

# FROM DUAL-USE AWARENESS TO DUAL-USE AGENCY: EPISTEMIC DISTANCE AND THE STRUCTURAL LIMITS OF ETHICAL RESPONSIBILITY IN AI RESEARCH

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Contemporary AI research is increasingly implicated in military and surveillance systems, often through indirect civilian pathways rather than explicit defense programs. Existing responses to this dual-use challenge have focused primarily on improving researcher awareness and ethical reflection. In this paper, we argue that such approaches overlook a more fundamental constraint: researchers frequently lack the institutional capacity to influence how their work is ultimately deployed. We introduce the concept of a dual-use agency gap to describe this mismatch between ethical responsibility and practical control. We further argue that this gap is produced by epistemic distance—the growing separation between sites of AI research and sites of deployment. Drawing on recent scholarship on surveillance pipelines, foundation models, and professional responsibility, we show how contemporary research infrastructures systematically limit researchers’ ability to contest or refuse downstream military and surveillance uses. We conclude by outlining modest, community-level interventions aimed at preserving researcher agency and supporting peace-oriented AI research grounded in harm prevention.

## 1 INTRODUCTION

Artificial intelligence research has long been shaped by military and security priorities, from early work in computer vision to contemporary developments in large-scale machine learning systems Crawford (2021); Dobson (2023). While some AI research is explicitly funded or developed for defense purposes, much of today’s militarization occurs through civilian research pathways with dual-use potential. Research artifacts produced in academic and industrial settings—models, datasets, benchmarks—are routinely incorporated into military and surveillance systems without direct involvement from the researchers who developed them Kalluri et al. (2023).

As a result, many researchers remain only partially aware of how their work is deployed in conflict or security contexts, or of the legal and ethical implications of those deployments. Schwartz et al. (2022); Fereidooni & Heidt (2024). Public and policy debates have increasingly focused on these issues, including concerns about autonomous weapons, mass surveillance, and the use of foundation models in intelligence and targeting. infrastructures Future of Life Institute (2016); Khlaaf et al. (2024); Romansky (2025) However, within mainstream machine learning venues, researchers have limited opportunities to collectively articulate positions on militarization or to meaningfully influence downstream uses of their work.

Current responses within the research community tend to emphasize individual awareness, disclosure, and ethical reflection Goodfriend (2023); Khlaaf et al. (2024). While important, these approaches implicitly assume that researchers who recognize dual-use risks are able to act upon that knowledge. In practice, this assumption often fails. Once research outputs enter broader dissemination pipelines, control over their use is typically transferred to institutions, intermediaries, or downstream actors. In this paper, we argue that the core challenge is not a lack of awareness, but a lack of agency. We conceptualize this condition as a dual-use agency gap: a structural separation between ethical responsibility and the capacity to intervene. We further argue that this gap is produced

054 by epistemic distance—the institutional and informational separation between research production  
055 and deployment contexts. By examining how epistemic distance operates within contemporary AI  
056 research pipelines, we aim to clarify why existing dual-use frameworks fall short and to identify  
057 avenues for agency-preserving research practices.

## 060 2 DUAL-USE AWARENESS AND ITS STRUCTURAL LIMITS

062 Dual-use concerns have received sustained attention in recent AI ethics and governance scholarship.  
063 Researchers have examined how ostensibly benign technologies can be repurposed for military or  
064 surveillance applications, and have proposed various strategies for mitigating associated risks Dob-  
065 son (2023); Loewenstein (2024); Romansky (2025). These strategies often emphasize the impor-  
066 tance of researcher awareness, transparency, and ethical deliberation.

067 Such efforts have been valuable in foregrounding ethical responsibility within AI research. How-  
068 ever, they tend to locate responsibility primarily at the level of individual cognition and intent. The  
069 underlying assumption is that informed researchers can adjust their behavior—by avoiding certain  
070 projects, modifying research designs, or choosing not to publish—thereby reducing harm.

071 Empirical accounts of AI research pipelines complicate this picture. As AI research becomes in-  
072 creasingly modular, reusable, and infrastructural, individual researchers exercise diminishing con-  
073 trol over how their contributions are recombined and redeployed Crawford (2021); Schwartz et al.  
074 (2022). Awareness of potential harm does not necessarily translate into the ability to prevent or con-  
075 test downstream use. In this context, ethical responsibility risks becoming aspirational rather than  
076 actionable.

077 This suggests a limitation in prevailing dual-use frameworks. By focusing on awareness without  
078 attending to institutional constraints, they obscure the conditions under which ethical agency is  
079 exercised—or foreclosed—within contemporary research ecosystems.

## 082 3 EPISTEMIC DISTANCE AND THE DUAL-USE AGENCY GAP

084 To better understand this constraint, we introduce the concept of epistemic distance. Epistemic  
085 distance refers to the structured separation between those who produce AI research and those who  
086 deploy it, encompassing differences in institutional authority, access to information, and decision-  
087 making power. Crucially, epistemic distance is not simply a matter of ignorance; it is a stable feature  
088 of how AI research is organized and disseminated.

089 Recent analyses of the surveillance AI pipeline illustrate how research outputs are absorbed into  
090 security infrastructures through layers of intermediaries, including private firms, government con-  
091 tractors, and state agencies Brenneis (2025). Similar dynamics are evident in the context of founda-  
092 tion models, where general-purpose systems developed for civilian applications are integrated into  
093 military and intelligence operations without transparency or researcher consent Loewenstein (2024).

094 Epistemic distance plays a central role in producing what we term the dual-use agency gap. Re-  
095 searchers may recognize that their work contributes to surveillance or military systems, yet re-  
096 main unable to verify specific deployments, influence institutional decisions, or withdraw partici-  
097 pation. Control over research outputs is often governed by funding agreements, intellectual property  
098 regimes, and publication norms that prioritize openness and reuse while offering few mechanisms  
099 for contestation.

100 The result is a redistribution of agency. Ethical responsibility remains rhetorically attached to re-  
101 searchers, while practical authority over deployment is exercised elsewhere. This gap is particularly  
102 consequential in cases where AI systems contribute to forms of structural harm, including pervasive  
103 surveillance and the targeting of marginalized populations. Kalluri et al. (2023); Fereidooni & Heidt  
104 (2024); Goodfriend (2023)

105 Reframing militarization as a problem of agency rather than intent shifts the analytical focus. It  
106 directs attention away from individual moral failure and toward the institutional conditions that  
107 shape what researchers can and cannot do in practice.

Table 1: Points of intervention for reducing the dual-use agency gap in AI research pipelines

PIPELINE STAGE	SOURCE OF EPISTEMIC DISTANCE	RESULTING AGENCY LOSS	POTENTIAL INTERVENTION
Funding & partnerships	Opaque downstream collaborators; shifting project scope	Researchers unable to refuse military redirection	Conditional funding clauses; scope-change notification norms
Model & dataset release	Open reuse without traceability	Loss of visibility into downstream deployment	Voluntary downstream-use registries; community tracking
Publication & benchmarking	Incentives favor reuse and scale over context	Ethical concerns detached from dissemination	Contextualized release statements; risk-aware publication norms
Deployment via intermediaries	Institutional control held by firms or state actors	Researchers excluded from deployment decisions	Collective contestation mechanisms through professional bodies

## 4 TOWARD AGENCY-PRESERVING RESEARCH PRACTICES

If the dual-use agency gap is structurally produced, then responses must extend beyond individual ethical reflection. Without proposing comprehensive solutions, we outline several modest intervention points that align with existing research practices.

First, greater transparency around downstream military and surveillance adoption could help reduce epistemic distance. This might include community-maintained records of significant security integrations of widely used models or datasets. Second, funding and collaboration agreements could incorporate provisions that recognize researchers’ ability to decline continued involvement when projects shift toward military deployment. Third, professional communities could develop shared norms for collective response, enabling researchers to contest harmful uses without acting in isolation.

These measures frame peace-oriented AI research not as opposition to technological development, but as an effort to align research governance with harm prevention. Similar approaches have emerged in other scientific fields with long-standing dual-use concerns, where professional norms and institutional safeguards play a central role in mediating ethical risk.

### 4.1 INTERPRETING INTERVENTION POINTS ACROSS THE RESEARCH PIPELINE

Table 1 illustrates how epistemic distance manifests at multiple stages of the AI research pipeline, and how the resulting loss of agency is not confined to any single moment of research production or dissemination. Instead, the dual-use agency gap emerges cumulatively, as control is progressively transferred away from researchers through funding arrangements, publication norms, and institutional intermediaries.

A key implication of this analysis is that agency-preserving interventions need not be comprehensive or centralized in order to be effective. Rather than attempting to eliminate dual-use risk altogether, the intervention points identified in Table 1 emphasize restoring limited but meaningful forms of visibility, contestability, and collective response at moments where epistemic distance is greatest. Importantly, these interventions operate within existing research practices, such as funding agreements, release norms, and professional governance structures.

The table also highlights an asymmetry between responsibility and control. While ethical responsibility is typically assigned to researchers across all pipeline stages, the capacity to influence downstream deployment diminishes most sharply at the point of intermediation, where decisions are made by firms or state actors beyond the reach of individual researchers. This asymmetry helps explain why calls for increased awareness or disclosure often fail to translate into material change, and why collective mechanisms—rather than individual ethical action—are particularly salient for peace-oriented AI research.

By making explicit where and how agency is lost, the table serves not as a prescriptive checklist, but as a diagnostic tool for identifying leverage points within contemporary AI research ecosystems. In this sense, it complements existing dual-use frameworks by shifting attention from researcher intent to the institutional conditions that shape the consequences of research dissemination.

## 5 DISCUSSION: IMPLICATIONS FOR AI FOR PEACE RESEARCH

### 5.1 PEACE AS HARM PREVENTION, NOT POST-HOC MITIGATION

Much of the emerging discourse on “AI for peace” focuses on developing applications intended to counteract violence, support humanitarian action, or mitigate the harms of conflict. While such efforts are valuable, the analysis in this paper suggests a more foundational challenge. When militarization occurs through ordinary research pipelines, peace-oriented AI cannot be treated solely as a matter of downstream application choice.

The dual-use agency gap reframes peace as a question of harm prevention within research infrastructures, rather than post-hoc harm mitigation. If researchers lack the capacity to influence how their work is appropriated into military and surveillance systems, then the distinction between “peaceful” and “harmful” applications becomes increasingly tenuous. Peace-oriented research, on this view, requires attention to how epistemic distance is produced and maintained, and to how agency can be preserved at earlier stages of the research lifecycle.

This framing aligns peace-building not with opposition to AI development, but with the design of institutional conditions under which research dissemination does not systematically foreclose ethical intervention. In this sense, peace becomes a property of research governance, not merely of research intent.

### 5.2 REPOSITIONING THE ML RESEARCHER IN MILITARIZED PIPELINES

A recurring assumption in dual-use debates is that researchers occupy a privileged epistemic position from which they can anticipate and manage downstream risks. The concept of epistemic distance challenges this assumption. As AI research becomes increasingly infrastructural, researchers are often positioned upstream of deployment decisions, with limited visibility and authority once their work enters broader ecosystems.

For AI for peace initiatives, this suggests the need to rethink the role assigned to the machine learning researcher. Rather than treating researchers as individual ethical gatekeepers, peace-oriented frameworks may be better served by recognizing researchers as participants in collective, institutionally constrained systems. From this perspective, mechanisms that enable collective contestation, refusal, or withdrawal are not ancillary ethical tools, but central components of responsible research practice.

This repositioning also helps explain why calls for individual ethical responsibility frequently fail to translate into material change. Without corresponding shifts in institutional design, ethical reflection risks becoming decoupled from action. Addressing the dual-use agency gap therefore requires not only normative commitment, but shared professional infrastructures capable of supporting agency in practice.

## 6 CONCLUSION

Efforts to address AI militarization have largely focused on improving researcher awareness and ethical reflection. While necessary, these approaches are insufficient in the absence of mechanisms that allow researchers to act on their concerns. By introducing the concept of the dual-use agency gap and identifying epistemic distance as its structural cause, this paper reframes militarization as a problem of institutional design rather than individual intent.

Addressing this gap requires careful attention to how research infrastructures allocate control and responsibility. Modest, community-level interventions aimed at preserving researcher agency can contribute to peace-oriented AI research grounded in harm prevention, professional responsibility, and respect for international law.

## GENERATIVE AI USAGE DISCLOSURE

The authors used large language models as a limited assistive tool only during the editing process. Specifically, AI assistance was used to support language refinement, structural clarity, and LaTeX formatting. All substantive arguments, conceptual framing, literature selection, interpretation, and claims are the sole responsibility of the authors. No AI system was used to generate original research ideas, empirical findings, or references. All cited sources were selected, verified, and interpreted by the authors in accordance with ICLR 2026 policies on large language model usage.

## ETHICS STATEMENT

This paper presents a conceptual and critical analysis of AI research paradigms and their relationship to peace, militarization, and governance. It does not involve human subjects, personal data, or experimental deployment of AI systems. As such, no institutional ethics approval was required.

The work is motivated by concerns about the downstream use of general-purpose AI research in military, security, and surveillance contexts. The analysis is intended to support harm prevention, responsible research governance, and peace-oriented outcomes. The paper does not promote, optimize, or enable military or weapons-related AI systems, and it avoids operational or tactical detail that could be repurposed for such uses.

## REFERENCES

- Alexander Brenneis. Assessing dual-use risks in ai research: Necessity, challenges and mitigation strategies. *Research Ethics*, 21(2):302–330, 2025.
- Kate Crawford. *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press, New Haven, CT, 2021.
- James Elkins Dobson. *The Birth of Computer Vision*. University of Minnesota Press, Minneapolis, MN, 2023.
- Sina Fereidooni and Vincent Heidt. The fallacy of precision: Deconstructing the narrative supporting ai-enhanced military weaponry. In *Harms and Risks of AI in the Military*. Springer, 2024.
- Future of Life Institute. Autonomous weapons open letter: Ai & robotics researchers. <https://futureoflife.org/open-letter/open-letter-autonomous-weapons-ai-robotics/>, 2016. Accessed January 2026.
- Sam Goodfriend. Algorithmic state violence: Automated surveillance and palestinian dispossession in hebron’s old city. *International Journal of Middle East Studies*, 55(3):461–478, 2023.
- Pratyusha Ria Kalluri, William Agnew, Myra Cheng, Katherine Owens, Luca Soldaini, and Abeba Birhane. The surveillance ai pipeline. *arXiv preprint arXiv:2309.15084*, 2023.
- Heidy Khlaaf, Sarah Myers West, and Meredith Whittaker. Mind the gap: Foundation models and the covert proliferation of military intelligence, surveillance, and targeting. *arXiv preprint arXiv:2410.14831*, 2024.
- Antony Loewenstein. *The Palestine Laboratory: How Israel Exports the Technology of Occupation around the World*. Verso Books, London, 2024.
- Sebastian Romansky. Lessons from the eu on confidence-building measures around artificial intelligence in the military domain. [https://www.sipri.org/sites/default/files/2025-05/eunpdc\\_no\\_97.pdf](https://www.sipri.org/sites/default/files/2025-05/eunpdc_no_97.pdf), 2025.
- Sebastian Schwartz, Leif Genrich Guntrum, and Christian Reuter. Vision or threat—awareness for dual-use in the development of autonomous driving. *IEEE Transactions on Technology and Society*, 3(3):163–174, 2022.