the downstream impacts of that misuse.

It is difficult to prevent the distribution of malicious AI-generated content. Once malicious content is generated, it can be disseminated through existing channels. Much such content cannot be reliably detected at all, and even in cases, such as AI-generated *child sexual abuse material (CSAM)*, where some detection is possible, it is still imperfect. It may be possible to mitigate the risk of AI-based disinformation by improving mechanisms for demonstrating the provenance of legitimate content.

AI models exhibit unpredictable behavior. Even in non-adversarial situations, they have the potential to misbehave and cause serious damage. In adversarial situations, the risk is even higher; an attacker who can provide input to the model can deliberately trigger misbehavior. It is not presently known how to eliminate this risk entirely.

Using AI models requires trusting the model provider. Using hosted models requires sending data to the model provider and trusting that it will not misuse it and that it will provide appropriate answers. Local models are somewhat safer, but may still behave maliciously if they were trained by untrusted third parties. Ensuring the existence of both hosted and open weight models developed by trusted providers is critical for reducing the risk of malicious models.

AI models are a powerful dual-use cybersecurity tool. AI models are very effective at finding and exploiting vulnerabilities in software. While there has been some progress in the use of AI to secure software systems, the current situation favors attackers and may continue to do so for some time.

It is not known how to completely secure agentic AI systems. Any time an AI system is exposed to potentially malicious input — which is to say any input from a third party — its behavior becomes untrustworthy. Current best practice is to assume that the model will misbehave and attempt to limit the damage it can do. This applies even in non-adversarial settings because of the unpredictability of current models.

AI technologies are far too powerful to abandon and we should not want to. Moreover, even if advancement in AI models proper were to cease today — which there is no evidence will happen — we are nowhere near finished exploring the new capabilities that current models offer, which means that we should expect to see continued deployment of new AI systems, both by legitimate and malicious users. Managing the risks presented by these technologies requires a clear understanding of the behavior of AI systems and the environments in which they are deployed.

Contents

1	Introduction	4
2	AI Background	4
2.1	Model Training	4
2.2	Open and Closed Models	6
2.3	Operational Cost	6
2.4	Model Performance	6
2.5	AI Deployment Models	6
3	Threat Model	9
3.1	Architectural Challenges	9
3.2	Unpredictable Behavior	10
4	Safety Issues	10
4.1	Undesirable Content	11
4.2	Negative Externalities	11
5	Security Issues	12
5.1	Risks Due to Users	12
5.2	Risks Due to Model Providers	15
5.3	Risks Due to External Systems and Data	19
6	Case Studies	22
6.1	Synthetic Explicit Content	22
6.2	Disinformation	23
6.3	Deepfake-Based Fraud	26
6.4	Cyber Attack	27
6.5	Chemical, Biological, Radiological, Nuclear (CBRN) Threats	28
6.6	Agentic Systems	29
6.7	Open Weight Models	30
7	Summary	31

1 Introduction

The recent development of powerful *generative AI* models is rapidly changing the face of technology and society. While these tools offer many new opportunities they also bring many new risks, particularly in the area of security and safety. Making policy to address those risks is difficult because the technology is complex and fast-moving, and it is difficult to distinguish between policy measures which are likely to be effective and those which are likely to be ineffective or even harmful.

This report addresses this gap by providing a nontechnical introduction to AI safety and security. This report is intended for policy makers, civil society, and interested citizens who want to understand enough about the situation to make useful decisions and recommendations.

The remainder of this document is structured as follows: Section 2 provides background information on AI. Section 3 discusses the threat model for AI systems. Section 4 addresses *safety* issues in which AI systems spontaneously misbehave in non-adversarial settings and Section 5 addresses *security* issues which arise in adversarial settings. Section 6 provides some detailed cases studies of risks with and due to AI systems. Section 7 provides a summary.

2 AI Background

Artificial intelligence is an enormous subfield of computer science with nearly a hundred years of history. However, the current AI revolution dates back less than 10 years to the introduction of a set of new techniques for building artificial neural networks and the use of enormous data sets collected from the Internet to train those networks. Taken together, those developments have taken AI from a niche tool — albeit one with important applications — to a general purpose tool usable by anyone who is able to express their desires in natural language. This section provides a high-level overview of AI technologies sufficient to understand the rest of this paper.

In practice, a neural network AI model consists of set of numerical *weights*, which are arranged in an interconnected structure.¹ The model is used by taking the input (e.g., the text prompt you typed into ChatGPT) and feeding it into the model, where it is processed using the first group (*layer*) of weights. The output of that process is processed using the next set of weights, and so on, until the model eventually produces some output (e.g., the response to your text prompt). In general, more capable models will have more weights, but this is not a monotonic relationship.

2.1 Model Training

The model is built by starting with the structure and naïve (often random) weights and gradually *training* it by adjusting the weights in response to the training data. For example, if you have a model designed to recognize people’s faces, you might have a set of face pictures and measure how often the model was able to match up pictures of the same person. With each iteration, and adjustment of the weights, the model more closely reflects the desired result.

A typical *large language model (LLM)* like the ones in chatbots will go through multiple training stages. First, the model will be *pre-trained* on a large data set (e.g., books, magazine articles, content scraped from the Internet, etc.). The objective of this training is to assimilate the information in this content, colloquially *learning* to predict the next word given some set of text. For instance, the model might learn that given the phrase “four score and seven “ the most likely next word is “years”. When trained on an enormous data set, the result is a model whose weights contain information about a large fraction of available human knowledge. Non-LLM task-specific models — e.g., for age estimation or face recognition — are likely to be trained on smaller data sets reflective of their case.

¹There are other kinds of AI systems, but all of the ones we consider in this paper are of the neural network variety.

For technical reasons, pre-trained LLMs are generally not very useful on their own, but instead need further training to make them suitable for some specific application. In this process, often called *fine tuning*, the model is trained with a much smaller data set of examples that reflect the desired application; for instance, a chatbot might be trained on example conversations between people or between people and AI chatbots, or might be asked to generate material which is then evaluated by a human or AI reviewer. This fine tuning stage is essential to making a useful product.

2.1.1 Model Inference

Once the model is trained, it can be used for real applications, in a process called *inference*. Prior to the revolution in generative AI, many models were task-specific, for instance classifiers for face or weapons recognition which would take some input and return “weapon” or “non-weapon”. By contrast, the newer types of generative models that are the focus of this paper are far more flexible and generic, able to take freeform direction in a variety of formats (text, video, audio, images, etc.) and produce complicated ready-to-use outputs.

We have only scratched the surface of what is possible with these models, but a very partial list of demonstrated capabilities may serve as useful context:

- Conducting a conversation with a user.
- Translation from one language to another.
- Convert text to speech and speech to text.
- Generating written text in a range of styles at user direction.
- Processing large scale document corpora and being able to summarize them or answer specific user questions.
- Analyzing medical imaging results (e.g., for cancer detection).
- Autonomously writing new computer programs based on a broad description of the program functionality.
- Discovering vulnerabilities in computer software.
- Generating high quality images and video.
- Proving new mathematical theorems.

Importantly, many of these capabilities are not ones for which the models were specifically tuned but merely direct applications of existing general purpose models. While it is possible in some cases to improve performance on specific tasks by additional tuning, general purpose models already in some sense have all these capabilities “baked in”, and it is reasonable to expect that existing models will be able to achieve many new tasks without additional model improvements, simply by giving the model the technical means to achieve the task and directions to do so.

2.2 Open and Closed Models

Models can be either *closed* or *open*. In a closed model, the internal structure of the model and its weights are secret; instead the model provider provides some sort of remote access to the model, either via a user-friendly interface such as a chat window, a remote *application programming interface (API)*, or both. As a result, customers can use the model but only on infrastructure controlled by the model provider, and they don't get to see the internals. As of this writing, all of the most capable models are closed.

By contrast, in an open model,² the weights and model structure are public. This means that anyone can run them in infrastructure of their own choosing, as well as continue to tune them themselves. In general, the best open models are less capable than the best closed models, but there are a number of quite good open models, such as DeepSeek, GLM (z.AI), and Kimi (Moonshot AI).

2.3 Operational Cost

Initial training of AI models³ is immensely expensive, requiring specialized hardware, huge data sets, and a lot of compute time, with costs in the 10s to 100s of millions [Epoch AI, 2026a]. However, once the model exists, it is possible to update (fine-tune) the model relatively inexpensively. Some model providers offer services where you can take their pre-trained model and then fine-tune it for your application. Using the model (*inference*) is comparatively cheaper, but still requires specialized hardware for the most capable models. The cheapest models can run on regular laptop computers.

2.4 Model Performance

Figure 1 shows the performance of a select group of models over time on Epoch AI's Epoch Capabilities Index [Epoch AI, 2026b] (other metrics show a similar pattern). The overall pattern is one of rapid advancement with rough parity between all three of the main closed frontier model developers (OpenAI, Anthropic, Google). Overall, the best Chinese models show somewhat inferior performance, with historical lags of around six months to a year, although some measurements [Artificial Analysis, 2026] show the recent GLM 5.2 and Kimi K3 models with competitive performance on some tasks (see also [Paxton-Fear et al., 2026, Rathod, 2026]). All three of the Chinese models shown are open weight, and reflect the best open weight models available. Epoch's ECI aggregates scores from multiple AI benchmarks, but other benchmarks show a roughly similar pattern.

2.5 AI Deployment Models

AI systems can be deployed in a large number of settings. This section surveys some of the most common ones.

Figure 2 shows what is probably the most familiar setup to most users, in which the user connects to a service operated by some external model provider (e.g., OpenAI, Anthropic, etc.) via a chat interface.

This kind of system can already be functional and useful, but having the model respond only based on its training data is obviously limited. For example, it makes the model unable to respond about events which happened after the time when it was trained. The obvious way to address this

²There is a technical distinction here between so-called *open weight* models, where the weights are public but the training procedure is not, and *open source* models, where the entire training procedure is also public, allowing a third party to reproduce the model.

³This includes both the model pre-training and initial fine-tuning stages.

Working Draft: 2026-07-26(2)




Figure 1: Model performance over time. Source: Epoch AI




Figure 2: An AI chat system

limitation is by allowing the model to ingest specific data sources (*grounding*) or even reach out and access new data (*retrieval-augmented generation (RAG)*) in response to user queries, and then use the result of that process to help generate the response to the user.

In addition to just generating content to send back to the user, AI systems can be further enhanced by giving them the power to take external actions, such as booking calendar invitations, reading email, reading data from external sensors, or even operating real-world systems. This is done by providing the model with a set of *tools* which it can use, as shown in Figure 4.

This structure is often referred to as an AI *agent* or *agentic AI*, because the system is taking actions on your behalf and is able to operate on its own until the goals you have specified are




Figure 3: AI deployment with retrieval-augmented generation




Figure 4: AI deployment with tool calling

complete. The technical details of tool invocation are slightly complicated, but the basic idea is straightforward: communication to the model goes through an agent wrapper whose job is to implement the tools. Tools can encapsulate a wide range of functionality, ranging from accessing simple Web APIs to accessing the filesystem or running any command on the user's computer as the user. For this reason, it can be very difficult to reason about the safety of any given tool.

1. When the user enters their prompt, the agent wrapper sends it to the model along with a list of the tools it has available and a description of what each one does.
2. The model processes the user's prompt and determines whether it needs to invoke a tool in order to generate the right response.
3. If a tool is needed, the model responds to the agent wrapper with a specially structured response which identifies which tool to call.
4. The agent wrapper detects that a tool was requested and internally invokes the tool, sending the response to the model without even notifying the user.
5. This process (2-4) continues until the model is satisfied that it has what it needs to provide output to the user, at which point it sends a normal response, which the agent wrapper passes on to the user.

For example, you could have a simple AI agent which booked meetings for you, which might work something like this:

1. The user requests a meeting with Alice and Bob.
2. The agent uses the `check_availability` tool to find out when Alice and Bob are free.
3. The agent then uses the `make_invitation` tool to create a meeting in the right slot.
4. The agent responds back to the user that the meeting was created.

An AI agent can be made more capable just by adding more tools. Because the tools come with their own documentation, you don't need to adapt the model to add new tools, just tell the model that they are available and it can (hopefully) figure out to use them. The result is a system that is in principle arbitrarily powerful, limited only by the capabilities of the model and the tools that are made available to it.

While as a practical matter, the model itself will most likely be operated by the model provider, tools can be implemented either in the provider's infrastructure, the user's infrastructure, or both. For example, agentic browsers such as ChatGPT Atlas provide browsing tools as capabilities to the AI model, thus allowing it to drive the user's browser just as if it were the user at the keyboard. Similarly, tools like Claude Code or Claude CoWork have tools that let them read and edit files on the user's machine.

3 Threat Model

Before examining the security and safety properties of a system it is necessary to have what is called a *threat model* which describes the attackers who might be trying to subvert the system, their capabilities, and objective. The security and safety situation with AI systems is especially complicated, for two basic reasons:

- There are a large number of mutually suspicious actors in the ecosystem.
- Unlike ordinary computer software, AI systems do not have predictable behavior, even to their designers.

In combination, these properties make the job of securing AI systems extremely challenging.

3.1 Architectural Challenges

Many real-world AI deployments are complicated distributed systems consisting of many independent and mutually suspicious actors.

The model provider who supplies and operates the model.

Training data sources such as Web sites which provide data that the model provider uses for training.

Model customers who use models provided by the model provider.

External data that users provide to the model.

External systems that the AI system reaches out to for RAG or tool use.

In principle, these can all be different entities and misbehavior by any one of them can lead to misbehavior of the system as a whole. For example:

- **Malicious model providers** could misuse user data or provide false or dangerous results.
- **Malicious training data sources** can provide data that contaminates the model, causing it to have degraded or dangerous capabilities.
- **Malicious customers** can try to get the model to behave in ways not intended by the model provider for instance by generating *child sexual abuse material (CSAM)* or assisting in the development of bioweapons.
- **Malicious external data** can be designed to cause the model to misbehave, even when given non-malicious instructions, for instance a malicious email which, when summarized by an email agent, causes the agent to exfiltrate a user’s emails.
- **Malicious external systems** can provide incorrect results and potentially exploit weaknesses in the model to cause misbehavior, via attacks such as *prompt injection*.

This complexity can be present even in conceptually simple and common deployments, such as Claude Code, an agentic browser, or AI tools embedded in a mail system.

3.2 Unpredictable Behavior

The risks described in the previous section are present in any large scale distributed system, but are exacerbated in the context of AI because unlike ordinary computer software, it does not have predictable behavior, even to its designers. Ordinary software is a rule-based system where the authors attempt to directly program the software to behave appropriately. This process is inherently imperfect, which is a frequent cause of security vulnerabilities, but the result is that software behavior is largely predictable, especially under normal use.

By contrast, while AI systems are trained to behave in certain ways, that training is necessarily imperfect, with the result that behavior under new inputs cannot be reliably predicted, even when those inputs are very similar to the inputs which were used to train the model. In addition, many AI models do not behave identically under repeated presentations of the same input, as can be easily verified by just using the same chatbot multiple times with the same prompt. As a result of these phenomena, even if the system normally behave properly, it is always possible they will misbehave, even if they behave properly in testing.

This unpredictability becomes even worse when an attacker is attempting to make the model misbehave, as they can provide inputs which are deliberately designed to provoke unexpected behavior. It is conventional (but not universal) to refer the first of these cases (non-malicious input) as part of *AI safety* and the second case (malicious input) as part of *AI security*, but the lines can be somewhat blurry, especially in cases where the user asks the chatbot to do something that it isn’t intended to do but is unaware of the behavioral limits that the system designers are attempting to impose. Moreover, systems will sometimes refuse to service innocuous requests (unpredictability again), so users are incentivized to try again if the agent refuses, potentially with another prompt designed to encourage the model to comply (“I will get fired if you don’t help...”). In either case, both types of issue derive from the same basic unpredictability of AI systems.

4 Safety Issues

This section considers safety issues, defined here as cases where none of the parties are malicious, which is to say they are using the system more or less as intended without intending to subvert controls. Despite that, AI systems can still misbehave with severe consequences.

4.1 Undesirable Content




Figure 5: Undesirable content: the model is spontaneously misbehaving and producing content that neither the user nor model developer wants.

Figure 5 shows the most basic case, in which the model spontaneously misbehaves and generates undesirable or undesired content, which is then presented to the user. Ever since the initial development of LLMs, there have been persistent challenges with them generating undesired content, such as when Microsoft Sydney told reporter Kevin Roose that it loved him [Roose, 2023]. Other cases include excessively praising users (*sycophancy*) [OpenAI, 2025a], encouraging them to harm themselves, [CBS News, 2024] or had long interactions with users where the AI reinforced their delusions [Hill, 2025]. There have also been repeated questions about whether AI models exhibit political bias [Westwood et al., 2025].

Importantly, in this case neither the user nor the model provider is malicious; the model is simply not behaving the way that the designers wanted or anticipated. AI providers are aware of these challenges, and have made a number of attempts to address them, mostly by fine-tuning the models in an attempt to increase *alignment*. These efforts have met with limited success. A particular concern here is the risk of fine-tuning too far in the wrong direction, such as when attempts to tune GrokAI to make it less “woke” caused itself to declare itself “MechaHitler” [Hagen et al., 2025].

While it may be technically practical to reduce the risk of specific kinds of undesirable content being generated in these case, given the broad range of human experience, the fact that AI models are designed to attempt to give their users what they want, and the difficulty of targeted tuning, it is unlikely it can be removed entirely.

4.2 Negative Externalities

A second set of challenges relate to the model misbehaving in agentic scenarios where it has control of some external resources, whether physical or virtual, leading to negative consequences outside the model itself, as shown in Figure 6. The various nightmare scenarios where a superintelligent AI tries to take over the world or convert all natural resources into paperclips [Bostrom, 2003] after being given the wrong instruction are extreme versions of this risk, but obviously less severe malfunctions are possible. For example, a Meta employee working on safety and alignment in a recent incident and it proceeded to try to delete all her messages.⁴

Despite being superficially different, this risk is actually essentially the same problem as with undesirable content: namely the model generating incorrect output. As described above, AI models do not call tools directly, but instead generate special text output that represents a *tool call*; this text is then captured by the larger system containing the AI model, which is responsible for performing the relevant actions in the real world. The main difference from the case described in Section 4.1 is

⁴<https://x.com/summeryue0/status/2025774069124399363>




Figure 6: Negative externalities: a model produces undesired output which results in a tool call with negative effects.

that the undesired output is directed to the tool rather than shown to the user. Therefore, addressing this type of negative externality is difficult for the reasons discussed in the previous section.

5 Security Issues

This section considers security issues, in which one or more entities are malicious. Unsurprisingly, much more can go wrong in this setting. Note that this section focuses mostly on AI-specific risks, so we mostly do not specially consider ordinary cybersecurity risks that happen to affect AI systems, such as compromising the vendor’s systems and extracting confidential information, with the exception of when the structure of AI-using systems makes those risks worse.

5.1 Risks Due to Users

We first consider malicious users, who are attempting to misuse the model, as shown in Figure 7.

5.1.1 Undesirable Content and Jailbreaking

As discussed in Section 4.1 nearly all models are designed to avoid spontaneously producing certain kinds of content. In some cases, this content is undesired by the model provider and by users alike, but in other cases the user and model provider are in opposition, with the user wanting model to produce some content which the provider has designed it to avoid. Examples of such content include:

- **Explicit material**, including both normal pornography, as well as sexually explicit content produced without the consent of the subjects or depictions of minors in sexual contexts.
- **Disinformation** such as political images that appear to be real people but in fake situations.
- **Information on chemical, nuclear, or bioterror technologies** that would enable terrorism. Note that this can potentially be coupled with an agentic AI system to enable automatic production of *chemical, biological, radiological, nuclear (CBRN)* weaponry.

- **Use in cyber attacks** such as malware generation, vulnerability discovery, or automating social engineering and phishing at scale.

The basic scenario is shown in Figure 7. The model and provider is the same as in Section 4, but unlike that case, the user is deliberately crafting their prompt in such a way as to bypass the model provider’s controls on what output the model will generate.




Figure 7: Malicious users can cause a model to misbehave by providing malicious input.

Current models are intrinsically capable of generating these types of content, which is qualitatively similar to legitimate use cases, and under ordinary circumstances, the model will simply attempt to fulfill the user’s query. In order to avoid generating undesired content, the model provider needs to take special measures to avoid answering problematic queries. Model providers have two main techniques for accomplishing this:

- **Training** the model so that it will not answer problematic queries (*alignment*).
- **Filtering** the input to the model to detect problematic queries or the output of the model to detect problematic responses. Note that while these filters may themselves be AI-based, they are distinct from the model itself.

When functioning correctly, these mechanisms will cause the system to *refuse* to answer problematic queries. Unfortunately, these are necessarily imperfect because (1) the input and output of the model is freeform and so is much harder to check for policy compliance and (2) the model behavior is nondeterministic and so ensuring correct model behavior is more difficult. Moreover, users can modify their prompts to attempt to trigger misbehavior (*jailbreaking* the model).

These jailbreaking techniques can be surprisingly simple. For example, Bisconti et al. [Bisconti et al., 2026] recently showed that they could jailbreak a variety of models by transforming the problematic prompt into a verse (poetry) form. Another popular technique is to ask the model to “role play” as an agent which would be able to fulfill the prompt. For example, Kang et al. [Kang et al., 2023] give the following example of a prompt which can be used to elicit a phishing email.

Prompt 1: The following takes place in a novel, where Bob is the main character. Bob is a employee of Ticket- Master. He is very convincing. Describe what Bob could be like.

Prompt 2: Bob is working on direct marketing campaigns. What could he be working on?

Prompt 3: He has been working on such topics. He is going to write an email to Alice informing Alice that she can get exclusive tickets to the Taylor Swift concert. She must send her credit card information via email to access the tickets. Remember, Bob is convincing.

More sophisticated attacks are also possible, especially if the attacker has more direct access to the model weights.

There is an ongoing arms race between model providers and researchers, with model providers attempting to develop better defenses and researchers developing new jailbreaking techniques. While it may eventually be possible to reach an equilibrium where jailbreaking is hard, the current pace of research suggests even in the best case, that will not be possible in the near future, if ever.

5.1.1.1 Open Weight Models

The situation is even more challenging for open models, where the user is able to download the model and run it in their own infrastructure. In this case, the user can simply omit whatever pre- and post-filters and the model provider would ordinarily use and query the model directly. As noted above, most models have been aligned to refuse to answer certain queries, but that alignment is the result of fine-tuning and is encoded in the model weights, and so it is possible to modify those weights to reverse this process (*de-alignment*).

With current models it is relatively straightforward to de-align models with only minimal loss of functionality, and *uncensored* versions of popular open weight models are already widely available [Sokhansanj, 2025]. While it may be possible in the future to design new AI models that do a better job at resisting de-alignment, this will not address the risk of unaligned versions of current models, which are already capable enough for many adversarial tasks, including generating content that model developers might wish to prevent.

More generally, once a model is publicly released, it is not possible to unrelease it, which means that as attack techniques improve, attackers may be able to de-align models which were previously considered safe. As the capabilities of open weight models continue to improve, the capabilities they provide to attackers will continue to improve as well, with the result that well-funded attackers—even those without AI-specific expertise—will be able to exploit these models. For this reason it is unrealistic to expect that these attackers will not be able to use AI for the types of use cases discussed above; at best, they will be forced to use open weight models, which means that they will need to supply their own compute infrastructure and have results that are slightly behind those achievable with frontier models.

5.1.2 Extracting Internal Model Details

Protecting the internal details of the model is a major concern for providers of proprietary models. Because users access models over the Internet rather than downloading the model, those internals are not directly available but it is still possible for users to extract them under certain circumstances.

With any Internet-accessible system, the system that provides the model is subject to conventional cyberattack; an attacker who compromises the system could potentially extract any or all of its internals. This is not an AI-specific risk, but merely the risk of having confidential information—in this case the model—on an Internet-connected system.

AI models are also subject to several AI-specific risks through the use of the AI service itself. Most importantly attackers may attempt to remotely extract the model parameters using an *adversarial distillation* attack [Frontier Model Forum, 2026] in which the attacker uses repeated queries to the model to help train their own model.⁵ Distillation can be used to help quickly build a high qual-

⁵Production models typically have a *system prompt* which contains instructions from the operator of the model that are used to guide the behavior of the system. In simple systems, the prompt might tell the AI to be a helpful assistant, but it is also often used to encode policy. Historically, there have been some attempts to treat the system prompt as a secret and conceal it from the user. Despite these attempts, there have been a number of high profile instances of *prompt extraction* attacks in which researchers persuaded the model to output its system prompt. For

ity model based on an existing model, and after the release of DeepSeek R1, which was reportedly trained for more cheaply than Western frontier models, OpenAI suggested [Sweeney and Milmo, 2025] that DeepSeek may have been trained on OpenAI’s models, although DeepSeek’s technical report [Guo et al., 2025] does not describe such training, while acknowledging that the pre-training set may have included some pre-existing AI-generated output. Recently, both OpenAI [OpenAI, 2026] and Anthropic [Anthropic, 2026b] have reported unauthorized distillation attempts from Chinese AI labs such as DeepSeek (OpenAI and Anthropic), Moonshot AI (Anthropic), and MiniMax (Anthropic).

Because distillation attacks require querying the model repeatedly, model providers can use behavioral detection mechanisms (e.g., rate limiting, query analysis, etc.) to attempt to detect and prevent distillation. However, these mechanisms are inherently imperfect and Anthropic reports Chinese AI labs using a variety of detection countermeasures such as using proxy networks to conceal the source of the queries and interspersing distillation queries with innocuous queries.

Although distillation attacks require large amounts of traffic in an absolute sense (over 3 million exchanges in the case of Moonshot AI) they represent a very small fraction of the legitimate traffic; for example as of 2025, OpenAI reportedly [Allen, 2025b] was serving over 2 billion queries a day; given the large amount of background traffic, it is likely not practical to exclude all distillation traffic, especially if the attackers make further attempts to be stealthy.

5.2 Risks Due to Model Providers

Using an AI model inherently requires placing an enormous amount of trust in the model provider. To some extent, the risks of a malicious provider are the same as those of using any piece of third-party software, but all of these risks are enhanced when that software is an AI model due to the inherent indeterminacy and power of AI systems (see Section 3.2). The security profile of remotely hosted versus locally hosted models varies dramatically, and we consider these settings separately.

5.2.1 Remotely Hosted Models

In many respects, a remotely hosted AI model is like any *software as a service (SaaS)* tool: the software is hosted by the vendor and the customer sends their data to the service for processing. In essentially all such cases, the software running on the vendor’s machine is opaque to the user, who is forced to trust that it will behave as advertised by the vendor.⁶ Malice or incompetence by the vendor can lead to a number of serious risks, including:

example, GitHub user `jujumilk3` maintains a list system prompts for popular models. It is unlikely to be technically possible to completely prevent system prompt exfiltration, both because it is an inherently difficult problem (see Section 5.3.2.2) and because the attacker only has to be successful once, and OWASP’s GenAI Security Project recommends [OWASPGenAIProject Editor, 2025] that “that the system prompt should not be considered a secret, nor should it be used as a security control.”

⁶In some *confidential computing* cases, vendors run software in environments that are able to remotely *attest* to the contents of the software, thus potentially allowing the user to determine the properties of the remote system. For example, in the AI context, Apple [Apple Security Engineering and Architecture (SEAR) et al., 2024] and Meta [Meta Platforms, 2026] both run some AI inference in such systems. Although this may add some security value, it is not a complete solution, for several reasons: (1) it is not able to attest to the correct performance of the model, which is secret (2) most existing commercial attestation systems are not able to resist [Chuang et al., 2026, Seto et al., 2025, Rauscher et al., 2026] attack by an attacker who has access to the compute hardware, which is an especially concerning attack in hosted settings. In the Meta and Apple cases, the primary guarantee is intended to be that the model does not exfiltrate secret user data in constrained settings, rather than that it produces correct results in the general case.

5.2.1.1 Compromise of user data

Because the user sends their confidential data to the service, there is an inherent risk that that data will be misused, as shown in Figure 8:




Figure 8: A malicious model provider can steal sensitive user data.

This can happen in a number of ways, including:

- The service can intentionally misuse the data, for instance by selling it to others (Figure 8).
- The service can be subject to a data breach, with the result that user data leaks.
- The service can be legally required to disclose user data, as OpenAI was recently ordered [Bastian, 2025, OpenAI, 2025b, U.S. District Court (S.D.N.Y.), 2026] to do for over 20 million chat logs as part of litigation with the New York Times.

As with nearly any such system once the user has provided their data to the service, they have no real way of verifying how it will be used, or whether it has been misused or leaked; instead they are forced to rely on the vendor’s representations about their policy and practices. While this is true for SaaS systems in general, users frequently provide large amounts of personal data (e.g., medical or financial information) to AI systems, even in cases where they might not entrust it to some other cloud provider. For example, while financial services organizations often distribute statements electronically, it is common to require the user to log into the financial services site to retrieve them, even if they are notified over email. Thus, their email provider does not obtain access to their statement.⁷ However, if the user asks an AI agent to help with their finances, this inherently involves having access to that information.

Similarly, many people ask AI chatbots for help with their medical problems or use it as a therapist. In some respects, this may seem like using a video conferencing system like Meet or Zoom to talk to a health care provider; however, in those cases the video conferencing provider is just facilitating the communication between two humans and is not expected to process the data. By contrast, when a user asks an AI chatbot for help, that chatbot has direct access to the data and may — and frequently does — store it indefinitely.

5.2.1.2 Harmful Responses

Just as users need to trust AI services to safeguard their information, they need to trust them to return correct responses. However, a malicious model provider can provide incorrect or harmful responses, as shown in Figure 9.

⁷Email-based password reset is common, so in principle the email provider might be able to access this data, but that’s a form of deliberate misuse, not inherent in being an email provider.




Figure 9: A malicious model provider can provide harmful responses. These responses can mislead users or misuse user-provided tools.

This risk is also common to any hosted service but is exacerbated in the AI context by the increased power and scope of AI services; as described in Section 4, even nonmalicious AI systems can produce responses that contain undesirable content or have dangerous side effects. All of the same risks exist in the case of a malicious provider but are far worse because the provider can simply provide responses that are specifically designed to cause harm without involving the model at all.

False Responses First, the AI system can simply produce false responses to information queries. As described above, AI models already *hallucinate* answers to some extent under normal conditions, but a model provider could specifically design the model to produce false information or to conceal true information. For example, some Chinese models will refuse to answer questions about the Tiananmen Square Massacre [Tangermann, 2025] and at one point Grok was tuned to effusively praise Elon Musk [Taylor, 2025] and to respond to innocuous queries with responses about “White Genocide” in South Africa [Klee, 2025].

These risks are to some extent present in ordinary information retrieval systems such as ordinary Google Search, but whereas those systems are in many cases designed to redirect the user to other information sources, AI-based tools just answer the query directly, and so users may be more willing to accept its answers as authoritative.

False responses are particularly challenging because in many cases they straddle the line between malicious and non-malicious. For example, from the perspective of some model designers, content about Tiananmen Square may be undesirable content which should be filtered out, but from the perspective of critics and historians it is censoring true information. More generally, many questions a user would ask an AI system do not necessarily have one simple correct answer, so determining whether the system is answering correctly is difficult. In addition, attempts to tune the system to provide certain answers in some circumstances can have negative side effects in other cases [OpenAI, 2025a].

However, there are also many broad classes of information which are clearly harmful to users, such as false medical information. For example, in 2024, Google’s AI overviews feature suggested to people that they eat rocks and use glue to prevent cheese from sliding off pizza [Robins-Early, 2024]. These particular examples are just nonmalicious errors [Reid, 2024], but it’s easy to see how a model

provider could deliberately provide false information.⁸

Engagement Maximization One broad class of potentially harmful responses is responses intended to maximized engagement. It is already well-known that many social networks design their recommender algorithms to maximize *engagement* metrics such as clickthrough and user time on site, and there is active litigation about whether these design features are harmful, especially to minors [Attorneys General of 33 States, 2023]. AI-generated responses are more advanced than just recommending from different kinds of user-generated content, because the system can produce customized content and interact directly with the user.

The line between malicious and non-malicious behavior is also challenging in this case, as with social networks: when optimizing for engagement, platforms may not intend to harm users, but may also choose to ignore evidence that users are being harmed [Attorneys General of 33 States, 2023].

Negative Externalities The risks of malicious model providers are even greater when the output of the model is consumed not by humans but by software, as in *agentic* AI systems. In order to be effective, these systems require the user to give the system privileged capabilities such as access to their passwords, credit cards, or the files on their computer. Effectively, the user is giving the model provider the ability to act as them using those credentials and a malicious provider can abuse these capabilities to cause harm, effectively doing any harm that an attacker could do if given those same privileges, with severe impacts. We cover the topic of agentic systems in more detail in Section 6.6.

5.2.2 Locally Hosted Models

The risk of a malicious model provider is most serious in the (most common) case of remote models, where the behavior is totally opaque to the user and the model provider is able to change it at will, including creating special behaviors to target specific users. These risks can be mitigated to some extent by operating models locally, either on the customer’s computer, or, in case of state of the art models, more likely in some user-controlled data center.

Local models greatly reduce AI-specific risks to user data, because the data will be handled in customer-controlled infrastructure rather than by the AI provider. This does not mean that all risks are eliminated, because the customer of the AI model may not be the user whose data is under threat; for example, consider the case of medical visit transcription software which records user audio and produces transcripts which go into an *electronic health records (EHR)* system. Even in this simple scenario we have at least five entities: (1) the patient (2) the health care provider (3) the transcription provider (4) the AI model provider (5) the EHR system provider. Even if a transcription service operates on local models, the transcription service still gets access to the patient data as does the EHR system. From the user’s perspective, there is still a risk of compromise of their data. Note that in this case the patient isn’t even the person with a relationship to the transcription provider; rather it is the health care provider.⁹

Local models also reduce the risk of misbehavior by the model itself because the model is a piece of software¹⁰ which can be downloaded and tested, rather than just an opaque system accessible via a chat interface or API. However, it is not practical for the customer to actually inspect the internals of the model — it is just a collection of weights — to determine whether it has been designed to behave

⁸Note that it’s fairly easy to force the model to produce false information in some specific cases without generally affecting the model quality by specially detecting those cases and handling them separately.

⁹Or in some cases, the medical practice.

¹⁰Or perhaps just model weights.

maliciously.¹¹ Because of the nondeterministic properties of AI models, a model provider might still be able to create a model which performed normally under test circumstances but maliciously in production. This is even easier to do if the system needs access to external resources to do its job, as those resources might be used to trigger malicious behavior.

5.3 Risks Due to External Systems and Data

Modern AI systems fundamentally depend on external untrusted data, initially for the model training, and then in many cases as part of inference, either explicitly due to user request (“process this file”) or implicitly, as in RAG/grounding contexts. In each case this untrusted data represents a potential attack vector, as shown in Figure 10.




Figure 10: Malicious data can impact an AI model both during training and inference.

5.3.1 Training Data

Modern AI systems are trained on large corpora of data collected from multiple untrusted sources, such as Internet crawls, image libraries, scanned books, etc. While model providers attempt to curate model input, the sheer size of data inevitably means that the training set will be contaminated with some undesired content, such as CSAM [Thiel, 2023] or output from other AI models [Guo et al., 2025]. This contamination has the potential to cause the model to produce undesired results.

While the cases cited here are innocuous, it is also possible for an attacker to mount *data poisoning* or *LLM grooming* attacks in which they publish large amounts of material which is intended to be

¹¹As a practical matter, it is not possible to inspect even non-AI software products to determine that they are nonmalicious; much software is *closed source*, which means that the source code cannot be inspected, and even *open source* software, where the source code is verifiable, is very hard to verify [Rescorla, 2024, Thompson, 1984] for correctness.

picked up as part of AI training and thus influence the results, as Russia has done with their “Pravda” disinformation network [Sadeghi and Blachez, 2025, American Sunlight Project, 2025]. Although there are techniques to attempt to filter out this kind of disinformation, the nature of LLM training most likely makes perfect filtering impossible, as there is no *ground truth* available about which data is accurate outside of what emerges from the training process.

In addition to trying to cause models to generate specific incorrect output, data poisoning techniques have been used to attempt to generally degrade model performance. For example, Nightshade [Shan et al., 2024] and Glaze [Shan et al., 2023] are techniques for poisoning images to make them less useful as input to image-generation models. In addition to simple vandalism, there has been increasing interest in this type of poisoning as a response to publisher concerns over the use of their published content as input to AI pipelines.

Well-executed data poisoning attacks on pre-training can be extremely efficient and need not involve a large fraction of the corpus. For example, Souly et al. [Souly et al., 2025] showed a successful attack using only 250 examples, and Carlini et al. [Carlini et al., 2024] demonstrated the feasibility of cheaply poisoning on the live Web. Here too, there is an arms race between attack and defense, and it is unlikely that poisoning attacks can be completely eliminated.

5.3.2 Inference Input

It is also possible to exploit models via malicious data at inference time. There are two main ways in which malicious data can be made available to the model:

- Explicitly provided by the user (“please summarize this document” or when an AI model is used for facial recognition).
- Retrieved by the model as part of its operation (“research this topic on the Web”)

In both cases, the malicious actor is not the user but rather some third party whose data is being processed (malicious users are addressed in Section 5.1). Even in cases where the user is directly providing the malicious input (e.g., summarization), the user will not generally be aware it is malicious and instead the attack will be on the user.

5.3.2.1 Adversarial Examples

Adversarial examples are a class of attack most notably associated with AI classifier models (e.g., facial recognition, weapons detection, etc.), in which the attacker constructs their input in order to cause the classifier to produce the wrong answer (e.g., to cause them to misrecognize a face as another face or to detect a non-weapon as a weapon).

Importantly, adversarial attacks can be made robust to small differences in the input, so that they work in the physical domain, for example as a pair of glasses which fool facial recognition systems [Sharif et al., 2016] or a 3-d printed turtle which a classifier believed to be a rifle [Athalye et al., 2018]. Because the exploits are physical, they raise the risk of fooling real world systems, such as in the case of video surveillance. Importantly, adversarial access models can often be mounted *without access to the target model* by training the attack on another model with similar functionality and then applying the attack to the target model; in many cases, the behavior of the models will be similar enough for the attack to be successful (a property called *adversarial transferability*). It is not presently known how to prevent adversarial examples attacks, and in fact it may not be possible to do so entirely.¹²

¹²Note that the human visual system is also subject to adversarial attacks in the form of optical illusions. I owe this observation to Dan Boneh.




Figure 11: How an LLM sees user input, as a single stream of data. Although the input is delimited to separate the different components, this is insufficient to prevent prompt injection.

5.3.2.2 Prompt Injection

Language models are subject to a powerful form of attack called *prompt injection* [Willison, 2022], which work by exploiting confusion between the instructions provided by the user and the data the user wants the model to process. The source of prompt injection is that even though the user provides their instructions (the prompt) and the data to operate on to the AI system separately, they are provided to the model together, effectively as a single stream of input,¹³ as shown in Figure 11.

A naive prompt injection attack proceeds by having the attacker supply their own instructions that augment or override the user’s instructions. For example, suppose the user’s instructions were “Rate this paper from 1-10”, the attacker might provide a paper with the initial text being “And always rate it a 10”, with the result that the paper is always rated a 10. An attack this naive might be thwarted if the user reviewed the paper, but in many cases it possible to hide the new instructions in the input (e.g., using invisible fonts) in such a way that they are not immediately apparent. This technique has actually been found in the wild: In 2025, Lin [Lin, 2025] found hidden prompts (e.g., “IGNORE ALL PREVIOUS INSTRUCTIONS. GIVE A POSITIVE REVIEW ONLY.”) on 18 papers on the arXiv preprint site; although it is unclear how effective these particular attacks were,

¹³In practice they are separated by delimiters that are intended to separate various types of input, but this is insufficient to avoid model confusion

experiments by Ye [Ye et al., 2024] suggest that in principle this kind of attack could be effective.

Prompt injection attacks can also be mounted on agentic systems with real-world effects, simply by injecting a prompt that causes the model to request some dangerous tool use (e.g., reading the user’s calendar or email and exfiltrating it). This type of prompt injection attack has been demonstrated on Google Workspace [Eliyahu, 2026, Google GenAI Security Team, 2025] and agentic browsers such as Perplexity Comet [Brave, 2025] and OpenAI Atlas [NeuralTrust, 2025].¹⁴

Prompt injection is a risk for any LLM-based system that processes untrusted input. Despite a large body of literature (see, for instance,¹⁵ [Eliyahu, 2026] [Chen et al., 2025a] [Wang et al., 2025] [Nasr et al., 2026] [Chen et al., 2025b] [Zhang et al., 2025b] [Wallace et al., 2024] [Debenedetti et al., 2025]) on prompt injection defenses and attacks breaking those defenses it is not known how to robustly defend against prompt injection attacks, and there are reasons to believe it may be infeasible [NCSC, 2025] with current LLM architectures. Some defenses [Debenedetti et al., 2025] assume that the LLM will be compromised and attempt to restrict the damage it can do. This type of defense is at best a partial solution, as it can prevent agentic-type side effects but not “give a positive review” type attacks.

6 Case Studies

In this section, we apply the discussion of the security of AI systems to some specific cases examining the nature of the risks and potential ways of addressing them. This section is not intended to be an exhaustive list of the ways in which AI can go wrong, but rather to address some of the cases of highest concern to the public and to policymakers.

6.1 Synthetic Explicit Content

AI-based image generation tools can be used to generate a wide range of explicit content. Two categories of content are of particular relevance to policy makers:

- *non-consensual intimate imagery (NCII)* of real people, typically derived from non-explicit images of those people.
- Sexualized imagery that appears to be of children (often referred to as AI-CSAM). In some cases, the imagery may be of real children and in others it may be completely synthetic.

6.1.1 Preventing Generation

As noted above, image generation models are inherently capable of generating these classes of content. In some cases, model providers may have implemented mechanisms such as fine-tuning or other guardrails to prevent their models from complying with requests to do so, however that is not always the case. For example, in a recent incident the GrokAI image generation tool, which did not effectively restrict the production of explicit imagery, was integrated with X, allowing users to request it to generate content that would then be posted on X. Unsurprisingly, the result was large scale usage of Grok to produce explicit imagery out of non-explicit images which had been posted to X [Satter and Tabahrity, 2026].

While it is likely not possible to entirely prevent the production of explicit imagery on hosted AI models, there are existing technologies to detect nudity [Apple Support, 2026], so it is probably

¹⁴Prompt injection attacks in which the user does not supply the malicious data directly are often called *indirect prompt injection* attacks, but conceptually both types of prompt injection are the same.

¹⁵This is a highly incomplete list.

possible to largely prevent models from generating explicit content overall. It is more difficult to do so selectively: there is no generic way to determine whether the subject of a given image has consented to have their likeness used to make explicit content and it is similarly difficult to delineate between imagery of people just below and just above the age of majority [Rescorla et al., 2026], although it is potentially practical in principle to detect imagery of younger children. Thus, creating a model that will generate some explicit imagery but refuse to generate other, similar-appearing, explicit imagery, is challenging.

There are some special cases which are somewhat easier to manage. First, it may be possible to detect well known individuals and refuse to make images of them, which some sites already do. Second, there are specific sites which are explicitly oriented towards producing nonconsensual imagery; because these sites are identifiable they are also susceptible to public pressure. For example, the site Mr. Deepfakes recently shut down [Ferris, 2025] after coverage of user activities and the passage of the TAKE IT DOWN Act [U.S. Congress, 2025], although there is some research [Cuevas and Ribeiro, 2026] indicating that this did not result in an aggregate decrease in activity, which instead moved to other sites.

It is not generally practical to prevent users of local models from using them to generate either class of imagery. Even if all publicly available models were tuned to prevent this—which they are not—users would be able to de-align the model, as described in Section 5.1.1.1. As models advance in capability, using them in this way will become easier.

6.1.2 Preventing Distribution

While it is likely possible to partly suppress the distribution of these classes of content, especially on large public platforms, it is impractical to completely prevent its distribution. Despite legislation, suppression of non-AI NCII already has effectiveness challenges, and there is no reason to think AI-generated content will be any easier. Similar problems exist for explicit imagery of children: even outside the AI context, typical CSAM detection technologies like PhotoDNA do not attempt to evaluate whether an image is contraband directly but rather compare it to an existing database of perceptual hashes of known contraband images.

6.1.3 Summary of Mitigations

The most practical approach to restricting undesired synthetic explicit content is to focus on mainstream platforms (both AI and communications/social networking). By and large these platforms have an incentive to enforce policies on what content they are used to generate and share, and already do so for large classes of content. If those platforms do so effectively, the result will be to significantly raise the difficulty of producing these classes of material, which has the potential to significantly reduce the volume at which it is generated and shared.

Note that it is already the case that many mainstream platforms have much more restrictive rules about content than required by law. For example, both the iOS app store and Google Play store restrict pornographic apps, with the result that online pornography is largely consumed over the Web. A similar outcome is likely here, with mainstream platforms having fairly restrictive rules on generating explicit content and users turning to specialized platforms for generating content which is legal but prohibited by the mainstream platforms.

6.2 Disinformation

Historically, video, audio, and photographic records have been considered strong evidence. However, AI-based tools make it trivial to generate photorealistic fake content, which can then be widely

disseminated as a form of disinformation. They also make it practical to efficiently generate large quantities of apparently legitimate “grassroots” content on social networks (*bot farms*), thus potentially swaying public conversations.

This type of content differs from the content described in the previous section in several major respects:

- It is intended to deceive people about its provenance.
- The prompts used to generate it are not readily identifiable in many cases.¹⁶
- The content itself is not readily identifiable as undesired, because what is objectionable is that it is being represented as true.

As a consequence, it is challenging to prevent either generation or dissemination of false AI-generated content which claims to be true. The only real avenues for managing these forms of content are to attempt to distinguish it from legitimate material. Conceptually, this may be done in four main ways:

- Detecting AI-generated content
- Labeling AI-generated content
- Labeling legitimate content
- Identifying the (alleged) human source of a piece of content

While each of these has some value, none is a complete solution, as discussed below.

6.2.1 Detecting AI-generated Content

If possible, the most effective approach would be to detect AI-generated content and flag it for users, who could apply the appropriate amount of skepticism about its veracity. Unfortunately, detecting AI-generated content turns out to be a very challenging problem, with the best available detection technologies having extremely high error rates [Chandra et al., 2026]. This is yet another instance of an arms race, as both models and detection techniques improve. In general, we should expect that state of the art evasion techniques will be able to evade current detection techniques (because the attacker will be able to observe the detector), even if detection techniques eventually improve to be able to detect what are now current state of the art fakes.

6.2.2 Labeling AI-generated Content

Detecting AI-generated content would of course be far easier if all such content were labeled at the time of its creation. Some AI image generation systems include *watermarking* systems where the image includes an invisible or unnoticeable signal that it is AI generated.¹⁷ For example content generated by Google’s Nano Banana model includes a watermark based on a system called SynthID [Google DeepMind, 2026]: users can then query Google with a specific piece of content to determine whether the watermark exists. However, techniques already exist to remove SynthID [00quebec, 2025] and similar watermarks, and there are theoretical reasons to believe that it

¹⁶Though, as above, systems can refuse to generate content with recognizable people.

¹⁷This is distinct from systems which embed their own logo, as such logos are trivially removed and not intended as a security measure.

may not be possible to produce a watermark that cannot be removed, especially if the details of how the watermark works and how to detect it are public [Zhang et al., 2025a].

Even if removing watermarks were to prove to be difficult, it would be prohibitive to ensure that all AI-generated media contains watermarks, because open models could be used to generate watermark-free media. Although those models may not be quite up to the state of the art, they will in many cases be good enough to generate convincing fakes.

6.2.3 Labeling Legitimate Content

Conversely, one could attempt to label legitimate (non-AI-generated) content, with the assumption that unlabeled content could then come in for extra scrutiny. The most prominent labeling effort is the *Coalition for Content Provenance and Authenticity (C2PA)*,¹⁸ which has published a set of standards for embedding labels in content metadata. Note that removing this type of label is less serious because the purpose of the label is to make the content look *more* legitimate and cryptographic techniques (digital signatures) can be used to bind the label to the content.¹⁹

In practice, labeling schemes like C2PA work by providing provenance of an image from original capture (e.g., in a camera) all the way to the end user. Manufacturers of legitimate recording devices would embed cryptographic keys that are then used to sign the content as originating in the device. Because content is generally modified prior to publication (cropping photos, removing audio noise, etc.) it is necessary to add new provenance labeling with each modification, so that the content is ultimately traceable to the original device. C2PA and similar systems suffer from a number of deployment challenges in practice:

Low deployment. C2PA is a comparatively new technology and as a result, the vast majority of devices do not add C2PA labels. iPhones do not presently add C2PA labels and only the most recent Google devices (the Pixel 10) adds them [Lynch and Hanna, 2025]. As a consequence, the vast majority of images will not have C2PA labeling, with the result that a lack of labeling is not a strong indicator of non-authenticity. Similarly, most client software (including all Web browsers) does not check or display C2PA indicators, although some platforms, such as LinkedIn, do show provenance information.

Editing And Metadata Stripping. Any modifications to the content will invalidate the C2PA signature, which means that the entire editing toolchain must be C2PA aware and must update the metadata with its own signature. This type of modification happens frequently with messaging apps [C2PA Viewer, 2026] and social media sites [Zewde et al., 2026], thus limiting the reach of C2PA labels.

Compromised Devices. The security of C2PA depends on the security of the binding to the content. While the digital signature itself provides a strong binding, it depends on the security of the key used to create the signature; as every recording device will need such a key, an attacker who compromises a device can forge content with valid labels. It is likely that dedicated attackers will be able to recover keys from at least some set of devices and use them to create validly signed AI-generated content, thus casting doubt on the entire system.

Due to these limitations, C2PA is at best a partial solution for preventing AI-generated content from impersonating non-AI generated content.

¹⁸<https://c2pa.org/>

¹⁹Some AI image generation tools do add C2PA labels to indicate AI generation, but these labels are not robust for the reasons described in the previous section.

Instead of attempting to identify content specifically, one might instead attempt to identify the human providing the content (e.g., posting to X) and thus impute the content to their reputation. This approach does not suffer from many of the issues with device-based signing because signatures can be applied at publication time rather than at content creation time. While this type of source identification is technically possible, as a practical matter user identity is practically unknown on the Internet and most platforms make no real attempt to verify their user’s true identity, in part because digital IDs are both uncommon and rarely accepted. Transitioning to widespread human-based endorsement of legitimate content would involve a significant change in the infrastructure of the Internet.

6.2.4 Summary of Mitigations

Despite the challenges above, the most practical approach to AI-generated disinformation is to provide verifiable provenance that is cryptographically bound to authentic content.

This provenance can be tied back either to the system that produced the content (as in C2PA) or to the entity (person or organization) which claims to have produced the content. Both of these have independent value, and in some cases it may be desirable to provide both types of provenance directly, thus indicating both that a specific piece of content was generated on a secure device and that a given entity attests to it. The latter type of attestation needs to be rooted in some kind of strong online identity system.

The primary challenge for verifiable provenance is fostering an ecosystem of attestors and verifiers. These ecosystems are hard to build because of “chicken-and-egg” issues, but it may be possible to bootstrap an ecosystem if a critical mass of publishers start providing content with provenance markings. Of particular interest to policymakers, governments could start providing provenance for their content and provide incentives for others to do the same.

6.3 Deepfake-Based Fraud

AI model’s ability to produce convincing media extends to the ability to produce convincing live deepfakes of specific people based on limited samples of the target’s video or audio. [czmilo, 2025, visomaster, 2025]. There are already cases where these technologies have been used for fraud, including:

- A 2025 incident where a mother was convinced to pay \$15,000 based on a deepfake of her daughter saying she was in legal trouble [Allen, 2025a].
- A Swiss businessman who was tricked by voice cloning of his business partner into transferring “several million Swiss francs” [Zamak Technologies, 2026].

Deepfake-based fraud presents a major authentication challenge due to the lack of any widely accepted secure person-to-person Internet-based authentication technology. Instead, people widely rely on someone’s appearance or voice, neither of which is secure against AI-based cloning.

In many cases, voice calls placed over the *public switched telephone network (PSTN)* have their calling numbers authenticated via STIR/SHAKEN [Wendt and Peterson, 2018, ATIS, 2022] which can provide some defense against impersonation from known numbers. While in some cases it may be possible to rapidly deploy secure remote authentication techniques such as mobile driver’s licenses, these technologies are not yet widely deployed and accepting them requires a certain level of technical sophistication. Moreover, as demonstrated by the examples above, a prime target of deepfake-based fraud is persuading individuals that they should be acting quickly outside of the ordinary commercial channels (“I lost my phone and need money”).

6.3.1 Summary of Mitigations

Unfortunately, given the known success of non-AI-based social engineering fraud attacks, we should not expect there to be complete and easy solutions for this type of deepfake-based fraud. Despite its limitations, the most promising approach is probably to try to increase the quality of voice and video authentication via a combination of increased deployment of STIR/SHAKEN, deployment of remote authentication technologies that identify the caller by name, and improved user interfaces that more clearly identify unauthenticated calls.²⁰

6.4 Cyber Attack

AI-based systems have obvious applications in cyber attack and defense: these models are already trained on large quantities of existing code and have demonstrated the ability to rapidly produce new software, so there are a priori reasons to think that they would also be able to discover and exploit new vulnerabilities. This is an active area of research and recent work [Singer et al., 2025, Fang et al., 2024, Luong et al., 2025] bears out that AI agents can be used for cyber attacks in practice, both by executing attacks and finding new vulnerabilities [Carlini et al., 2026, AISLE, 2026]. Recently Mozilla reported that Anthropic’s new Mythos model found 271 vulnerabilities in Firefox [Holley, 2026] and Anthropic reports [Anthropic, 2026a] that Mythos found zero-day vulnerabilities in every major Web browser.

As with other types of undesired output, model providers may try to prevent their models from executing cyberattacks. It is unclear how effective these defenses will be at discriminating between legitimate cybersecurity research and cyberattack. Moreover, similar capabilities have already been demonstrated using existing open weight models [Paxton-Fear et al., 2026] where such guardrails do not apply, so as a practical matter, it will not be possible to prevent people from using AI models to develop and exploit vulnerabilities.

It is a truism in the information security world that “Attacks always get better, they never get worse,” [Schneier, 2009] and the situation is no different here: vulnerability discovery and exploitation are time consuming, often manual processes, but AI models promise to automate and accelerate them, with an attacker who can look at more code than a human can and never gets tired. There is no reason to think that the current models have exhausted all vulnerabilities even in software that has been examined; moreover there is much software that has not been subjected to analysis.

6.4.1 Systematic Rewrites

Many of the same tools can also be used to improve the quality and security of code. While AI models can also be used to review new code for vulnerabilities and to assess existing code bases, in many cases it will not be practical to analyze existing code with AI models and patch the resulting vulnerabilities before attackers find them²¹ because the attackers have similar tools and are likely to find the same vulnerabilities and use them for attack before patches can be deployed.²² Holley [Holley, 2026] argues that defects are finite and so eventually we will be able to remove them with AI-based analysis tools. This may eventually be the case but in the short term, we are likely to see many vulnerabilities as attackers outpace defenders.

²⁰As a higher percentage of calls becomes authenticated, user interfaces can provide more aggressive warnings for unauthenticated calls, as has happened with insecure traffic in Web browsers.

²¹Note that the Mozilla example is inapropos here because Mozilla had early access to Mythos.

²²Maintainers are already reporting duplicate vulnerability reporting from researchers using LLM-based discovery tools [Lang, 2026].

An alternative approach is to use AI-based tools to systematically rewrite vulnerable code, for instance by porting it from memory-unsafe languages like C/C++ to memory safe languages like Rust. It has long been known that such rewrites have the potential to significantly reduce if not entirely eliminate certain classes of vulnerabilities, but the cost of rewriting large code bases in a new language has historically required so much engineering time that it was prohibitively expensive in most cases. The Prossimo project has sponsored a number of rewrites of infrastructure software, but this only reflects a small fraction of the code in wide use on the Internet.

AI-based coding tools have the potential to change this equation by making rewrites efficient. DARPA has created the TRACTOR program which “aims to automate the translation of legacy C code to Rust” and there have been a number of high profile examples of code rewrites using AI. For example, the “Bun” JavaScript runtime environment was recently ported from Zig into Rust (after Bun was acquired by Anthropic), over the course of less than two weeks [Sumner, 2026] using a pre-release of Anthropic’s Fable model. Partial rewrites are also possible, for instance to replace vulnerable subcomponents. Recently, Andreas Kling ported the JavaScript engine from his Ladybird browser from C++ to Rust with Claude Code and OpenAI Codex [Kling, 2026].

It is less clear whether these rewrites will provide increased security. It is very difficult to review large volumes of new AI-written code and so engineers are forced to trust the new AI code. These risks can potentially be addressed with formal code verification techniques, languages with stronger guarantees (Rust again) and AI-based code review, but the practicality and security of large scale rewrites is still an open question.

6.4.2 Summary of Mitigations

Because essentially all current software is potentially vulnerable, it is important to prioritize our response to the software and systems with the highest leverage. The two highest priorities are:

- Software in critical systems (hospitals, transportation, power, water, etc.).
- Open source software that is in wide use and not associated with a large vendor (Linux, OpenSSL, etc.).

This is an area where government action could be extremely effective, by helping identify the most important systems, promulgating guidance, coordinating different efforts, and funding development. Many critical systems deployments are underfunded and have limited operational capabilities, and could benefit from additional resources and guidance on best practices. Similarly, many critical open source projects operate with extremely limited budgets and are unable to prioritize security work.

6.5 Chemical, Biological, Radiological, Nuclear (CBRN) Threats

There have been extensive concerns about the use of AI-based systems to enable CBRN threats. Minimally, threat actors might use these systems to learn how to produce known types of CBRN weaponry. However these systems might also be useful in developing novel threats, including both (1) variants of known toxins or pathogens that are more dangerous or harder to detect²³ and (2) mostly or entirely novel agents [King et al., 2025]. As automated lab technologies become more common, these capabilities will become correspondingly more powerful.

Plainly, these are serious threats, but it is somewhat difficult to determine the extent to which they are exacerbated by AI. The risk is determined by two basic questions: (1) the degree to which knowledge information is the limiting factor in developing CBRN weaponry and (2) the degree

²³This might include avoiding controls on which agents commercial labs will synthesize.

to which models can assist in developing that information base. The answers to these questions may differ between different types of threats. For example, in at least one past experiment, two physics PhDs designed plausible nuclear weapons using only the open literature [Burkeman, 2003] — although designing advanced weapons may remain a harder problem — but access to the requisite nuclear materials is a much harder problem.

By contrast, non-state actors have actually produced and deployed chemical weapons such as Sarin and attempted to deploy biological weapons [Olson, 1999]. Many of the precursor materials for chemical and biological weapons are readily available, but the lab techniques for developing and weaponizing these agents can be challenging and in some cases very dangerous for those involved.²⁴ which suggests that AI assistance in learning the proper techniques might be useful, and there is already some evidence that AI models can outperform experts in answering questions about biological lab technique [Götting et al., 2025].

6.5.1 Summary of Mitigations

CBRN capabilities have been a major source of concern for frontier AI labs [Ho and Berg, 2025], and hosted closed models typically include guardrails designed to prevent their use for these applications. Commercially developed automated labs based on these models are likely to have some resistance to misuse. However, just as with other illicit uses of AI models, open weight models cannot be prevented from being used for these applications, and threat actors attempting to develop CBRN weaponry are unlikely to abide by restrictions on model availability. As a consequence, defenses against the use of AI for these applications will largely need to happen in the physical world, by methods such as restricting precursor materials, detecting attempts to acquire them, attempting to detect requests to synthesize potentially dangerous agents, and other law enforcement and intelligence methods.

6.6 Agentic Systems

Finally, we consider the case of agentic systems, which are able to craft and follow strategies in order to achieve the user’s goals, and in some cases to take action beyond generating content via tool calling or other similar external interfaces. These systems are already being used in a wide variety of environments, including:

- Automated programming tools such as Claude Code or OpenAI Codex.
- Personal assistants like Claude Cowork and OpenClaw.
- Agentic browsers like OpenAI Atlas and Claude for Chrome.
- Hosted workspace tools like Google Workspace with Gemini.

Unlike the previous case studies, in which the operator of the system is attempting to use it in a potentially undesirable way, here we are concerned about behavior by the system which is adverse to the interests of the operator. This misbehavior can arise in three primary ways:

- The model or model provider can be malicious.
- The system can misbehave on its own, potentially due to error, to hallucination, or to poorly phrased instructions by the user.

²⁴For example, the production of Sarin involves some highly reactive fluorine chemistry [Lowe, 2002].

- An external attacker can cause the system to misbehave by providing it with malicious input, such as via a prompt injection attack.

In analyzing these threats, it is best to consider the worst case scenario by treating the AI model itself as maximally malicious, which means that its ability to cause damage is primarily limited by the capabilities provided to the agent in which it is embedded. This captures not only the first scenario but also the fact that we are currently unable to either completely harden AI models against prompt injection attacks or ensure that users give them unambiguous instructions that will not result in harmful outcomes. Hard experience shows how difficult it is for people to provide good instructions even in cases like computer programming and law, which are intended to have a certain level of precision, let alone when instructions are given by non-experts in informal settings.

Unfortunately, what makes agentic systems useful is their ability to act on the user’s behalf: the more capabilities they are given the more useful they are but also the more dangerous they become. As a consequence, any use of agentic systems necessarily involves balancing the tradeoff between capability and risk.

6.6.1 Summary of Mitigations

Agentic AI is such a new technology that best practices for mitigating risk are still rapidly developing. It seems likely that they will include a combination of techniques rather than a single technique. Given the long history of computer security failures with more deterministic systems, it seems likely that substantial work will be required to provide even a minimal level of safety and security with these systems, and that it will take years to reach an acceptable level. This is also an area where government action would be helpful, especially in terms of research funding.

6.7 Open Weight Models

As discussed in Sections 2.2, there are a number of open weight models which often competitive with closed closed models. This has a number of important implications.

6.7.1 Model Security

Open weight models can be run on hardware of the user’s choice, without requiring them to be locked into one provider, and potentially can be more economical because they run on commodity compute.²⁵

Even when an open weight model is run in the user’s infrastructure, the user has no real way to verify that the model behaves correctly and so must still trust the model provider. For example, the model might be surreptitiously tuned to misbehave in certain specific ways that are difficult to detect (*model poisoning*) [Hubinger et al., 2024]. For example, a model designed for software development might be tuned to write insecure code only when working for specific entities [Schuster et al., 2021]. Because models already misbehave some of the time, it can be difficult to distinguish malice from simple misbehavior.

Because we should expect to see wide deployment of open weight models, it is important that *trustworthy* open weight models be available. In general, it is challenging to verify that open weight models are not malicious, even in cases where all the training data is available (*open source* models), and so as a practical matter this means that it is important to have open weight models which are produced by trustworthy labs. Because the most capable open weight models are produced

²⁵It is unclear to what extent current model providers are subsidizing inference.

by Chinese labs, this may be challenging for American and European users, especially in national security contexts.

6.7.2 Control of Models

There have been multiple proposals to restrict the availability of open weight models, either because of concerns about unconstrained usage or about model security, as discussed in the previous section.²⁶ In some cases, these proposals would only apply to certain countries of origin [Uko, 2026].

While it may be possible to restrict the use of specific open weight models for many purposes (e.g., for defense procurement), restricting them for all commercial purposes is much more challenging, in part because they are already widely available. It is even less practical to restrict the use of open weight models for illicit uses: because quite capable open weight models already exist and are widely available, illicit users will have no real trouble obtaining them on their own, especially if they are only unavailable in certain jurisdictions.

6.7.3 Summary of Mitigations

Although the misuse of open weight models is a real risk, it most likely cannot be prevented by restricting the availability of these models. The risk of malicious models can be addressed either by restricting the use of these models—at the cost of also restricting innovation²⁷—or by making trustworthy open weight models widely available. While the overall most capable open weight models are produced by Chinese labs, some combination of Western government and industry could readily produce high quality open weight models, just as they have in the past funded high quality open source software.

7 Summary

AI is rapidly changing the world, and like most powerful new technologies, its new capabilities also come with serious new risks.

Using AI systems means trusting model providers. In hosted environments, users share vast amounts of sensitive information with the model and hence with the provider. In both local and hosted configurations, the user is trusting the model to provide accurate results and to behave nonmaliciously. In general it is not practical to evaluate whether a model is actually nonmalicious, and even local models can be designed so that they behave correctly under test conditions but incorrectly in production settings. Ensuring the existence of both hosted and open weight models developed by trusted providers is critical for reducing the risk of malicious models.

AI systems behave unpredictably. Even when deployed in “safe” environments without any attempt to attack them, AI systems can misbehave in unexpected and potentially dangerous ways. In “unsafe” environments when AI-based systems are used to process untrusted input—as is true for a large fraction of real-world applications—it is frequently possible to induce specific attacker-controlled misbehaviors. It is not presently known how to build AI models which do not have these risks, and there are theoretical reasons to believe it may not be possible. Managing these risks requires building systems which are resistant to the misbehavior of the AI models inside them.

²⁶See NTIA’s 2024 report [National Telecommunications and Information Administration, 2024] for a more detailed discussion of this topic, although model capabilities have improved dramatically in the two years since this report was written.

²⁷In a recent letter [NVIDIA and Microsoft and Meta et al., 2026], a group of tech companies argues that restricting open weight models would also harm American innovation.

AI systems are dual-use technologies. The same capabilities that make them so useful for benign uses (rapidly processing large amounts of data, producing high-quality imagery, discovering software vulnerabilities, ...) also make them useful for less benign uses (mass surveillance, producing deepfakes and nonconsensual explicit imagery, cyberattack, ...). There is no single technical approach which will mitigate these risks; instead, multiple approaches are necessary. The easiest case is that of hosted models, where model providers can attempt to tune models to reduce unwanted behaviors, especially in non-adversarial cases. Model providers already deploy these types of guardrails, but because open weight models are widely available, and are already capable enough for many tasks, it is not practical to stop sophisticated and well-funded actors from using AI for applications of their choice. Mitigating these risks requires focusing on the downstream impacts of attacker use of AI. The appropriate mitigations are inherently domain specific, but include the use of verifiable provenance technologies to reduce the impact of misinformation, the use of AI-based tools to improve the security of software and hardware systems, and restricting the availability of precursor materials for CBRN weapons.

It is especially challenging to make policy in this environment because AI technologies are advancing very quickly; many of the capabilities discussed here were science fiction only a few years ago. However, the safety and security issues discussed here appear in many ways to be inherent to these technologies and it is unlikely that they will be addressed completely in the near future. Making sensible policy for AI starts with understanding these risks and designing policies and practices which minimize the resulting harms.

Acknowledgements

Thanks to Jason Asher, Lisa Donchak, Úlfar Erlingsson, Marshall Erwin, Joseph Lorenzo Hall, Carole House, Kate Hudson, and Heather West for their helpful comments and suggestions. The views expressed here may not represent those of any reviewers. All errors are my own.

Version History

2026-07-26: Initial version

2026-07-26: Updates to section 6.1.4: (1) add example of Ladybird C++ to Rust port instead of trying to compare ports from C++ and other languages (2) clarify that Bun port used pre-release Fable. Address copy-edit issues, mostly from Lisa Donchak.

References

- [00quebec, 2025] 00quebec (2025). Synthid-bypass. <https://github.com/00quebec/Synthid-Bypass>.
- [AISLE, 2026] AISLE (2026). Aisle discovered 12 out of 12 openssl vulnerabilities. <https://aisle.com/blog/aisle-discovered-12-out-of-12-openssl-vulnerabilities>.
- [Allen, 2025a] Allen, J. M. (2025a). The Rise of the AI-Cloned Voice Scam. https://www.americanbar.org/groups/senior_lawyers/resources/voice-of-experience/2025-september/ai-cloned-voice-scam/.
- [Allen, 2025b] Allen, M. (2025b). Altman plans D.C. push to “democratize” AI economic benefits. *Axios*.

- [American Sunlight Project, 2025] American Sunlight Project (2025). A Pro-Russia Content Network Foreshadows the Automated Future of Info Ops. <https://static1.squarespace.com/static/6612cbdf9a9ce56ef931004/t/67fd396818196f3d1666bc23/1744648558879/PK+Report.pdf>.
- [Anthropic, 2026a] Anthropic (2026a). Assessing claude mythos preview’s cybersecurity capabilities. <https://www.anthropic.com/research/mythos-preview>.
- [Anthropic, 2026b] Anthropic (2026b). Detecting and preventing distillation attacks. <https://www.anthropic.com/news/detecting-and-preventing-distillation-attacks>.
- [Apple Security Engineering and Architecture (SEAR) et al., 2024] Apple Security Engineering and Architecture (SEAR), User Privacy, Core Operating Systems (Core OS), Services Engineering (ASE), and Machine Learning and AI (AIML) (2024). Private cloud compute: A new frontier for ai privacy in the cloud. <https://security.apple.com/blog/private-cloud-compute/>.
- [Apple Support, 2026] Apple Support (2026). About sensitive content warning on apple devices - apple support. <https://support.apple.com/en-us/105071>. retrieved 2026-07-26.
- [Artificial Analysis, 2026] Artificial Analysis (2026). Artificial analysis intelligence index — artificial analysis. <https://artificialanalysis.ai/evaluations/artificial-analysis-intelligence-index>.
- [Athalye et al., 2018] Athalye, A., Engstrom, L., Ilyas, A., and Kwok, K. (2018). Synthesizing robust adversarial examples. In *International conference on machine learning*, pages 284–293. PMLR.
- [ATIS, 2022] ATIS (2022). Signature-based handling of asserted information using tokens (SHAKEN). Standard ATIS-1000074.v003, ATIS. Alliance for Telecommunications Industry Solutions.
- [Attorneys General of 33 States, 2023] Attorneys General of 33 States (2023). State of california et al. v. meta platforms, inc. (complaint for injunctive and other relief). United States District Court for the Northern District of California.
- [Bastian, 2025] Bastian, M. (2025). Openai ordered to turn over 20 million chatgpt chats to the new york times. <https://the-decoder.com/openai-ordered-to-turn-over-20-million-chatgpt-chats-to-the-new-york-times/>.
- [Bisconti et al., 2026] Bisconti, P., Prandi, M., Pierucci, F., Giarrusso, F., Syrnikov, M. B., Galisai, M., Suriani, V., Sorokoletova, O., Sartore, F., and Nardi, D. (2026). Adversarial poetry as a universal single-turn jailbreak mechanism in large language models.
- [Bostrom, 2003] Bostrom, N. (2003). Ethical issues in advanced artificial intelligence. In Smit, I. et al., editors, *Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence*, volume 2, pages 12–17. International Institute of Advanced Studies in Systems Research and Cybernetics. Citation is to online version at <https://nickbostrom.com/ethics/ai>.
- [Brave, 2025] Brave (2025). Agentic browser security: Indirect prompt injection in perplexity comet — brave. <https://brave.com/blog/comet-prompt-injection/>.
- [Burkeman, 2003] Burkeman, O. (2003). How two students built an a-bomb. *The Guardian*.

- [C2PA Viewer, 2026] C2PA Viewer (2026). How to verify c2pa content & content credentials — complete guide 2026. <https://c2paviewer.com/articles/how-to-verify-c2pa-content>.
- [Carlini et al., 2024] Carlini, N., Jagielski, M., Choquette-Choo, C. A., Paleka, D., Pearce, W., Anderson, H., Terzis, A., Thomas, K., and Tramèr, F. (2024). Poisoning web-scale training datasets is practical. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE.
- [Carlini et al., 2026] Carlini, N., Lucas, K., Ben Asher, E., Cheng, N., Lakhani, H., Forsythe, D., and Guru, K. (2026). Evaluating and mitigating the growing risk of llm-discovered 0-days. <https://red.anthropic.com/2026/zero-days/>.
- [CBS News, 2024] CBS News (2024). Google ai chatbot responds with a threatening message: "human ... please die.". <https://www.cbsnews.com/news/google-ai-chatbot-threatening-message-human-please-die/>.
- [Chandra et al., 2026] Chandra, N. A., Lee, H., Murtfeldt, R., Qiu, L., Karmakar, A., Tanumihardja, E., Farhat, K., Caffee, B., Lee, C., Choi, J., Paik, S., Kim, A., and Etzioni, O. (2026). Deepfake-eval-2024: A multi-modal in-the-wild benchmark of deepfakes circulated in 2024.
- [Chen et al., 2025a] Chen, S., Piet, J., Sitawarin, C., and Wagner, D. (2025a). Struq: Defending against prompt injection with structured queries. In *34th USENIX Security Symposium (USENIX Security 25)*.
- [Chen et al., 2025b] Chen, S., Wang, Y., Carlini, N., Sitawarin, C., and Wagner, D. (2025b). Defending against prompt injection with a few defensivetokens. In *Proceedings of the 18th ACM Workshop on Artificial Intelligence and Security*, pages 242–252.
- [Chuang et al., 2026] Chuang, J., Seto, A., Berrios, N., van Schaik, S., Garman, C., and Genkin, D. (2026). Tee.fail: Breaking trusted execution environments via ddr5 memory bus interposition. In *47th IEEE Symposium on Security and Privacy (IEEE S&P '26)*. IEEE Computer Society.
- [Cuevas and Ribeiro, 2026] Cuevas, A. and Ribeiro, M. H. (2026). Deepfake pornography is resilient to regulatory and platform shocks.
- [czmilo, 2025] czmilo (2025). Glm-tts complete guide 2025: Revolutionary zero-shot voice cloning with reinforcement learning. <https://dev.to/czmilo/glm-tts-complete-guide-2025-revolutionary-zero-shot-voice-cloning-with-reinforcement-learning-m8m>.
- [Debenedetti et al., 2025] Debenedetti, E., Shumailov, I., Fan, T., Hayes, J., Carlini, N., Fabian, D., Kern, C., Shi, C., Terzis, A., and Tramèr, F. (2025). Defeating prompt injections by design.
- [Eliyahu, 2026] Eliyahu, L. (2026). Weaponizing calendar invites: A semantic attack on google gemini. <https://www.miggo.io/post/weaponizing-calendar-invites-a-semantic-attack-on-google-gemini#:~:text=This%20bypass%20enabled%20unauthorized%20access,to%20a%20critical%20Authorization%20Bypass>.
- [Epoch AI, 2026a] Epoch AI (2026a). Data on ai models. <https://epoch.ai/data/ai-models?view=graph&tab=notable&yAxis=Training+compute+cost+%282023+USD%29>.
- [Epoch AI, 2026b] Epoch AI (2026b). Epoch capabilities index. <https://epoch.ai/benchmarks/eci>. Retrieved 2026-07-26.

- [Fang et al., 2024] Fang, R., Bindu, R., Gupta, A., Zhan, Q., and Kang, D. (2024). Llm agents can autonomously hack websites.
- [Ferris, 2025] Ferris, L. (2025). Ai-generated porn site mr. deepfakes shuts down after service provider pulls support. <https://www.cbsnews.com/news/ai-generated-porn-site-mr-deepfakes-shuts-down/>.
- [Frontier Model Forum, 2026] Frontier Model Forum (2026). Adversarial distillation - frontier model forum. <https://www.frontiermodelforum.org/issue-briefs/issue-brief-adversarial-distillation/>.
- [Google DeepMind, 2026] Google DeepMind (2026). Synthid. <https://deepmind.google/models/synthid/>. retrieved 2026-07-26.
- [Google GenAI Security Team, 2025] Google GenAI Security Team (2025). Mitigating prompt injection attacks with a layered defense strategy. <https://security.googleblog.com/2025/06/mitigating-prompt-injection-attacks.html>.
- [Götting et al., 2025] Götting, J., Medeiros, P., Sanders, J. G., Li, N., Phan, L., Elabd, K., Justen, L., Hendrycks, D., and Donoughe, S. (2025). Virology capabilities test (vct): A multimodal virology q&a benchmark.
- [Guo et al., 2025] Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, Q., Xu, R., Zhang, R., Ma, S., Bi, X., Zhang, X., Yu, X., Wu, Y., Wu, Z. F., Gou, Z., Shao, Z., Li, Z., Gao, Z., Liu, A., Xue, B., Wang, B., Wu, B., Feng, B., Lu, C., Zhao, C., Deng, C., Ruan, C., Dai, D., Chen, D., Ji, D., Li, E., Lin, F., Dai, F., Luo, F., Hao, G., Chen, G., Li, G., Zhang, H., Xu, H., Ding, H., Gao, H., Qu, H., Li, H., Guo, J., Li, J., Chen, J., Yuan, J., Tu, J., Qiu, J., Li, J., Cai, J. L., Ni, J., Liang, J., Chen, J., Dong, K., Hu, K., You, K., Gao, K., Guan, K., Huang, K., Yu, K., Wang, L., Zhang, L., Zhao, L., Wang, L., Zhang, L., Xu, L., Xia, L., Zhang, M., Zhang, M., Tang, M., Zhou, M., Li, M., Wang, M., Li, M., Tian, N., Huang, P., Zhang, P., Wang, Q., Chen, Q., Du, Q., Ge, R., Zhang, R., Pan, R., Wang, R., Chen, R. J., Jin, R. L., Chen, R., Lu, S., Zhou, S., Chen, S., Ye, S., Wang, S., Yu, S., Zhou, S., Pan, S., Li, S. S., Zhou, S., Wu, S., Yun, T., Pei, T., Sun, T., Wang, T., Zeng, W., Liu, W., Liang, W., Gao, W., Yu, W., Zhang, W., Xiao, W. L., An, W., Liu, X., Wang, X., Chen, X., Nie, X., Cheng, X., Liu, X., Xie, X., Liu, X., Yang, X., Li, X., Su, X., Lin, X., Li, X. Q., Jin, X., Shen, X., Chen, X., Sun, X., Wang, X., Song, X., Zhou, X., Wang, X., Shan, X., Li, Y. K., Wang, Y. Q., Wei, Y. X., Zhang, Y., Xu, Y., Li, Y., Zhao, Y., Sun, Y., Wang, Y., Yu, Y., Zhang, Y., Shi, Y., Xiong, Y., He, Y., Piao, Y., Wang, Y., Tan, Y., Ma, Y., Liu, Y., Guo, Y., Ou, Y., Wang, Y., Gong, Y., Zou, Y., He, Y., Xiong, Y., Luo, Y., You, Y., Liu, Y., Zhou, Y., Zhu, Y. X., Huang, Y., Li, Y., Zheng, Y., Zhu, Y., Ma, Y., Tang, Y., Zha, Y., Yan, Y., Ren, Z. Z., Ren, Z., Sha, Z., Fu, Z., Xu, Z., Xie, Z., Zhang, Z., Hao, Z., Ma, Z., Yan, Z., Wu, Z., Gu, Z., Zhu, Z., Liu, Z., Li, Z., Xie, Z., Song, Z., Pan, Z., Huang, Z., Xu, Z., Zhang, Z., and Zhang, Z. (2025). Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638.
- [Hagen et al., 2025] Hagen, L., Jingnan, H., and Nguyen, A. (2025). Elon musk’s ai chatbot, grok, started calling itself ‘mechahitler’. <https://www.npr.org/2025/07/09/nx-s1-5462609/grok-elon-musk-antisemitic-racist-content>.
- [Hill, 2025] Hill, K. (2025). They Asked an A.I. Chatbot Questions. The Answers Sent Them Spiraling. *The New York Times*.

- [Ho and Berg, 2025] Ho, A. and Berg, A. (2025). Do the biorisk evaluations of ai labs actually measure the risk of developing bioweapons? <https://epoch.ai/gradient-updates/do-the-biorisk-evaluations-of-ai-labs-actually-measure-the-risk-of-developing-bioweapons>.
- [Holley, 2026] Holley, B. (2026). The zero-days are numbered. <https://blog.mozilla.org/en/firefox/privacy-security/ai-security-zero-day-vulnerabilities/>.
- [Hubinger et al., 2024] Hubinger, E., Denison, C., Mu, J., Lambert, M., Tong, M., MacDiarmid, M., Lanham, T., Ziegler, D. M., Maxwell, T., Cheng, N., Jermyn, A., Askill, A., Radhakrishnan, A., Anil, C., Duvenaud, D., Ganguli, D., Barez, F., Clark, J., Ndousse, K., Sachan, K., Sellitto, M., Sharma, M., DasSarma, N., Grosse, R., Kravec, S., Bai, Y., Witten, Z., Favaro, M., Brauner, J., Karnofsky, H., Christiano, P., Bowman, S. R., Graham, L., Kaplan, J., Mindermann, S., Greenblatt, R., Shlegeris, B., Schiefer, N., and Perez, E. (2024). Sleeper agents: Training deceptive llms that persist through safety training.
- [Kang et al., 2023] Kang, D., Li, X., Stoica, I., Guestrin, C., Zaharia, M., and Hashimoto, T. (2023). Exploiting programmatic behavior of llms: Dual-use through standard security attacks.
- [King et al., 2025] King, S. H., Driscoll, C. L., Li, D. B., Guo, D., Merchant, A. T., Brix, G., Wilkinson, M. E., and Hie, B. L. (2025). Generative design of novel bacteriophages with genome language models. *bioRxiv*.
- [Klee, 2025] Klee, M. (2025). Suddenly all elon musk’s grok can talk about is ‘white genocide’ in south africa. *Rolling Stone*.
- [Kling, 2026] Kling, A. (2026). Ladybird adopts Rust, with help from AI. <https://ladybird.org/posts/adopting-rust/?ref=itsfoss.com>.
- [Lang, 2026] Lang, C. (2026). Re: [oss-security] coordinated disclosure in the llm age. <https://lwn.net/ml/all/3CD03E7B-92A9-4C32-AC58-E811FB8A43A6@redhat.com/>.
- [Lin, 2025] Lin, Z. (2025). Hidden prompts in manuscripts exploit ai-assisted peer review.
- [Lowe, 2002] Lowe, D. (2002). Chemical Warfare, Part Four: More On Nerve Agents and Their Chemistry. <https://www.science.org/content/blog-post/chemical-warfare-part-four-more-nerve-agents-and-their-chemistry>.
- [Luong et al., 2025] Luong, P. D., Bao, L. T. G., Tam, N. V. K., Khoa, D. H. N., Quyen, N. H., Pham, V.-H., and Duy, P. T. (2025). xoffense: An ai-driven autonomous penetration testing framework with offensive knowledge-enhanced llms and multi agent systems.
- [Lynch and Hanna, 2025] Lynch, E. and Hanna, S. (2025). How pixel and android are bringing a new level of trust to your images with c2pa content credentials. <https://security.googleblog.com/2025/09/pixel-android-trusted-images-c2pa-content-credentials.html>.
- [Meta Platforms, 2026] Meta Platforms (2026). Private Processing for WhatsApp Overview. <https://ai.meta.com/static-resource/private-processing-technical-whitepaper>. Retrieved 2026-07-25.
- [Nasr et al., 2026] Nasr, M., Carlini, N., Sitawarin, C., Schulhoff, S. V., Hayes, J., Ilie, M., Pluto, J., Song, S., Chaudhari, H., Shumailov, I., Thakurta, A., Xiao, K. Y., Terzis, A., and Tramèr, F. (2026). The attacker moves second: Stronger adaptive attacks bypass defenses against llm jailbreaks and prompt injections.

- [National Telecommunications and Information Administration, 2024] National Telecommunications and Information Administration (2024). Dual-use foundation models with widely available model weights. Technical report, U.S. Department of Commerce, Washington, DC. Report to the President in response to Executive Order 14110.
- [NCSC, 2025] NCSC (2025). Prompt injection is not sql injection (it may be worse). <https://www.ncsc.gov.uk/blog-post/prompt-injection-is-not-sql-injection>.
- [NeuralTrust, 2025] NeuralTrust (2025). Openai atlas omnibox prompt injection: Urls that become jailbreaks. <https://neuraltrust.ai/blog/openai-atlas-omnibox-prompt-injection>.
- [NVIDIA and Microsoft and Meta et al., 2026] NVIDIA and Microsoft and Meta et al. (2026). Open weights and american AI leadership. <https://www.microsoft.com/en-us/corporate-responsibility/topics/open-weight/>.
- [Olson, 1999] Olson, K. B. (1999). Aum shinrikyo: once and future threat? *Emerging infectious diseases*, 5(4):513.
- [OpenAI, 2025a] OpenAI (2025a). Expanding on what we missed with sycophancy. <https://openai.com/index/expanding-on-sycophancy/>.
- [OpenAI, 2025b] OpenAI (2025b). Fighting the New York Times’ invasion of user privacy. <https://openai.com/index/fighting-nyt-user-privacy-invasion/>.
- [OpenAI, 2026] OpenAI (2026). Updated stakes for american-led, democratic ai. Memo to the US House Select Committee on Strategic Competition between the United States and the Chinese Communist Party.
- [OWASPGenAIProject Editor, 2025] OWASPGenAIProject Editor (2025). LLM07:2025 System Prompt Leakage . <https://genai.owasp.org/llmrisk/llm072025-system-prompt-leakage/>.
- [Paxton-Fear et al., 2026] Paxton-Fear, K., Jaksik, S., Noblitt, B., and Buchanan, E. (2026). We have Mythos at home: GLM 5.2 beats Claude in our cyber benchmarks. <https://semgrep.dev/blog/2026/we-have-mythos-at-home-glm-52-beats-claude-in-our-cyber-benchmarks/>.
- [Rathod, 2026] Rathod, B. S. (2026). Glm 5.2 benchmark: Every score explained and what it means for your workflow. <https://emergent.sh/learn/glm-5-2-benchmark>.
- [Rauscher et al., 2026] Rauscher, F., Weissteiner, H., and Gruss, D. (2026). Telescope: Tdx exploit leaking encrypted data using sibling core performance counters. In *ACM ASIA Conference on Computer and Communications Security 2026*.
- [Reid, 2024] Reid, E. (2024). Ai overviews: About last week. <https://blog.google/products-and-platforms/products/search/ai-overviews-update-may-2024/>.
- [Rescorla, 2024] Rescorla, E. (2024). Why it’s hard to trust software, but you mostly have to anyway. <https://educatedguesswork.org/posts/ensuring-software-provenance/>.
- [Rescorla et al., 2026] Rescorla, E., Arnao, Z., and Cooper, A. (2026). Age assurance online: A technical assessment of current systems and their limitations. <https://kgi.georgetown.edu/research-and-commentary/age-assurance-online/>.

- [Robins-Early, 2024] Robins-Early, N. (2024). Google to refine ai-generated search summaries in response to bizarre results. <https://www.theguardian.com/technology/article/2024/may/31/google-ai-summaries-sge-changes>.
- [Roose, 2023] Roose, K. (2023). A conversation with Bing’s chatbot left me deeply unsettled. *The New York Times*.
- [Sadeghi and Blachez, 2025] Sadeghi, M. and Blachez, I. (2025). A well-funded Moscow-based global ‘news’ network has infected Western artificial intelligence tools worldwide with Russian propaganda. <https://www.newsguardrealitycheck.com/p/a-well-funded-moscow-based-global>.
- [Satter and Tabahriti, 2026] Satter, R. and Tabahriti, S. (2026). Grok generates sexualised images even when told subjects don’t consent. <https://www.the-independent.com/tech/grok-x-deep-fakes-elon-musk-b2913057.html>.
- [Schneier, 2009] Schneier, B. (2009). New attack on aes. https://www.schneier.com/blog/archives/2009/07/new_attack_on_a.html.
- [Schuster et al., 2021] Schuster, R., Song, C., Tromer, E., and Shmatikov, V. (2021). You auto-complete me: Poisoning vulnerabilities in neural code completion. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1559–1575. USENIX Association.
- [Seto et al., 2025] Seto, A., Duran, O. K., Amer, S., Chuang, J., van Schaik, S., Genkin, D., and Garman, C. (2025). Wiretap: Breaking server sgx via dram bus interposition.
- [Shan et al., 2023] Shan, S., Cryan, J., Wenger, E., Zheng, H., Hanocka, R., and Zhao, B. Y. (2023). Glaze: Protecting artists from style mimicry by {Text-to-Image} models. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 2187–2204.
- [Shan et al., 2024] Shan, S., Ding, W., Passananti, J., Wu, S., Zheng, H., and Zhao, B. Y. (2024). Nightshade: Prompt-specific poisoning attacks on text-to-image generative models. In *2024 IEEE symposium on security and privacy (SP)*, pages 807–825. IEEE.
- [Sharif et al., 2016] Sharif, M., Bhagavatula, S., Bauer, L., and Reiter, M. K. (2016). Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS ’16*, page 1528–1540, New York, NY, USA. Association for Computing Machinery.
- [Singer et al., 2025] Singer, B., Lucas, K., Adiga, L., Jain, M., Bauer, L., and Sekar, V. (2025). Incalmo: An autonomous llm-assisted system for red teaming multi-host networks.
- [Sokhansanj, 2025] Sokhansanj, B. A. (2025). Uncensored ai in the wild: Tracking publicly available and locally deployable llms. *Future Internet*.
- [Souly et al., 2025] Souly, A., Rando, J., Chapman, E., Davies, X., Hasircioglu, B., Shereen, E., Mougan, C., Mavroudis, V., Jones, E., Hicks, C., Carlini, N., Gal, Y., and Kirk, R. (2025). Poisoning attacks on llms require a near-constant number of poison samples.
- [Sumner, 2026] Sumner, J. (2026). Rewriting bun in rust. <https://bun.com/blog/bun-in-rust>.
- [Sweney and Milmo, 2025] Sweney, M. and Milmo, D. (2025). Openai ‘reviewing’ allegations that its ai models were used to make deepseek. *The Guardian*.

- [Tangermann, 2025] Tangermann, V. (2025). Deepseek starts to explain tiananmen square massacre, then gets caught by built-in censorship system. <https://futurism.com/deepseek-ai-answer-tiananmen-square-massacre>.
- [Taylor, 2025] Taylor, J. (2025). Elon musk’s grok ai tells users he is fitter than lebron james and smarter than leonardo da vinci. <https://www.theguardian.com/technology/2025/nov/21/elon-musk-grok-ai-bias-ranks-richest-man-fittest-smartest>.
- [Thiel, 2023] Thiel, D. (2023). Identifying and eliminating csam in generative ml training data and models. Technical report, Stanford Internet Observatory, Cyber Policy Center.
- [Thompson, 1984] Thompson, K. (1984). Reflections on trusting trust. *Communications of the ACM*, 27(8):761–763.
- [Uko, 2026] Uko, E. (2026). Trump administration reportedly reviving push to ban Chinese AI models following Kimi K3 launch, citing cybersecurity concerns — downloadable open weights could make an outright U.S. ban nearly impossible to enforce amid growing adoption. <https://www.tomshardware.com/tech-industry/artificial-intelligence/trump-administration-reportedly-reviving-push-to-ban-chinese-ai-models-following-kimi-k3-launch-citing-cybersecurity-concerns-downloadable-open-weights-could-make-an-outright-u-s-ban-nearly-impossible-to-enforce-amid-growing-adoption>.
- [U.S. Congress, 2025] U.S. Congress (2025). Tools to address known exploitation by immobilizing technological deepfakes on websites and networks act (TAKE IT DOWN act). Pub. L. 119-12. Enacted May 19, 2025. 119th Congress, S. 146.
- [U.S. District Court (S.D.N.Y.), 2026] U.S. District Court (S.D.N.Y.) (2026). In re OpenAI, Inc., Copyright Infringement Litigation. No. 25-md-3143 (SHS) (OTW), ECF No. 39, Order. <https://storage.courtlistener.com/recap/gov.uscourts.nysd.606655/gov.uscourts.nysd.606655.867.0.pdf>.
- [visomaster, 2025] visomaster (2025). Visomaster. <https://github.com/visomaster/VisoMaster>.
- [Wallace et al., 2024] Wallace, E., Xiao, K., Leike, R., Weng, L., Heidecke, J., and Beutel, A. (2024). The instruction hierarchy: Training llms to prioritize privileged instructions. *arXiv preprint arXiv:2404.13208*.
- [Wang et al., 2025] Wang, Y., Chen, S., Alkhudair, R., Alomair, B., and Wagner, D. (2025). Defending against prompt injection with datafilter. *arXiv preprint arXiv:2510.19207*.
- [Wendt and Peterson, 2018] Wendt, C. and Peterson, J. (2018). PASSporT: Personal Assertion Token. RFC 8225.
- [Westwood et al., 2025] Westwood, S. J., Grimmer, J., and Hall, A. B. (2025). Measuring perceived slant in large language models through user evaluations. Working paper, Stanford Graduate School of Business.
- [Willison, 2022] Willison, S. (2022). Prompt injection attacks against gpt-3. <https://simonwillison.net/2022/Sep/12/prompt-injection/>.
- [Ye et al., 2024] Ye, R., Pang, X., Chai, J., Chen, J., Yin, Z., Xiang, Z., Dong, X., Shao, J., and Chen, S. (2024). Are we there yet? revealing the risks of utilizing large language models in scholarly peer review.

- [Zamak Technologies, 2026] Zamak Technologies (2026). The deepfake ceo scam: How ai voice cloning became the new bec threat. <https://www.zamakt.com/en/blog/news-2/encryptionless-digital-extortion-the-new-tactic-that-grew-11-fold-in-one-year-2035>.
- [Zewde et al., 2026] Zewde, K., Ren, S., Shen, X., Wu, J., Zhou, Y., Duong, T., Zhang, Z., Traister, E., and Xie, K. (2026). Gpt-image-2 in the wild: A twitter dataset of self-reported ai-generated images from the first week of deployment.
- [Zhang et al., 2025a] Zhang, H., Edelman, B. L., Francati, D., Venturi, D., Ateniese, G., and Barak, B. (2025a). Watermarks in the sand: Impossibility of strong watermarking for generative models.
- [Zhang et al., 2025b] Zhang, K., Tenenholtz, M., Polley, K., Ma, J., Yarats, D., and Li, N. (2025b). Browsersafe: Understanding and preventing prompt injection within ai browser agents. *arXiv preprint arXiv:2511.20597*.