Safety Co-Option and Compromised National Security: The Self-Fulfilling Prophecy of Weakened AI Risk Thresholds

HEIDY KHLAAF, AI Now Institute

SARAH MYERS WEST, AI Now Institute

Risk thresholds provide a measure of the level of risk exposure that a society or individual is willing to withstand, ultimately shaping how we determine the safety of technological systems. Against the backdrop of the Cold War, the first risk analyses, such as those devised for nuclear systems, cemented societally accepted risk thresholds against which safety-critical and defense systems are now evaluated. But today, the appropriate risk tolerances for AI systems have yet to be agreed on by global governing efforts, despite the need for democratic deliberation regarding the acceptable levels of harm to human life. Absent such AI risk thresholds, AI technologists—primarily industry labs, as well as "AI safety"-focused organizations—have instead advocated for risk tolerances skewed by a purported AI arms race and speculative "existential" risks, taking over the arbitration of risk determinations with life-or-death consequences, subverting democratic processes.

In this paper, we demonstrate how such approaches have allowed AI technologists to engage in "safety revisionism," substituting traditional safety methods and terminology with ill-defined alternatives that vie for the accelerated adoption of military AI uses at the cost of lowered safety and security thresholds. We explore how the current trajectory for AI risk determination and evaluation for foundation model use within national security is poised for a race to the bottom, to the detriment of the US's national security interests. Safety-critical and defense systems must comply with assurance frameworks that are aligned with established risk thresholds, and foundation models are no exception. As such, development of evaluation frameworks for AI-based military systems must preserve the safety and security of US critical and defense infrastructure, and remain in alignment with international humanitarian law.

1 INTRODUCTION

At the height of the Cold War and on the heels of the Manhattan Project, US nuclear researchers devised the first risk-analysis frameworks to evaluate the risk and safety mitigations of developing and deploying powerful technological systems, while ensuring the US maintained an edge amidst a nuclear arms race with the Soviet Union. These established risk analyses and thresholds provided the precedent against which defense and safety-critical systems continue to be evaluated today [71, 87, 88]. Yet parallels of a purported AI-arms race with China are nevertheless being used to justify the jettisoning of these time-tested frameworks, at the cost of both safety and security, in exchange for the unfettered development and deployment of AI technologies [3, 29, 60].

Establishing the accepted risk and dependability of AI systems has surfaced several challenges in the determination of how and if AI technologies can and need to be trusted in the context of safety-critical systems and accelerated military deployment. These challenges have often stemmed from a surfeit of hype, a lack of clear definitions, and seeming inability to clearly specify the precise intended application of AI systems. Global governance efforts, like those at the United Nations (UN), currently lack specificity, as the appropriate societal risk tolerances for AI systems have yet to be agreed on by civil society, legislatures, or other rulemaking and enforcement bodies. The absence of such concrete deliberation has given way for AI technologists to seize the determinations of said risk tolerances, and thus thresholds, going against long-established norms.

2 • Khlaaf and Myers West

Indeed, despite invoking comparisons that AI risk is on par with safety risks that were once posed with the coming of the nuclear age [50], and that have arguably resulted in the resilience of US military capabilities and critical infrastructure, **AI technologists are instead using the pretense of an AI arms race and speculative "existential" risks to discard the very risk and safety thresholds established by this era**. The purported, and even dispelled, cost-benefit justification for this discrepancy is that the accelerated and ubiquitous adoption of AI above all else provides a marker of the US's technological advantage and defense prowess over China and other adversaries [3, 29, 36, 60]. Not only are these presumptions unsubstantiated given the historical inaccuracies and questionable efficacy of AI-based systems [14, 26, 53, 94], but they also fully concede the determinations of risk tolerances and thresholds to AI technologists, reflecting an adoption of risk scales that are fundamentally at odds with democratic deliberative norms. Put bluntly, the militaristic push for AI adoption puts life-or-death decisions in the hands of technologists instead of governing and legislative bodies, often in ways that are opaque and unaccountable, while defying previous precedents established for novel technological advancements. These matters extend beyond influencing how AI is used on the battlefield: Military safety-engineering standards[1] can set precedents for the evaluation of civilian safety-critical systems (and vice versa) [54], thus shaping broader standards that define the infrastructures that members of the public interact with every day.

Ironically, this approach—while articulated as necessary for national security—not only shifts institutional risk framing to bolster the embedding of technological systems with autocratic values within critical US infrastructure, but also, **by allowing for accelerated AI adoption at the cost of lowered safety and security thresholds, may be precisely what disadvantages US military and technological capabilities against China**. As such, a slew of existing national-security concerns emanating from uses of AI models will be ripe for adversaries to exploit and compromise US national and defense infrastructure. If integrated into critical infrastructures, AI's persistent lack of performance may lead to strategic, operational, and tactical mistakes that could compromise the US's international standing, and have deadly and geopolitically consequential impacts. Indeed, a plethora of critiques have explored the destabilizing impact of the AI arms race on geopolitics, international relations, and AI regulation [8, 21, 35, 43, 81, 90, 97]. However, such critiques have not yet addressed how shifted risk thresholds diminish the role safety plays in the assurance of critical systems by weakening (purposely or inadvertently) long-established safety-engineering principles that would deem many AI systems unsafe and insecure for military use. As such, discarding rudimentary safety-engineering measures to evaluate AI-enabled military systems will imperil not only the safety and security of military personnel, but of the civilian populace as well.

In this paper, we examine the history of safety engineering and its parallels with today's AI arms race, highlighting the appropriation and co-option of safety engineering and verification and validation (V&V) terminology within the field of "AI safety" to justify unverifiable and unsubstantiated claims regarding AI capabilities—claims that substitute technical and regulatory terms such as "safety," "safety cases," and "red-teaming" with substantially revised definitions that hollow out assurance methodologies, ultimately compromising safety-critical and defense systems. As an alternative, we offer a grounded exploration of AI safety that reconnects this term with its disciplinary origins. This work seeks to help remedy an alarming gap in the field by bringing military AI systems to the forefront of the AI safety and security discussion. In particular, if evaluation of foundation models serve as the basis of prospective assessments with which emerging AI safety-critical and defense systems are to comply

---

[1]These are standards adopted by the military that define appropriate practices for guaranteeing system safety. For example, DOD MIL-STD-882 outlines the US DOD's approach to "eliminating hazards, where possible, and minimizing risks where those hazards cannot be eliminated," and defines thresholds for acceptable risk [67].

(i.e., lethal autonomous weapons systems[LAWs] or AI decision support systems[DSS]), then as they stand they must be aligned with and assured against risk thresholds that continue to maintain the safety and security of US critical and defense infrastructure, while remaining in line with international humanitarian law (IHL).

2 HISTORICAL DETERMINATIONS OF RISK THRESHOLDS AND AI

The safety[2] of technological systems is a well-established discipline spanning several fields that is often revisited with the advent of a novel technology, necessitating the adaptation of existing methodologies for newly introduced hazards. Establishing the safety of AI systems, and in particular foundation models, has reinvigorated a focus on safety methodologies such as risk assessments and safety cases in pursuit of the assurance of AI systems given a surge of interest in deploying them widely, including within national infrastructure [69]. Yet despite numerous efforts to adopt safety practices to address the unique hazards introduced by AI systems [4, 10, 28, 34, 51, 52, 79], these works have failed to address a key gap in the AI safety discussion: establishing the risk thresholds of AI systems that safety assurances are to be determined against. That is, adequately addressing **the question of "How safe is safe enough?" first requires establishing a societally accepted risk-benefit trade-off or tolerance that provides safety goals and thresholds against which technological systems can be assessed, a discussion largely absent from current AI governing efforts.**

In this section, we surface historical lessons on risk tolerances and thresholds established for technologies such as nuclear energy that provided a precedent against which modern-day safety-critical systems continue to be evaluated. We then examine how, in failing to determine safety goals that follow tried-and-true approaches, AI global governance efforts have allowed AI promoters to instead supplant risk tolerances with an AI-arms-race logic that ultimately allows for accelerated AI adoption at the cost of lowered safety thresholds. In the ensuing sections, we demonstrate how the absence of societally accepted risk tolerance for AI has also allowed AI technologists to wield safety and V&V terminology toward unverifiable or speculative risk thresholds and claims that intend to signal compliance with established safety and defense thresholds, whilst only displaying a veneer of safety.

To effectively demonstrate the gap in which AI safety frameworks fail to adequately address "How safe is safe enough?" and the safety goals needed to establish the acceptable risk to the public, we first revisit the foundational work of Chauncey Starr, a nuclear and electrical engineer and one of the pioneers of probabilistic risk analysis. Through his work on the Manhattan Project and subsequent efforts with Atomics International to pursue the commercialization of the generation of electricity by nuclear power, Starr formalized the first risk analysis frameworks to evaluate the risk and safety mitigations of developing and deploying powerful nuclear technological systems amidst a nuclear arms race with the Soviet Union [24, 32, 86, 92]. Working under geopolitical conditions that mirror our own, Starr was especially privy to the ongoing nuclear arms race at the time, and even believed that President Truman was correct in his decision to use atomic weapons against Japan; Starr himself had helped develop the "Little Boy" bomb [76]. **Despite this, Starr's formulation of risk thresholds and analysis significantly deviates from the arms race calculus being touted today by AI technologists.**

---

[2]We use the term "safety" in the traditional systems-engineering sense, meaning the prevention of a system from impacting its environment in an undesirable or hazardous way, typically aiming to protect human lives, the environment, or even military assets. This is distinct from the meaning of the term as used in AI, which we will address in more detail in the following sections.

Specifically, Starr defined risk analysis to be the study of the relevant cause-and-effect relationships that give rise to safety risks, to describe their magnitudes and distributions, and to identify mitigations for risk reduction, which concerns the character of risk and the social significance of risk identified [88]. He posited that risk analysis would require the consideration of the societal evaluation of risk and, as such, interpreting public attitudes and values. This characterization was intended to emphasize that risk analysis concerns at once questions that can be addressed scientifically and questions involving public policy concerns (i.e., societal safety)[3]—where **risk tolerance** characterizes the degree, amount, or volume of risk that a society or individual will withstand, while **risk threshold** measures the level of risk exposure above which action must be taken to address risks proactively, and below which risks may be accepted. Starr noted that risk acceptance, and thus thresholds, are inseparable from risk perception and evaluation.

Starr observed two primary, and often conflicting, societal value systems of risk that reflect different types of social costs to risks: **societal risk** and **individual risk**. In modern-day literature, societal risk has been formalized to concern the probability of an event with many fatalities; individual risk has been formalized to concern the annual probability of death of a single person. Starr provided a more intuitive and example-driven distinction, where an individual value system "charges an increasing social cost to each successive disability day." Under such a system, the social cost of one individual being disabled or harmed for a year is considered greater than the cost of 365 separate individuals each disabled or harmed for one day.

With regard to selecting either societal or individual value systems of risk, a determination of risk tolerance and thresholds is intended to be driven by societal decision-makers that bridge the gap between the perceived and assessed risks and benefits of a novel technological system. Starr pointedly noted that autocratic powers and military organizations overindex societal risk thresholds in the aim of preserving the welfare of these regimes rather than an individual risk value system, whereas democratic governments use a mix of both scales to determine thresholds of safety. These dynamics extend to the role of technologists themselves in the establishment of risk thresholds: Starr noted that in a democratic society, a technologist's role ends with providing the quantification of risk impacts, as has often been reflected by US advisory mechanisms that have allowed science advisers to shape, with limitations, US federal policy with key technologies such as nuclear [40, 56]. Otherwise, Starr characterized the placement of technocratic experts in decision-making positions that extend their role beyond risk assessors as indicators of an autocratic society [88].

Indeed, the historical determinations of safety thresholds for safety-critical systems within the US have largely been established through broad democratic deliberations involving elected representatives, regulators, and expert consensus-building that draws on both societal and individual value risks. This has included risk management to implement policies consistent with values as perceived by democratic institutions, and managing public risk perceptions through mitigations, including for nuclear systems even against the backdrop of the Cold War [41].[4]

---

[3]We note that modern-day uses of the term "risk assessment" may not address societal decisions carried out on systems for which risk thresholds have already been determined.

[4]An example of such threshold includes Starr's recommendation for nuclear power where "the upper bound be set at a risk level equivalent to those risks of routine living which are normally accepted, i.e., about 10^{-4} deaths per year per person (100 deaths/yr/million). The proposed lower design target is 10^{-8} (0.1 deaths/yr/million), about one-hundredth of the minimal risk from the natural hazards all people are exposed to" [86].

We find ourselves in dramatically different circumstances with today's AI development and deployment efforts. Rather than through democratic deliberation, AI technologists are exclusively seeking to establish the risk tolerances—and thus thresholds—by which AI systems are evaluated, utilizing skewed risk and safety interpretations at odds with the demands of social, political, and ethical civilian life [84], and with assurance practices established by experts like Starr that are critical to the safety and security posture of national and defense infrastructure. Under the pretense of an AI arms race and "existential" risks, "AI safety" proponents are working to accelerate the adoption of AI within defense and critical infrastructure at the cost of safeguarding implementation—justifying likely harm to both defense personnel and civilians by the need to presumably establish the US's technological advantage and defense prowess over China [3, 29, 36, 60]. Yet the parallels between nuclear proliferation and AI development have been found to be misrepresented and largely unfounded, not only because of the lack of definable success of AI within military applications that have relied on perception rather than fit-for-purpose capabilities, but also because AI is not itself a weapon; it is, rather, a decision support or automation function [26, 81]. Unsubstantiated claims aside, the prioritization of militarism in current AI policymaking in the US indicates a shift toward a societal value for risk determination rather than one of individual value, reflecting an adoption of risk scales associated with autocratic rather than democratic societies that would require democratic deliberation on the impacts of AI and acceptable thresholds for risks.

In such an autocratic society, Starr noted, the social costs of harms or accidents are determined (whether accurately or not) to be directly proportional to the estimated expected values of direct harm costs, and as such "a social value system relating to risk characteristics is not applied, and the decision process is rather automatic" [88]. We see this reflected in the recent flurry of AI-driven systems being deployed to purportedly automate bureaucracy and critical and defense governing decision-making [2, 16, 58, 69, 98]. Such processes implicitly position risk assessors (i.e., AI technologists) as decision-makers by eliminating the channels through which public attitudes and values can bridge the gap between the perceived and assessed risks and benefits of novel technological systems, laundering decision-making and risk determination to the very AI systems built by technologists. For example, lofty claims about potential AI use are being deployed to displace grounded expertise, such as in the case of the US Department of Government Efficiency (DOGE), where the promise that AI can administer large parts of the federal government has functioned to undermine and discredit the expertise of career civil servants and to present this as a technical transition, rather than as a political transfer of power [22, 48, 83]. This is further bolstered by the lobbying power of AI industry players who influence decision-makers within the US in ways that go beyond the limitations typically placed on US science advisers in shaping federal policy [12, 13]. **Starr's frameworks thus dispel claims that the acceleration of AI use within the national security apparatus would allow "democracies to maintain the lead": The means themselves are reflective of a transformation away from democratic use of technology and toward the very autocratic standards that these technologists warn against in a purported AI-arms race with China.**

Finally, the value judgments required to convert risk decision factors to a common AI risk threshold metric have been muddled by a surfeit of hype, a lack of clear definitions, and a concomitant absence of clearly specified claims about the precise intended application of AI systems. In neglecting to consider that safety frameworks can only be effective when operating within the parameters of a societally accepted risk value system, research and global governance efforts alike have allowed diffuse conversations regarding "AI safety" to substitute the concrete and precise language used in safety engineering with interpretations antithetical to their original meaning.

We note that various risk-perception literature, including Starr's, does not directly compare how technological system risks compare to human errors, as the safety-critical systems and processes under scrutiny (e.g., financial, health, nuclear risks) are not ones that human labor could substitute, and thus have their risk be compared against (as in the case for some sociotechnical systems). However, recent AI developments have been explicitly concerned with the automation of human labor (e.g., foundation models), leading some to demand that standardized risk analyses be overhauled with measures only leveraging average human performance as the acceptable risk threshold. Yet safety-critical risks concern the lack of harm or failures arising in an attempt to meet a system intent, and not the intent itself. **Given that AI failure modes are distinctly and unpredictably different from that of humans', their risk categories fall firmly within the scope of traditional risk frameworks aimed at safety-critical technological systems, regardless of whether or not an AI system is intended to automate human capabilities.** This issue has been explored in other capacities, including within the autonomous-driving literature, notably regarding flawed and misleading human performance metrics used to justify insufficient AI reliability [27, 44, 89]. Given the safety-critical context of AI use within defense, we thus exclude human performance as an appropriate or feasible risk metric for use in our analysis, and revisit the impacts of such metrics on benchmarks in later sections.

In the next section, we demonstrate how these skewed risk tolerances have allowed for the co-option of safety methodologies through permitting the use of unverifiable claims, or claims that are not indicative of the accepted safety thresholds that have governed other technologies. We also explore how, in neglecting to apply safety thresholds that have been established since the nuclear arms race, global governance efforts have allowed for implicit acceptance of arbitrary risk tolerances and thresholds, as determined by those with self-interest in the development and deployment of AI-based systems. As a result, this safety revisionism will not only undermine and fail to preserve national security and its necessary compliance with IHL, but will imperil civilian safety-critical systems to the detriment of the public.

3   SAFETY REVISIONISM AND IMPLICATIONS FOR NATIONAL SECURITY

Within the AI discourse, the term "safety" has come to have a multitude of definitions that vary according to the context and the community. Yet these definitions seldom overlap with the broader meaning of "safety" used within safety-critical fields, including in defense, and may in fact directly contradict it. "Safety" in AI has become conflated with "alignment" or "capabilities" [5] that aim to steer AI systems toward human-oriented intents and goals and prevent a "model's capabilities from proliferating broadly" [4, 6, 17, 23, 28, 78]. These metrics are not only subjective given their reliance on hypothetical scenarios untethered to empirical assessments, but fundamentally conflate safety properties with system intent [51]. Safety must center the protection of human lives, the environment, and critical assets. This is achieved by ensuring that the system is prevented from impacting its environment in an undesirable or hazardous way, including due to the system intent itself. Safety is also context-sensitive and there are no standard assurance or assessment approaches for "general" safety (as posited by AI literature), since doing so would render the evaluation of risk and safety intractable [10, 51, 59]. Ultimately, the objective of safety is "freedom from risk which is not tolerable"[ISO/IEC Guide 51:1999, definition 3.1]; or, in other words, residual risk is reduced as far as reasonably tolerable or practicable.[6]

---

[5]The use of the term "capabilities" in AI is not to be confused with its definition in the systems engineering context, which is the ability to execute a specified course of action given a set of requirements.
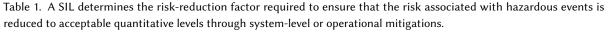
[6]This is a principle known as "as low as reasonably practicable" or ALARP.

The discrepancy is not an issue of pedantry. Alignment and capabilities measures are simply not concerned with the system hazards arising out of specified intent, but rather with a system's intent itself, either through the speculative ascription of autonomy to foundational models and their ability to inflict harm, or through concerns of how deployed models may be "broadly" misused. Concerns of speculative capabilities aside (given their lack of verifiability and thus applicability to safety frameworks), considerations of how novel technologies may be misused are not novel, whether they regard nuclear proliferation or use of software to commit cyber crimes [33, 41, 47, 63, 72]. **Although adapting systematic safety thinking to capability evaluations can provide constructive outcomes [52, 59], these are not context- or safety-specific hazards or failure modes as traditionally defined under safety-engineering frameworks. As such, their conflation with safety risks impacts the efficacy and relevance of security and safety engineering techniques necessary for the assurance of AI within mission-critical systems (e.g., defense)** Yet the supplanting of traditional safety assessments by broad "capabilities evaluation" frameworks is precisely what AI technologists and labs are proposing for evaluations in their pursuit to deploy foundation models within safety-critical and defense applications for use in DSS, intelligence, surveillance, target acquisition, and reconnaissance (ISTAR), or even within counter-unmanned aircraft systems [5, 53, 65, 74, 77, 91, 93]. Indeed, the absence of societally accepted AI risk tolerances has allowed AI technologists to leverage both the altered and unfounded definition of safety, jointly with skewed AI-arms tolerances, to bolster the deployment of frameworks that deviate from well-established military standards necessary for the functionality of defense systems.

To demonstrate how this conflation undermines safety, consider that the traditional definition of safety was first conceived to ensure the performance of safety-critical infrastructure and defense systems [59], where dependability properties (e.g., functionality, performance, reliability, operability, robustness, availability, and even security-informed-safety [7, 46]) have since become fundamental requirements for safety-critical applications. Despite this, recent rhetoric that "the AI future is not going to be won by hand-wringing about safety. It will be won by building, from reliable power plants to the manufacturing facilities that can produce the chips of the future" [16, 37], effectively demonstrate how the AI "alignment" and speculative "capabilities" narrative has distanced the term "safety" from its context, despite its historical centrality and necessity in assuring critical and defense systems. **This enables safety to be rendered as "hand-wringing," working to the detriment of methodologies that critically ensure the functionality, reliability, operability, robustness, and availability properties of safety critical systems**—the same "reliable power plants" Vance calls for—and precisely demonstrating the harms in revisionism of safety and security terminology.

Deviations from traditional safety terms not only break with established safety thresholds and corresponding safety and dependability properties, but render safety techniques ineffective, as has also been demonstrated with the use of "safety cases" as redefined by AI technologists. Traditionally, a safety case is a structured argument supported by evidence intended to justify that a system is acceptably safe for a specific application in a defined operating environment. A safety case is often used as a regulatory mechanism for the basis of approvals [49]. Safety cases, and safety analyses more generally, are intended to model causal dependencies between faults and failures, and functional capabilities, in the case of safety of the intended functionality [59]. In defense, the use of a safety case enables the understanding of the cumulative and interrelated risks associated with the use or deployment of a system [66]. Yet AI labs have redefined this methodology with concepts such as "safety cases with red-team validation" and "affirmative cases." The former is defined as "a robustness target set based on a safety

Safety Integrity Level	Safety	Probability of Failure on Demand	Risk Reduction Factor
SIL 4	>99.99%	0.001% - 0.01%	100,000 - 10,000
SIL 3	99.9% - 99.99%	0.01% - 0.1%	10,000 - 1,000
SIL 2	99% - 99.9%	0.1% - 1%	1,000 - 100
SIL 1	90% - 99%	1% - 10%	100 - 10

Table 1. A SIL determines the risk-reduction factor required to ensure that the risk associated with hazardous events is reduced to acceptable quantitative levels through system-level or operational mitigations.

case considering factors like the critical capability and deployment context," while the latter is intended to support claims regarding "models not autonomously attempt[ing] to strategically undermine our safety measures or cause large-scale catastrophe" [6, 28]. Both definitions are at odds with traditional safety cases and risk thresholds.

For one, a safety case is not intended to produce any "targets" or thresholds, as defined by Google DeepMind's "Frontier Safety Framework" [28], but rather is the mechanism to substantiate that a system meets a defined claim, one that is typically derived from a risk threshold. As previously noted, risk thresholds are derived from risk tolerances that characterize the degree, amount, or volume of risk that a society or individual will withstand for a specific application. As such, these risk thresholds are established through risk analyses as initially formalized by Starr [87, 88] and not through safety cases. In Table 1, we provide an example of risk thresholds expressed in Safety Integrity Levels (SIL), a formal measure of expected system safety performance. These revised definitions thus blur the distinction between the processes through which risk thresholds are defined (i.e., risk analyses) and the safety analyses through which risk thresholds are substantiated (e.g., safety cases). This lack of distinction undermines not only the soundness of these methodologies, but also the independence between the democratic deliberation of risk tolerances and the technologists that provide the quantification of risk impacts, as noted by Starr [88].

DeepMind's framework also casts safety cases as a mitigation measure, when safety cases are in fact the primary activity that enables the exploration of how a system impacts its environment in an undesirable or hazardous way (e.g., hazard analysis), and for the demonstration of how a system implementation and hazard mitigations enable a system to meet an acceptably safe and predetermined risk threshold for a specific application. Other works have likely misinterpreted criticism of confirmation bias in safety cases with the inability to consider hazards and harms [6, 23], yet this criticism itself has long been addressed through the use of "defeater" claims [11]. Indeed, safety cases encompass all the necessary activities to substantiate or counter safety claims, and, as such, substitution of well-established safety activities within a case (e.g., through Hazard and Operability Studies, Fault Tree Analysis, Failure Modes and Effects Analysis) with ill-defined methodologies such as "red-teaming" [42, 51], will only undermine a safety case's efficacy in determining that a critical system is acceptably safe.

Finally, proposed "frontier safety frameworks" have supplanted traditional safety claims with assertions concerning "capabilities risks" [6, 28, 39] that fail to justify their safety targets and involve deliberative or consensus-building approaches to establish an accepted balance between societal and individual value risk scales.[7] Indeed, a

---

[7]Recent works have addressed similar shortcomings and have attempted to refine the expressiveness of "general" safety arguments to be more conducive to verifiable argumentation [34]. However, this work similarly lacks concrete risk thresholds that foundation models can be assured against, and problematically cedes the acceptability of said risk tolerances to "AI advisors."

safety case must demonstrate that the safety targets are well-founded and that the safety functions implemented within a system meet said targets with consideration of the anticipated hazards associated with operating a system, neither of which are demonstrated within these proposed frameworks. Examples of such claims include AI models "not autonomously attempt[ing] to strategically undermine our safety measures" or "not caus[ing] large-scale catastrophe" [6, 23, 28], without further clarification on the quantitative thresholds they present [45] nor specificity on what "catastrophe," "undermine," or "autonomously" mean. These claims not only utilize the problematically revised definition of "safety," but lack the formalization and specificity of risk-tolerance quantification that would be necessary to produce verifiable risk-reduction arguments (see Table 1) and the system-level evidence to substantiate them beyond hypothetical scenarios. Although safety case notations such as Claims, Arguments, Evidence (CAE) can be utilized toward non-safety goals and as a tool for general reasoning (i.e., assurance cases), the use of the term "safety case" in the context of the assurance of AI for defense and national security signals an alignment with established military standards and terminology that proposed "frontier safety frameworks" simply do not satisfy. Indeed, safety in defense and safety-critical systems would necessitate context, and as such **the safety of AI systems cannot be effectively evaluated in the abstract, and assessments must address either specific use cases and contexts or deployment given a range of determined risk thresholds.**

But absent AI risk thresholds established through democratic deliberation, and military initiatives that place AI technologists in positions to form corresponding AI evaluation frameworks [3, 6, 38, 55, 91], AI technologists have sought to substitute traditional safety frameworks with the aforementioned ill-defined "capabilities" or "alignment" counterparts that deviate from well-established military standards necessary for the functionality of defense systems and ultimately cede the determination of risk tolerance to AI technologists. Put bluntly, this places decision-making regarding the acceptable risk tolerance for AI, including number of deaths, in the hands of AI technologists themselves (a signifier of an autocratic risk value system). Yet they have insufficiently shown the responsibility and rigor necessary for this task in the context of national security and where risk tolerance needs to be accountable to IHL.

3.1 Impacts on National Security

This safety revisionism has significant consequences for the defense and national security apparatus. Not only do revised definitions of "safety" and "safety cases" deviate from well-established military standards necessary for the functionality of defense systems, but the speculative nature of "AI safety" claims often hinges on unfounded assumptions regarding the capabilities and "intent" of AI systems that may be leveraged to justify the deployment of risky systems. For one, the misuse of terminology intended to signal compliance with established safety and security defense practices can mislead duty holders[8] regarding the claims that an AI system complies, and can lead to approval of AI systems that display only a veneer of safety. Furthermore, unsubstantiated and even dispelled [21, 26, 81, 90] claims of capabilities have already been used to advocate for a skewed cost-benefit calculus, where the accelerated adoption of AI above all else, including the dependability of these systems, has been touted as a marker for the US's technological advantage and defense prowess over China and other adversaries [3, 29, 36, 60, 74]. Weakened terminology and arms-race justification come at a clear cost, as identified risks and failure modes continue to be unveiled [18–20, 64, 96, 100, 101] despite purported claims that "potential" AI systems will eventually eliminate these risks.

---

[8]A defense term for those responsible for the management of risk.

Rhetoric that has attempted to prioritize security threats over safety hazards will only escalate national-security failures, as security and safety are closely related and their interactions introduce complex consequences of failure and interdependencies [7]. Examples of such risks include safety hazards emanating from the inability to prevent personally identifiable information from contributing to the use and proliferation of defense AI capabilities such as ISTAR by adversaries [53]. They also include interconnected security hazards, where it has been demonstrated that the usage of foundation models within military settings inherently expands the attack vectors military systems are susceptible to, including the defense infrastructures they interface with, particularly given new AI supply-chain vulnerabilities. These supply-chain threats can lead to novel attack surfaces that allow unauthorized access and changes in data, model parameters, and intellectual property given a lack of separation between commercial models and those deployed within military applications [18–20, 53, 64, 100, 101]. These vulnerabilities pose defense and national security risks due to the significant increase of attack surfaces that AI-based systems present, which can inadvertently or maliciously accelerate these systems toward a state of imprecise and unlawful targeting and a loss of human control, impacting both the safety and security of defense infrastructure and also creating conditions that may harm defense personnel and give rise to potential violations to IHL [70].

If appropriate military standards and safety procedures had been applied in pilots that have sought to integrate foundation AI models in defense use, then these risks would be easily unveiled. This has not been the approach taken by the US Department Of Defense (DOD) Task Force Lima, which was established to create an evaluative framework for the deployment of foundation models in defense contexts. Rather than create a framework stemming from traditional safety and defense practices, Task Force Lima took on a "capabilities evaluation" approach that was then concluded through the creation of an Artificial Intelligence Rapid Capabilities Cell (AI RCC) designed for "rapid user-centric experimentation"—though details on the methodologies that will be used for system testing are slim. Similar approaches have been taken by other governments, such as the UK's "AI Security Institute" [38, 39], despite important assertions of caution by leaders in the defense community; for example, the Navy's chief information officer, Jane Overslaugh Rathbun, noted that commercial models have "inherent security vulnerabilities" and are "not recommended for operational use cases" [80]. In the next section, we further address the chasm between "capabilities evaluations" of foundation AI models and their inability to meaningfully demonstrate their safety, security, and fitness for purpose by examining inconsequential evaluation criteria that have been scoped too narrowly or have little to no relevance for defense applications, and ultimately break with defense V&V norms.

4   THE FACADE OF CAPABILITIES VIA MANUFACTURED METRICS

The principles of evaluation for defense and safety-critical applications hinge on the well-established formalization of Test & Evaluation, Validation & Verification (TEVV), where verification ensures that the system meets safety specifications as designed, while validation confirms that the system fulfills its intended purpose in its operational environment without causing unacceptable risk [68]. That is, TEVV processes are intended to evaluate the functional requirements of a system, including its reliability, operability, robustness, and availability, while also ensuring that a system meets its safety requirements intended to avert risk. However, evaluation and benchmarking criteria for foundation AI models have significantly deviated from these tried-and-true approaches in favor of ill-defined methodologies such as "red-teaming" and general evaluations [55, 65] that are scoped too narrowly or have little to no relevance for safety-critical applications [42, 51], whilst continuing to be propped up

by a flawed arms-race calculus. Such an approach will ultimately compromise strategic, operational, and tactical defense missions given their proposed use to support system-specific claims that they cannot substantiate.

Once again, the absence of enforceable risk thresholds has led to an implicit acceptance of arbitrary risk tolerances and assessments as determined by those with self-interest in the development and deployment of AI-based systems, without independent scrutiny. The combination of corporate secrecy shielding AI developers with military classification shielding deployed use exacerbates this problem. The trend toward the use of "general" benchmarking for defense applications, wherein a set of arbitrary tasks is intended to indicate the generalizable capabilities and risks of an AI model, demonstrates a profound misunderstanding of the foundations of TEVV, specifically that safety-critical and defense systems cannot be substantiated without application-specific requirements [10, 51, 59, 67, 85].

Consider type certification in aviation, where evaluations are carried out for the approval of a particular vehicle design under specific airworthiness requirements (e.g., Federal Aviation Administration 14 CFR part 21). There is no standard assessment approach for "generic" vehicle types across all domains given varying requirements and accepted risk threshold depending on the deployment environment of a vehicle. It would thus be contrary to established practices—and unproductive—to presume that evaluating foundation models against puzzle-solving tasks or a subset of generic tasks would be generalizable to the performance of systems designed for mission-critical tasks such as command and control (C2), weapons development, intelligence gathering, target acquisition, and reconnaissance. For the evaluation of hazards, "red-teaming" in AI has come to refer to probing foundational models for hazardous or harmful outputs, which are then used to update the model safeguards latently [31, 73, 75]. Note that "red-teaming" is not defined according to its traditional security context.[9] The equivalent technique under TEVV that was likely intended for use is boundary- or stress-testing, a verification technique that aims to test edge cases or fringe inputs that may lead to unknown failure modes and potential hazards. However, just as with the revisionism of "safety" and "safety cases," AI "red-teaming" fails to consider the context against which potential harms and hazards are to be measured, and as such cannot substantiate safety and security requirements intended to avert risk as would be required by TEVV. Furthermore, the safeguards that result from red-teaming exercises have been consistently demonstrated to be insufficient and brittle, bringing into question their applicability in safety-critical environments [42, 100, 101].

The correct selection and use of TEVV methodologies is paramount to the assessment of novel military technologies in assuring their fitness for purpose to prevent casualties whilst maximizing the destruction of intended adversaries within the bounds of IHL. Yet novel military technologies alone, from the "electronic battlefields" in Vietnam to drone warfare in the Middle East, have never been able to replace a realistic assessment of what military force can achieve [15]. Today, similar challenges arise with proposed AI-based defense technologies that not only remain untested, but are highly prone to failures amid purported and contradictory claims that AI-driven systems minimize civilian casualties [1, 9, 15, 30]. As previously noted, AI failure modes are distinctly and unpredictably different from human failure modes and, as such, the assessment for their fitness would not only require affirmation of their functionality (i.e., validation), but also assurance that AI models meet safety requirements intended to avert risk that may arise due to their failures (i.e., verification). Indeed, AI systems pose an intractable number of risks relative to previous technological counterparts given their nondeterministic

---

[9]In cybersecurity, the intent of a red-teaming exercise is to realistically test an organization's capability to detect and respond to a staged adversarial attack, and to assess and validate security posture and attack resilience.

nature, and the trend to train AI models for general tasks (i.e., foundation models) rather than military-exclusive functionality only expands the failure modes that must be accommodated for further [51, 53].

Yet given that AI proponents believe that the accelerated and ubiquitous adoption of AI above all else would demonstrate the US's technological advantage and defense prowess over China, they position generalized benchmarks as a manufactured metric that serves to illustrate broad capabilities as an expression of such an advantage, a weak substitute for robust evaluation of what these systems can do when put into use. These sentiments are further bolstered by drawing on parallels to the nuclear arms race against the Soviet Union, and applying the same logic to US adversaries by fueling concerns that China may outpace the US's AI military capabilities through its investments in AI [99]. Yet the parallels between the nuclear arms race and AI have largely been found to be misrepresented, not only due to the lack of definable success of AI within military applications that have relied on promises and perception rather than capabilities, but because AI is not a weapon itself; it is, instead, a support or automation function. Indeed, unlike AI models, physical military systems like nuclear weapons can produce tangible illustrations of their capabilities. On the other hand, AI capabilities claims are difficult to substantiate not only due to their software-based nature [26], but to a history of opaque and misleading claims that conceal the role of human involvement, making it difficult to ascertain and validate actual AI capabilities of state actors [82].

Further challenges arise given that deep neural networks (DNNs), which form the basis of current foundation models, cannot solve tasks outside of their data distribution and training data sets. Their inability to handle novel scenarios that would arise from the fog of war thus raises serious questions about whether they can be successful in military settings. Appropriate evaluation metrics for AI's fitness of purpose must consider how unreliable AI models may be in dynamic and uncertain environments, especially in situations not previously encountered, to reflect the true nature of warfare [26]. However, known deployments of military AI systems, including uses of foundation models by the Israel Defense Forces (IDF) in Gaza, have demonstrated a trend of fielding undertested AI systems in the battlefield for target recommendations, despite evidence of the constitutive lack of accuracy and known failure modes [1, 9, 30, 95].

Field testing is the process of evaluating the performance, reliability, and suitability of defense technologies under real-world conditions through war games, modeling, simulation, technical assessments, and, finally, live-force operations [25]. Once deployed, technologies are subjected to ongoing maintenance testing to ensure the continued reliability of their fitness for purpose. However, despite notable evidence of foundational models' poor accuracy with non-English languages, especially for Arabic [61, 62], forces such as the IDF continue to use foundation models to guide in the identification of targets, the most critical of defense operations, through means such as the faulty transcription and translation of intercepted communications in Gaza [9]. To make matters worse, the use of foundation models may feed into other AI-driven systems, as with the IDF's use of Gospel, Lavender, and Where's Daddy, which have contributed to the significant civilian death toll in Gaza through the fallible collection and use of intelligence information for targeting purposes, and which may even subject their users to violations of IHL [70].

Indeed, the nested use of AI within decision-support systems will ultimately lead to compounding errors that propagate and accumulate through each subsequent phase of a kill chain [70], including for evaluating impacts of attacks required for field maintenance testing to ensure that technological outputs align with strategic and tactical defense objectives. Consider, for example, the IDF's algorithm predicting how many civilians might be

affected by attacks through the use of data from nearby cell towers to tally the number of civilians in an area [30]. This analysis in particular took "no account of whether a cellphone might be turned off or had run out of power or of children who wouldn't have a cellphone." Deploying AI models with profoundly low accuracy rates for target identification, compounded with flawed algorithms that are incapable of predicting the impact of AI systems' decisions, not only impairs the ability to verify and validate AI-driven outcomes, but will result in consequences that may not be aligned with strategic and tactical defense objectives. This includes the potential loss of human control over the use of force in an armed conflict, and likely violations of IHL [94].

Ultimately, the insistence on rapidly deploying foundation models in the highest-stakes defense applications (e.g., the DOD's AICC), despite their demonstrable failure modes and clear lack of fitness for warfighting activities, illustrates how general benchmarking contributes to a manufactured narrative where the purported increase of broad (and unrelated) model capabilities have been falsely positioned as better performance indicators, antithetical to tried-and-true field-testing metrics. Moreover, AI industry players further resist any form of ex ante scrutiny by positioning the kinds of evaluation and validation frameworks as overly burdensome, despite their wide utility across sectors [57]. This has enabled systems to be deployed that perform at levels far below the risk thresholds that would be deemed acceptable through standardized safety processes in other domains.

5 LOOKING AHEAD: AN AGENDA TO CORRECT COURSE

The establishment of societally acceptable safety goals and tolerances to determine the risk thresholds against which technological systems can be assessed is an arduous task that can take years or decades to establish. Despite the nearly insurmountable challenges posed by such tasks, there are precedents to look to, such as the risk determinations established for nuclear systems against the backdrop of the Cold War, which have provided invaluable safety and dependability goals through which defense and safety-critical systems have been evaluated. This worthwhile effort has helped the US establish its technological advantage and defense prowess over its adversaries. But rather than follow such historical norms that necessitated broad democratic deliberations, AI technologists—primarily industry labs, as well as "AI safety"-focused organizations—have staunchly advocated for a skewed cost-benefit justification driven by the unsubstantiated claims of an AI arms race and speculative "existential" capabilities risks. Under such pretenses, specifically for the use of AI in defense and national-security applications, AI technologists have sought to substitute traditional safety frameworks with ill-defined "capabilities" or "alignment" counterparts that deviate from well-established military standards necessary for the functionality of defense systems. This "safety revisionism" may be precisely what disadvantages US military and technological capabilities against China or other adversaries by allowing for the accelerated AI adoption at the cost of lowered safety and security thresholds.

As evidenced by the slew of existing national-security concerns emanating from uses of AI models that are ripe for adversaries to exploit and compromise, ceding the determination of risk tolerance to AI technologists ultimately leads to the degradation of existing defense safety and security practices. Furthermore, this places decision-making regarding the acceptable societal risk tolerance for AI, including number of deaths and the corresponding safety methodologies (i.e., via dissemination of co-opted safety frameworks), in the hands of AI technologists themselves. As noted by Chauncey Starr, a pioneer of probabilistic risk analysis, such patterns are signifiers of an autocratic risk value system, a far cry from messaging that the acceleration of AI use within the

US national-security apparatus would allow "democracies to maintain the lead" [3, 60]. The utilization of said skewed risk thresholds within national security, in combination with the demonstrable brittleness of foundational models [96, 101], further calls into question AI uses that may subject their users to violations of international humanitarian law [70], bring harms to both defense personnel and civilians in active war zones, and that ultimately normalizes fraught and inaccurate performance of AI in safety-critical contexts. Such normalization may imprint and set a precedent of low safety thresholds for civilian-critical infrastructure going forward (e.g., policing, energy, border governance, and so on). Indeed, as defense and civilian-safety critical standards are inextricably linked [54], lowered safety thresholds for AI use within national security may imperil the safety of civil society given that military standards can set the precedent under which civilian safety-critical systems are to be evaluated.

Situating themselves as the arbiters of life or death, AI technologists have selected risk thresholds that utilize skewed risk and safety interpretations at odds with the demands of social, political, and ethical civilian life, and with established assurance practices critical to the safety and security posture of national and defense infrastructure. To remedy this gap, global governance efforts must establish consensus-building bodies that congregate civil society, legislatures, and judiciary bodies to establish democratically deliberated risk tolerances that require the consideration of the societal evaluation of risk and, as such, interpreting public attitudes and values. Governing bodies must move away from hollow "responsible" and "trustworthy" frameworks that concern broad and hypothetical capabilities and AI "alignment," and maneuver toward defining AI risk thresholds that address "How safe is safe enough?" That is, they must produce AI safety objectives to safeguard defense and safety-critical systems in line with established norms, ultimately reckoning with the number of societally accepted fatalities and hazards that would require considerations such as to whom safety is afforded, under what contexts (e.g., human rights law versus IHL), and the exploration of concrete and substantiated AI risks.

Finally, policymakers and military stakeholders must preserve established military standards and risk thresholds, which have long ensured the safety and dependability of defense and critical systems, in spite of pressure to adopt manufactured metrics that serve to illustrate broad AI capabilities as an expression of military technological advantage. "Capabilities evaluations" and "red-teaming" are a weak substitute for existing TEVV frameworks that serve to evaluate a system's fitness for purpose in line with strategic and tactical defense objectives. Indeed, with the emergence of increasingly powerful AI defense systems, from LAWS to AI DSS, we must ensure their deployment is not only aligned with risk thresholds that continue to maintain the safety and security of critical and defense infrastructure, but that they remain in line with IHL. The deadly and geopolitically consequential impacts of AI within military applications brings with it existential risks that are very real and present; as they stand, the safety interpretations put forward by AI technologists will only result in a significant civilian death toll with no meaningful oversight or validation regarding their use, reliability, or safety.

REFERENCES

[1] Yuval Abraham. 'Lavender': The AI machine directing Israel's bombing spree in Gaza, April 2024. URL https://www.972mag.com/lavender-ai-israeli-army-gaza/.

[2] Jack Aldane. Canada launches first AI strategy for federal public service - Global Government Forum, April 2025. URL https://www.globalgovernmentforum.com/canada-launches-first-ai-strategy-for-federal-public-service/.

[3] Dario Amodei and Matt Pottinger. Trump Can Keep America's AI Advantage. *Wall Street Journal*, January 2025. URL https://www.wsj.com/opinion/trump-can-keep-americas-ai-advantage-china-chips-data-eccdce91.

[4] Markus Anderljung, Joslyn Barnhart, Anton Korinek, Jade Leung, Cullen O'Keefe, Jess Whittlestone, Shahar Avin, Miles Brundage, Justin Bullock, Duncan Cass-Beggs, Ben Chang, Tantum Collins, Tim Fist, Gillian Hadfield, Alan Hayes, Lewis Ho, Sara Hooker, Eric Horvitz, Noam Kolt, Jonas Schuett, Yonadav Shavit, Divya Siddarth, Robert Trager, and Kevin Wolf. Frontier AI Regulation: Managing Emerging Risks to Public Safety. *arXiv*, 2023. doi: 10.48550/ARXIV.2307.03718. URL https://arxiv.org/abs/2307.03718.

[5] Anduril. Anduril Partners with OpenAI to Advance U.S. Artificial Intelligence Leadership and Protect U.S. and Allied Forces, December 2024. URL https://www.anduril.com/anduril-partners-with-openai-to-advance-u-s-artificial-intelligence-leadership-and-protect-u-s/.

[6] Anthropic. Anthropic's Responsible Scaling Policy, September 2023. URL https://www.anthropic.com/news/anthropics-responsible-scaling-policy.

[7] National Protective Security Authority. Security-informed safety, January 2023. URL https://www.npsa.gov.uk/security-informed-safety.

[8] Haydn Belfield and Christian Ruhl. Why policy makers should beware claims of new 'arms races'. *Bulletin of the Atomic Scientists*, July 2022. URL https://thebulletin.org/2022/07/why-policy-makers-should-beware-claims-of-new-arms-races/.

[9] Michael Biesecker, Sam Mednick, and Garance Burke. As Israel uses US-made AI models in war, concerns arise about tech's role in who lives and who dies. *AP News*, February 2025. URL https://apnews.com/article/israel-palestinians-ai-technology-737bc17af7b03e98c29cec4e15d0f108.

[10] Robin Bloomfield and John Rushby. Assurance of AI Systems From a Dependability Perspective. *arXiv*, August 2024. doi: 10.48550/arXiv.2407.13948. URL http://arxiv.org/abs/2407.13948.

[11] Robin Bloomfield, Kate Netkachova, and John Rushby. Defeaters and eliminative argumentation in assurance 2.0. *arXiv*, 2024. doi: 10.48550/arXiv.2405.15800. URL https://arxiv.org/abs/2405.15800.

[12] Brendan Bordelon. How a billionaire-backed network of AI advisers took over Washington. *POLITICO*, October 2023. URL https://www.politico.com/news/2023/10/13/open-philanthropy-funding-ai-policy-00121362.

[13] Brendan Bordelon. Key Congress staffers in AI debate are funded by tech giants like Google and Microsoft. *POLITICO*, March 2023. URL https://www.politico.com/news/2023/12/03/congress-ai-fellows-tech-companies-00129701.

[14] Vincent Boulanin. The Risks of Integrating Generative AI into Weapon Systems. Technical report, GC REAIM Expert Policy Note Series, April 2025.

[15] Michael Brenes and William D. Hartung. Private Finance and the Quest to Remake Modern Warfare, June 2024. URL https://quincyinst.org/research/private-finance-and-the-quest-to-remake-modern-warfare/.

[16] Tom Bristow. Britain dances to JD Vance's tune as it renames AI institute. *POLITICO*, February 2025. URL https://www.politico.eu/article/jd-vance-britain-ai-safety-institute-aisi-security/.

[17] Daniel S. Brown, Jordan Schneider, Anca D. Dragan, and Scott Niekum. Value Alignment Verification. *arXiv*, June 2021. doi: 10.48550/arXiv.2012.01557. URL http://arxiv.org/abs/2012.01557. arXiv:2012.01557.

[18] Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwag, Florian Tramèr, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models. In *Proceedings of the 32nd USENIX Conference on Security Symposium (SEC '23). USENIX Association, USA, Article 294, 5253–5270*, 2023.

[19] Nicholas Carlini, Matthew Jagielski, Christopher A. Choquette-Choo, Daniel Paleka, Will Pearce, Hyrum Anderson, Andreas Terzis, Kurt Thomas, and Florian Tramèr. Poisoning Web-Scale Training Datasets is Practical. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 407–425, San Francisco, CA, USA, May 2024. IEEE. ISBN 9798350331301. doi: 10.1109/SP54263.2024.00179. URL https://ieeexplore.ieee.org/document/10646610/.

[20] Nicholas Carlini, Daniel Paleka, Krishnamurthy Dj Dvijotham, Thomas Steinke, Jonathan Hayase, A. Feder Cooper, Katherine Lee, Matthew Jagielski, Milad Nasr, Arthur Conmy, Itay Yona, Eric Wallace, David Rolnick, and Florian Tramèr. Stealing Part of a Production Language Model. *arXiv*, March 2024. URL https://arxiv.org/abs/2403.06634v2.

[21] Stephen Cave and Seán S. ÓhÉigeartaigh. An AI Race for Strategic Advantage: Rhetoric and Risks. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 36–40, New Orleans LA USA, December 2018. ACM. ISBN 9781450360128. doi: 10.1145/3278721.3278780. URL https://dl.acm.org/doi/10.1145/3278721.3278780.

[22] Brian J. Chen. Dispelling Myths of AI and Efficiency. *Data & Society*, 2025. URL https://datasociety.net/library/dispelling-myths-of-ai-and-efficiency/?trk=feed_main-feed-card_feed-article-content.

[23] Joshua Clymer, Nick Gabrieli, David Krueger, and Thomas Larsen. Safety Cases: How to Justify the Safety of Advanced AI Systems. *arXiv*, March 2024. doi: 10.48550/arXiv.2403.10462. URL http://arxiv.org/abs/2403.10462.

[24] U.S. Nuclear Regulatory Commission. Safety Goals for Nuclear Power Plant Operation. U.S. Nuclear Regulatory Commission, May 1983. URL https://www.nrc.gov/docs/ML0717/ML071770230.pdf.

[25] The United States Marine Corps. Marine Corps Warfighting Laboratory: Experiment Division. URL https://www.mcwl.marines.mil/Divisions/Experiment/.

[26] Mary (Missy) Cummings. The AI that Wasn't There: Global Order and the (Mis)Perception of Powerful AI. *Policy Roundtable: Artificial Intelligence and International Security, Texas National Security Review*, June 2020. URL https://tnsr.org/roundtable/policy-roundtable-artificial-intelligence-and-international-security/.

[27] M.L. Cummings and Ben Bauchwitz. Identifying Research Gaps through Self-Driving Car Data Analysis. *IEEE Transactions on Intelligent Vehicles*, pages 1–10, 2024. ISSN 2379-8904. doi: 10.1109/TIV.2024.3506936. URL https://ieeexplore.ieee.org/document/10778107/?arnumber=10778107.

[28] Google DeepMind. Introducing the Frontier Safety Framework, April 2025. URL https://deepmind.google/discover/blog/introducing-the-frontier-safety-framework/.

[29] Bill Drexel and Hannah Kelley. Behind China's Plans to Build AI for the World. *POLITICO*, November 2023. URL https://www.politico.com/news/magazine/2023/11/30/china-global-ai-plans-00129160.

[30] Elizabeth Dwoskin. Israel built an 'AI factory' for war. It unleashed it in Gaza. *The Washington Post*, December 2024. URL https://www.washingtonpost.com/technology/2024/12/29/ai-israel-war-gaza-idf/.

[31] Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, Andy Jones, Sam Bowman, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Nelson Elhage, Sheer El-Showk, Stanislav Fort, Zac Hatfield-Dodds, Tom Henighan, Danny Hernandez, Tristan Hume, Josh Jacobson, Scott Johnston, Shauna Kravec, Catherine Olsson, Sam Ringer, Eli Tran-Johnson, Dario Amodei, Tom Brown, Nicholas Joseph, Sam McCandlish, Chris Olah, Jared Kaplan, and Jack Clark. Red Teaming Language Models to Reduce Harms: Methods, Scaling Behaviors, and Lessons Learned. *arXiv*, November 2022. URL http://arxiv.org/abs/2209.07858.

[32] Paul M. Grant. Chauncey Starr (1912–2007). *Nature*, 447(7146):789–789, June 2007. ISSN 1476-4687. doi: 10.1038/447789a. URL https://www.nature.com/articles/447789a.

[33] Ibrahim Habli and John Alexander McDermid. AI Safety: Navigating the Expanding Landscape of Potential Harms. *Safety-Critical Systems Club Newsletter*, May 2024. URL https://eprints.whiterose.ac.uk/id/eprint/213035/.

[34] Ibrahim Habli, Richard Hawkins, Colin Paterson, Philippa Ryan, Yan Jia, Mark Sujan, and John McDermid. The BIG Argument for AI Safety Cases. *arXiv*, March 2025. doi: 10.48550/arXiv.2503.11705. URL http://arxiv.org/abs/2503.11705.

[35] Justin Haner and Denise Garcia. The Artificial Intelligence Arms Race: Trends and World Leaders in Autonomous Weapons Development. *Global Policy*, 10(3):331–337, September 2019. ISSN 1758-5880, 1758-5899. doi: 10.1111/1758-5899.12713. URL https://onlinelibrary.wiley.com/doi/10.1111/1758-5899.12713.

[36] Dan Hendrycks and Eric Schmidt. The Nuclear-Level Risk of Superintelligent AI. *TIME*, March 2025. URL https://time.com/7265056/nuclear-level-risk-of-superintelligent-ai/.

[37] Coffee House. Read: JD Vance's full speech on AI and the EU. *The Spectator*, February 2025. URL https://www.spectator.co.uk/article/read-jd-vances-full-speech-on-ai-and-the-eu/.

[38] AI Security Institute. Tackling AI Security Risks to Unleash Growth and Deliver Plan for Change. *Department for Science, Innovation and Technology*, February 2025. URL https://www.gov.uk/government/news/tackling-ai-security-risks-to-unleash-growth-and-deliver-plan-for-change.

[39] Geoffrey Irving. Safety Cases at AISI. *AI Security Institute*, August 2024. URL https://www.aisi.gov.uk/work/safety-cases-at-aisi.

[40] Sheila Jasanoff. *The Fifth Branch*. Harvard University Press, 1994. URL https://www.hup.harvard.edu/books/9780674300620.

[41] Sheila Jasanoff and Sang-Hyun Kim. Containing the Atom: Sociotechnical Imaginaries and Nuclear Power in the United States and South Korea. *Minerva*, 47(2):119–146, June 2009. ISSN 0026-4695, 1573-1871. doi: 10.1007/s11024-009-9124-4. URL http://link.springer.com/10.1007/s11024-009-9124-4.

[42] Elliot Jones, Mahi Hardalupas, and William Agnew. Under the radar? examining the evaluation of foundation models. *Ada Lovelace Institute*, July 2024. URL https://www.adalovelaceinstitute.org/wp-content/uploads/2024/09/Ada-Lovelace-Institute-Under-the-radar-230924.pdf.

[43] Amba Kak and Sarah Myers West. 2023 Landscape: Confronting Tech Power. *AI Now Institute*, April 2023. URL https://ainowinstitute.org/wp-content/uploads/2023/04/AI-Now-2023-Landscape-Report-FINAL.pdf.

[44] Nidhi Kalra and Susan M. Paddock. Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability? Technical report, RAND Corporation, April 2016. URL https://www.rand.org/pubs/research_reports/RR1478.html.

[45] Atoosa Kasirzadeh. Measurement challenges in AI catastrophic risk governance and safety frameworks. *Tech Policy Press*, Sep 2024. URL https://www.techpolicy.press/measurement-challenges-in-ai-catastrophic-risk-governance-and-safety-frameworks/.

[46] Raj kamal Kaur, Babita Pandey, and Lalit Kumar Singh. Dependability analysis of safety critical systems: Issues and challenges. *Annals of Nuclear Energy*, 120:127–154, October 2018. ISSN 0306-4549. doi: 10.1016/j.anucene.2018.05.027. URL https://www.sciencedirect.com/science/article/pii/S030645491730213X.

[47] Danielle Kehl, Andi Wilson, and Kevin Bankston. Doomed to Repeat History? Lessons from the Crypto Wars of the 1990s. URL http://newamerica.org/cybersecurity-initiative/policy-papers/doomed-to-repeat-history-lessons-from-the-crypto-wars-of-the-1990s/.

[48] Makena Kelly. Elon Musk Ally Tells Staff 'AI-First' Is the Future of Key Government Agency. *Wired*, 2025. ISSN 1059-1028. URL https://www.wired.com/story/elon-musk-lieutenant-gsa-ai-agency/.

[49] Tim Kelly. A Systematic Approach to Safety Case Management. *SAE Transactions*, 113:257–266, 2004. ISSN 0096-736X. URL https://www.jstor.org/stable/44699541.

[50] Heidy Khlaaf. How AI Can Be Regulated Like Nuclear Energy. *TIME*, October 2023. URL https://time.com/6327635/ai-needs-to-be-regulated-like-nuclear-weapons/.

[51] Heidy Khlaaf. Toward Comprehensive Risk Assessments and Assurance of AI-Based Systems. *Trail of Bits*, March 2023. URL https://www.trailofbits.com/documents/Toward_comprehensive_risk_assessments.pdf.

[52] Heidy Khlaaf, Pamela Mishkin, Joshua Achiam, Gretchen Krueger, and Miles Brundage. A Hazard Analysis Framework for Code Synthesis Large Language Models. *arXiv*, July 2022. URL https://arxiv.org/abs/2207.14157v1.

[53] Heidy Khlaaf, Sarah Myers West, and Meredith Whittaker. Mind the Gap: Foundation Models and the Covert Proliferation of Military Intelligence, Surveillance, and Targeting. *arXiv*, October 2024. doi: 10.48550/arXiv.2410.14831. URL http://arxiv.org/abs/2410.14831.

[54] Steven Lapsley. Nato use of civil standards, October 2018. URL https://www.dsp.dla.mil/Portals/26/Documents/Publications/Conferences/2018/2018%20International%20Standardization%20Workshop/20181031-Item4-UK_NATO_UseOfCivilStandards-IntlStdznWorkshop_Lapsley.pdf?ver=2018-11-06-151834-190.

[55] Task Force Lima. Task Force Lima Executive Summary. Technical report, US Department of Defense, December 2024.

[56] Donald A. MacKenzie. *Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance*. MIT Press, 1990. ISBN 9780262132589.

[57] Gianclaudio Malgieri and Frank A. Pasquale. From transparency to justification: Toward ex ante accountability for ai. *SSRN Electronic Journal*, 2022. doi: 10.2139/ssrn.4099657.

[58] Rowena Mason and Rowena Mason Whitehall editor. AI should replace some work of civil servants, Starmer to announce. *The Guardian*, March 2025. ISSN 0261-3077. URL https://www.theguardian.com/technology/2025/mar/12/ai-should-replace-some-work-of-civil-servants-under-new-rules-keir-starmer-to-announce.

[59] John McDermid, Yan Jia, and Ibrahim Habli. Upstream and Downstream AI Safety: Both on the Same River? *arXiv*, December 2024. doi: 10.48550/arXiv.2501.05455. URL http://arxiv.org/abs/2501.05455.

[60] Madhumita Murgia. Anthropic's Dario Amodei: Democracies must maintain the lead in AI. *Financial Times*, December 2024. URL https://www.ft.com/content/e75e3388-4700-413d-ab67-778410c2d977.

[61] Tarek Naous and Wei Xu. On The Origin of Cultural Biases in Language Models: From Pre-training Data to Linguistic Phenomena. *arXiv*, January 2025. doi: 10.48550/arXiv.2501.04662. URL http://arxiv.org/abs/2501.04662.

[62] Tarek Naous, Michael J Ryan, Alan Ritter, and Wei Xu. Having Beer after Prayer? Measuring Cultural Bias in Large Language Models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16366–16393, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.862. URL https://aclanthology.org/2024.acl-long.862/.

[63] Arvind Narayanan and Sayash Kapoor. AI as Normal Technology. *Columbia University, The Knight First Amendment Institute*, 2025. URL https://www.aisnakeoil.com/p/ai-as-normal-technology.

[64] Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A. Feder Cooper, Daphne Ippolito, Christopher A. Choquette-Choo, Eric Wallace, Florian Tramèr, and Katherine Lee. Scalable Extraction of Training Data from (Production) Language Models, 2023. URL

https://arxiv.org/abs/2311.17035.

[65] NIST. U.S. AI Safety Institute Establishes New U.S. Government Taskforce to Collaborate on Research and Testing of AI Models to Manage National Security Capabilities & Risks, November 2024. URL https://www.nist.gov/news-events/news/2024/11/us-ai-safety-institute-establishes-new-us-government-taskforce-collaborate.

[66] Ministry of Defence. Implementation of improved Air System Safety Case Regulation - RA 1205, 2019. URL https://www.gov.uk/government/news/implementation-of-improved-air-system-safety-case-regulation-ra-1205.

[67] The U.S. Department of Defense. Standard Practice: System Safety. Technical Report MIL-STD-882E, The U.S. Department of Defense, May 2012.

[68] U.S. Department of Defense. DoD Instruction 5000.61, "DoD Modeling and Simulation (M&S) Verification, Validation, and Accreditation (VV&A),". U.S. Department of Defense, September 2024. URL https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500061p.pdf.

[69] Department of Homeland Security. Artificial Intelligence Safety and Security Board, 2025. URL https://www.dhs.gov/artificial-intelligence-safety-and-security-board.

[70] International Committee of the Red Cross. 2024 ICRC IHL Challenges Report | ICRC, September 2024. URL https://www.icrc.org/en/report/2024-icrc-report-ihl-challenges.

[71] David Okrent. The Safety Goals of the U.S. Nuclear Regulatory Commission. *Science*, 236(4799):296–300, April 1987. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.3563510. URL https://www.science.org/doi/10.1126/science.3563510.

[72] Amy Orben and J. Nathan Matias. Fixing the science of digital technology harms. *Science*, 388(6743):152–155, April 2025. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.adt6807. URL https://www.science.org/doi/10.1126/science.adt6807.

[73] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *arXiv*, March 2022. URL https://arxiv.org/abs/2203.02155v1.

[74] Palantir. Palantir IR, 2024. URL https://investors.palantir.com/news-details/2024/Anthropic-and-Palantir-Partner-to-Bring-Claude-AI-Models-to-AWS-for-U.S.-Government-Intelligence-and-Defense-Operations/.

[75] Ethan Perez, Saffron Huang, Francis Song, Trevor Cai, Roman Ring, John Aslanides, Amelia Glaese, Nat McAleese, and Geoffrey Irving. Red Teaming Language Models with Language Models. *arXiv*, February 2022. doi: 10.48550/arXiv.2202.03286. URL http://arxiv.org/abs/2202.03286.

[76] POWER. Chauncey Starr: A personal memoir, June 2007. URL https://www.powermag.com/chauncey-starr-a-personal-memoir/.

[77] Primer. Introducing Primer Delta: Our next-gen AI-native platform to transform information overload into decision advantage, April 2023. URL https://primer.ai/business-solutions/introducing-primer-delta-our-next-gen-ai-native-platform-to-transform-information-overload-into-decision-advantage/.

[78] Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. Fine-tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To! *arXiv*, October 2023. URL https://arxiv.org/abs/2310.03693v1.

[79] Inioluwa Deborah Raji and Roel Dobbe. Concrete Problems in AI Safety, Revisited. *arXiv*, December 2023. doi: 10.48550/arXiv.2401.10899. URL http://arxiv.org/abs/2401.10899.

[80] Jane O. Rathbun. Don guidance on the use of generative artificial intelligence and large language models. *Department of Navy Chief Information Officer*, 2023. URL https://www.doncio.navy.mil/ContentView.aspx?id=16442.

[81] Heather M. Roff. The frame problem: The AI "arms race" isn't one. *Bulletin of the Atomic Scientists*, 75(3):95–98, May 2019. ISSN 0096-3402, 1938-3282. doi: 10.1080/00963402.2019.1604836. URL https://www.tandfonline.com/doi/full/10.1080/00963402.2019.1604836.

[82] Charles Rollet. Fintech Founder Charged with Fraud after "AI" Shopping App Found to be Powered by Humans in the Philippines. *TechCrunch*, Apr 2025. URL https://techcrunch.com/2025/04/10/fintech-founder-charged-with-fraud-after-ai-shopping-app-found-to-be-powered-by-humans-in-the-philippines/.

[83] Eryk Salvaggio. Anatomy of an AI Coup. *Tech Policy Press*, February 2025. URL https://techpolicy.press/anatomy-of-an-ai-coup.

[84] Elke Schwarz and DePaul University. Silicon Valley Goes to War: Artificial Intelligence, Weapons Systems, and Moral Agency. *Philosophy Today*, 65(3):549–569, 2021. ISSN 0031-8256. doi: 10.5840/philtoday2021519407. URL http://www.pdcnet.org/oom/service?url_ver=Z39.88-2004&rft_val_fmt=&rft.imuse_id=philtoday_2021_0065_0003_0549_0569&svc_id=info:www.pdcnet.org/collection.

[85] Jonathan Spring and Divjot Singh Bawa. AI Red Teaming: Applying Software TEVV for AI Evaluations | CISA, November 2024. URL https://www.cisa.gov/news-events/news/ai-red-teaming-applying-software-tevv-ai-evaluations.

[86] Chauncey Starr. Risk Criteria for Nuclear Power Plants: A Pragmatic Proposal. *Risk Analysis*, 1(2):113–120, June 1981. ISSN 0272-4332, 1539-6924. doi: 10.1111/j.1539-6924.1981.tb01406.x. URL https://onlinelibrary.wiley.com/doi/10.1111/j.1539-6924.1981.tb01406.x.

[87] Chauncey Starr and Chris Whipple. Risks of Risk Decisions. In *Risk In The Technological Society*. Routledge, 1982. ISBN 9780429304873.

[88] Chauncey Starr, Richard Rudman, and Chris Whipple. Philosophical Basis for Risk Analysis. *Annual Review of Energy*, 1(1):629–662, November 1976. ISSN 0362-1626. doi: 10.1146/annurev.eg.01.110176.003213. URL https://www.annualreviews.org/doi/10.1146/annurev.eg.01.110176.003213.

[89] Jack Stilgoe. How can we know a self-driving car is safe? *Ethics and Information Technology*, 23(4):635–647, December 2021. ISSN 1572-8439. doi: 10.1007/s10676-021-09602-1. URL https://doi.org/10.1007/s10676-021-09602-1.

[90] Helen Toner, Jenny Xiao, and Jeffrey Ding. The Illusion of China's AI Prowess. *Foreign Affairs*, June 2023. URL https://www.foreignaffairs.com/china/illusion-chinas-ai-prowess-regulation-helen-toner.

[91] Defense Innovation Unit. Thunderforge Project: Integrating Commercial AI-Powered Decision-Making, 2025. URL https://www.diu.mil/latest/dius-thunderforge-project-to-integrate-commercial-ai-powered-decision-making.

[92] McElheny Victor K. 'Breeder' Reactor Plan Facing Delays. *The New York Times*, December 1973.

[93] Brandi Vincent. Inside Task Force Lima's exploration of 180-plus generative AI use cases for DOD. *DefenseScoop*, November 2023. URL https://defensescoop.com/2023/11/06/inside-task-force-limas-exploration-of-180-plus-generative-ai-use-cases-for-dod/.

[94] Jimena Sofía Viveros Álvarez. The risks and inefficacies of AI systems in military targeting support. *ICRC Humanitarian Law & Policy Blog*, September 2024. URL https://blogs.icrc.org/law-and-policy/2024/09/04/the-risks-and-inefficacies-of-ai-systems-in-military-targeting-support/.

[95] Human Rights Watch. Gaza: Israeli military's digital tools risk civilian harm, 2024. URL https://www.hrw.org/news/2024/09/10/gaza-israeli-militarys-digital-tools-risk-civilian-harm.

[96] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, Zac Kenton, Sasha Brown, Will Hawkins, Tom Stepleton, Courtney Biles, Abeba Birhane, Julia Haas, Laura Rimell, Lisa Anne Hendricks, William Isaac, Sean Legassick, Geoffrey Irving, and Iason Gabriel. Ethical and social risks of harm from Language Models. *arXiv*, 2021. doi: 10.48550/ARXIV.2112.04359. URL https://arxiv.org/abs/2112.04359.

[97] Jessica Chen Weiss. The China Trap: U.S. Foreign Policy and the Perilous Logic of Zero-Sum Competition. *Foreign Affairs*, August 2022. URL https://www.foreignaffairs.com/china/china-trap-us-foreign-policy-zero-sum-competition.

[98] Matteo Wong. DOGE's Plans to Replace Humans With AI Are Already Under Way, March 2025. URL https://www.theatlantic.com/technology/archive/2025/03/gsa-chat-doge-ai/681987/.

[99] Robert O. Work and Greg Grant. Beating the Americans at their Own Game. An Offset Strategy with Chinese Characteristics. *Center for a New American Security*, 2019. ISSN 2510-2648, 2510-263X. doi: 10.1515/sirius-2019-4022.

[100] Yiming Zhang, Javier Rando, Ivan Evtimov, Jianfeng Chi, Eric Michael Smith, Nicholas Carlini, Florian Tramèr, and Daphne Ippolito. Persistent Pre-Training Poisoning of LLMs. *arXiv*, October 2024. doi: 10.48550/arXiv.2410.13722. URL http://arxiv.org/abs/2410.13722.

[101] Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J. Zico Kolter, and Matt Fredrikson. Universal and Transferable Adversarial Attacks on Aligned Language Models. *arXiv*, December 2023. doi: 10.48550/arXiv.2307.15043. URL http://arxiv.org/abs/2307.15043.