Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

# Military and Public Attitudes Toward Autonomous Weapons: Stage-Specific Drivers and Value-Based Moderation

**Ho-Suk Yang**
Ulsan National Institute of Science and Technology
**Eunseon Park**
Ulsan National Institute of Science and Technology
**Seunghoo Jeong**
Korea Railroad Research Institute
**Ji-Bum Chung**
learning@unist.ac.kr

Ulsan National Institute of Science and Technology

**Article**

**Keywords:** autonomous weapon systems, acceptance, artificial intelligence, knowledge, value

**Posted Date:** October 6th, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-10547531/v1

**License:** This work is licensed under a Creative Commons Attribution 4.0 International License. Read Full License

**Additional Declarations:** No competing interests reported.

## Abstract

This study examines acceptance of autonomous weapon systems (AWS) among military personnel and general citizens by distinguishing between support for their development and adoption versus their deployment and use. Using survey data from South Korea (general public, N = 1,222; military personnel, N = 155), we constructed hierarchical and interaction regression models to examine how demographics, value orientations, perceptions of just war and risk, and attitudes toward and knowledge of AI shape acceptance across different stages. Results show that military personnel report higher support for AWS development but substantially lower support for use than public. For development acceptance, the strongest drivers are interest in AWS and perception of just war compliance. For use acceptance, perception of just war compliance and acceptance of algorithm-based decision-making are the most influential predictors. Notably, knowledge of AI is negatively associated with acceptance of AWS use. Furthermore, this negative association is moderated by value orientations: individuals with stronger post-materialist priorities exhibit a steeper decline in support for AWS use with increasing AI knowledge, whereas materialist individuals demonstrate comparatively stable patterns. These results indicate that AWS acceptance is stage-dependent, conditioned by stakeholder responsibilities and value-based divisions. Consequently, governance discussions should differentiate between development and deployment/use stages, placing greater emphasis on operational risk and ethical legitimacy.

## Highlight

- Military samples support AWS development more than public but support using it less.
- AWS acceptance is stage-dependent; perception of just war compliance is important.
- Greater AI knowledge is negatively associated with AWS use acceptance.
- Post-materialists show steeper drop in AWS use support with rising AI knowledge.

## 1. Introduction

Technological change has repeatedly reshaped the conduct of warfare and the evolution of weapon systems, from early projectile weapons to gunpowder, industrial-era armaments, and twentieth-century weapons of mass destruction (Arkin, 2017; Turchin et al., 2021). Contemporary advances in artificial intelligence (AI) extend this trajectory into a new phase, most notably through the emergence of autonomous weapon systems (AWS) (Fleming, 2025; Lucas, 2018; Russell, 2023). AWS constitute an emerging military technology whose diffusion trajectory depends on technical capability, societal acceptance, and governance feasibility (Arai and Matsumoto, 2024; Harwood, 2024; Sagan and Valentino, 2019). Unlike previous weapon innovations, AWS integrate AI into lethal decision-making processes, fundamentally challenging established norms of human control, accountability, and ethical legitimacy.

International governance debates have expanded but remain fragmented. UN attention intensified after a 2010 report warned that advances in unmanned systems were outpacing legal and ethical frameworks (Alston, 2010), and a 2013 report on lethal autonomous robotics raised concerns about control and accountability, recommending a moratorium (Heyns, 2013). These issues were subsequently addressed within the Convention on Certain Conventional Weapons (CCW), which has debated definitions, meaningful human control, humanitarian law compliance, and responsibility; however, despite the 2019 Guiding Principles, progress toward binding regulation has been limited. Parallel UN General Assembly resolutions have further reinforced a multi-track, rather than unified, governance structure (Bruun & Bo, 2025; Harwood, 2024).

Likewise, civil society actors have shaped these debates by framing AWS as an ethical and legal threat and advocating preventive regulation or prohibition. The Campaign to “Stop Killer Robots,” launched in 2012, has become a prominent transnational initiative mobilizing public attention and policy pressure (C.S.K.R., 2013). Within this landscape, “meaningful

human control" (MHC), promoted by groups such as Article 36 and referenced in several international discussions, has emerged as a central governance concept.

Meanwhile, the industrial and technological environment surrounding AWS has shifted. Until the late 2010s, major technology firms tended to adopt AI principles restricting weaponization, largely to maintain public trust and respond to regulatory pressure (de Laat, 2021; Jobin et al., 2019). By the mid-2020s, however, these prohibitions have increasingly been relaxed or reinterpreted, and corporate involvement has expanded in enabling domains such as surveillance, target identification, threat detection, and decision-support functions (DiBenedetto, 2024; Elias, 2025). This trend underscores that AWS governance cannot be understood only as an intergovernmental legal problem; it is also shaped by corporate policy, defense innovation ecosystems, and the diffusion of dual-use capabilities.

Against this backdrop, AWS generates acute tensions between military utility expectations and ethical, legal, and societal risk concerns (Arkin, 2017; Harwood, 2024; Sparrow, 2007). Proponents of AWS emphasize potential operational advantages: enhanced response speed, improved battlefield situational awareness, and mitigation of human limitations such as fatigue, emotion, and cognitive bias (Guo, 2025; Riesen, 2022). Some analysts further contend that algorithmic systems may strengthen compliance with international humanitarian law through precision targeting and systematic decision-support (Woodcock, 2024). However, critics counter that AWS deployment introduces substantial governance challenges—algorithmic bias, data errors, unpredictable system behavior in complex environments, and accountability gaps when machines execute lethal decisions (Sagan and Valentino, 2019; Trusilo, 2023).

Under these conditions, development and deployment of AWS depends critically on acceptance among key stakeholder groups. Existing research shows that AWS evoke strong intuitive opposition due to concerns about dignity, moral agency, and accountability (Horowitz, 2016; Rosendorf et al., 2024; Verdiesen et al., 2019). Yet attitudes are context-dependent: acceptance shifts with information about utility, risk, and mission context (Arkin and Moshkina, 2007; Horowitz, 2016). Despite these insights, the empirical literature remains fragmented. Many studies do not distinguish between support for development and support for operational use, nor do they systematically compare civilian publics and military personnel—potential users who bear direct responsibility and risk (Arai and Matsumoto, 2024).

Thus, this study addresses these gaps. First, it conceptualizes AWS acceptance as a multi-stage process, distinguishing between (1) development and adoption, (2) deployment and operational use, and (3) mission-specific contexts. Second, it compares acceptance determinants across civilians and active-duty military personnel, examining whether role-based responsibilities produce divergent evaluation logics. Third, it integrates ethical legitimacy perceptions (grounded in *jus in bello* principles), AI knowledge, and value orientations within a unified analytical framework, testing whether value priorities moderate acceptance under high-salience conditions. Drawing on survey data from South Korea, we estimate hierarchical regression models to examine how sociodemographic, value priorities, perceptions of just war execution and risk, and AI attitudes and knowledge shape acceptance across diffusion stages. The study addresses three research questions:

RQ1: How do acceptances of AWS differ between the general public and military personnel across development and use stages?

RQ2: How do ethical legitimacy perceptions (*jus in bello* principle) and value orientations shape AWS acceptance?

RQ3: Does knowledge of AI increase or decrease acceptance of AWS? What role do value orientations play in this relationship?

By demonstrating that AWS acceptance is stage-dependent, stakeholder-contested, and conditioned by ethical legitimacy and value orientations, this study contributes to emerging technology governance research by identifying factors that shape the political feasibility of regulation, deployment, and diffusion.

The remainder of the article proceeds as follows. Section 2 reviews prior research on AWS acceptance and develops the analytical framework. Section 3 describes research methods including survey design and questionnaires. Section 4 reports results, including stakeholder differences, stage-specific determinants, and value-based moderation patterns. Finally, this study discusses implications for emerging technology governance and concludes with limitations and directions for future research.

# 2. Literature review

## 2.1. Definition of AWS

Debates on AWS fundamentally stem from the difficulty of defining what constitutes “autonomy” in the context of military force and identifying the threshold at which weapon systems become qualitatively distinct from existing automated technologies (Guo, 2025; Taddeo and Blanchard, 2023; Verdiesen et al., 2019; Wood, 2023). Despite growing international attention, no standardized definition has yet been consolidated across international organizations, governments, or the scholarly community due to the intersection of competing geopolitical interests among nation-states and substantial commercial stakes in the defense industry (CCW, 2025; Harwood, 2024; Taddeo and Blanchard, 2023). Notwithstanding this definitional fragmentation, the UN Convention on Certain Conventional Weapons (CCW), the International Committee of the Red Cross (ICRC), and the US Department of Defense (DoD) converge on the core feature that AWS “can select and engage targets without human intervention” (CCW, 2025; Harwood, 2024; International Committee of the Red Cross, 2025; Wood, 2023). While numerous issues surrounding AWS exist, the most critical concerns center on three interconnected dimensions that distinguish AWS from conventional weapons: the definition of autonomy, adaptive capabilities, and the requirement for meaningful human control (Guo, 2025; Taddeo and Blanchard, 2023; Wood, 2023).

### 2.1.1. Automation and autonomy

Building on the computer ethics tradition emphasizing that conceptual gaps lead to regulatory gaps (Moor, 1985), Taddeo and Blanchard (2023) provide a crucial distinction between automation and autonomy based on purposiveness. While automatic actors lack goal-directedness and act as causal entities, such as landmines triggered by specific mechanical inputs, autonomous weapons function as teleological agents capable of carrying out tasks to achieve objectives (Taddeo and Blanchard, 2023). These systems can adjust their actions based on environmental feedback (automated artificial agents), formulate plans to accomplish goals (heuristic artificial agents), and possess the capacity to refine their behavior in response to changing circumstances (adaptive artificial agents). This purposiveness fundamentally differentiates autonomous systems from their automated predecessors (Taddeo and Blanchard, 2023; Tamburrini and Amoroso, 2020; Wood, 2023).

From the perspective of the OODA loop (Observe Orient Decide Act, Figure 1), autonomy implies that functions traditionally performed by human operators, ranging from target detection to engagement execution, may be increasingly delegated to algorithmic processes (Boyd, 1996). As engagement tempo exceeds human cognitive and deliberative capacities, human intervention may be constrained to *ex ante* programming or *ex post* oversight, effectively creating operational conditions characterized as human out of the loop.

Figure 1. Human-out-of-the-loop

* Source: Boyd (1996). Adapted by the authors.

### 2.1.2. Adaptive capabilities

A significant gap in the official definitions provided by the ICRC and the US Department of Defense is the absence of explicit reference to “adaptive capabilities” (Bode et al., 2023; Guo, 2025; Taddeo and Blanchard, 2023). The capacity of an artificial agent to modify its internal states without direct intervention from external actors establishes a defining criterion

for distinguishing between automated systems and autonomous systems (Taddeo and Blanchard, 2023; Wood, 2023). Adaptability has become an increasingly prevalent characteristic in contemporary AWS, underpinning their potential to handle complex and rapidly evolving operational scenarios (Guo, 2025). This character simultaneously generates unpredictability, opacity in decision-making processes, diminished controllability, and accountability gaps associated with the deployment of such agents. Consequently, the inclusion of adaptive capabilities in AWS definitions is essential, accompanied by regulatory frameworks that provide adequate technical specificity regarding these capabilities (Taddeo and Blanchard, 2023).

#### 2.1.3. Meaningful human control

The precise attribution of responsibility for AWS constitutes a critical determinant of both AWS legitimacy and compliance with just war principles (Himmelreich, 2019; Sparrow, 2007). A broad coalition of scholars and members of the international community argue that all weapon systems, including AWS, must be maintained under meaningful human control (MHC) to ensure compliance with the laws of war, enable proper attribution of responsibility, protect human dignity of potential victims, and mitigate risks to international stability (Amoroso and Tamburrini, 2021; CCW, 2025; International Committee of the Red Cross, 2025; Tamburrini and Amoroso, 2020).

Frameworks for MHC can be analyzed through three distinct approaches: uniform, differentiated, and prudential frameworks. Given the manifest limitations of uniform solutions, effective MHC frameworks must be both differentiated and principle-based, grounded in prudential considerations (Tamburrini and Amoroso, 2020). In managing, developing, and deploying AWS, humans must fulfill three essential roles: (1) serving as a “fail-safe actor” who prevents weapon operations from resulting in indiscriminate attacks that violate international humanitarian law (IHL); (2) functioning as an “accountability attractor” who establishes the legal conditions necessary for attribution of responsibility under international criminal law (ICL); and (3) acting as a “moral agency enactor” who ensures that decisions affecting the lives, physical integrity, and property of individuals involved in armed conflicts are made exclusively by moral agents, thereby mitigating human dignity concerns associated with autonomous targeting decisions (Amoroso and Tamburrini, 2021; Horowitz, 2016; Tamburrini and Amoroso, 2020).

Synthesizing these discussions, this study defines AWS as weapon technologies designed to detect, identify, and classify targets, and to determine whether and how to engage them, in a manner that can proceed after activation without continuous and direct human intervention (CCW, 2025). Although fully autonomous weapons have not yet reached widespread routine deployment, concrete cases such as the Turkish made Kargu 2 loitering munition have illustrated the growing feasibility of systems capable of coupling target selection and engagement without continuous operator input (Choudhury, 2021). Moreover, borderline AI enabled targeting recommendation systems, such as those reported in recent operations in Gaza, highlight how decision support technologies may generate automation bias that weakens meaningful human judgment even when humans remain formally in the loop (Abraham, 2024; Blanchard & Bruun, 2025).

-

### 2.2. Just war theory

AWS raise not only technological and governance challenges but also questions of ethical legitimacy, particularly regarding whether lethal decision-making can remain compatible with fundamental norms of armed conflict (Sparrow, 2007). To capture this dimension, this study draws on just war theory as a normative framework for evaluating perceptions of the acceptability of AWS in warfare (Himmelreich, 2019; Sparrow, 2007; Walzer, 2015). In the context of AWS, the most relevant dimension is *jus in bello*, which governs the ethical conduct of warfare through three core requirements: (1) distinction—the obligation to discriminate between combatants and civilians, ensuring that only legitimate military targets are engaged; (2) proportionality—the requirement that anticipated military advantage outweighs expected civilian harm; and (3) precaution—the duty to take all feasible measures to minimize incidental civilian casualties and damage (Sparrow, 2007; Taddeo and

Blanchard, 2023; Walzer, 2015). These principles establish the normative standards against which AWS deployment must be evaluated, as algorithmic targeting decisions must demonstrably satisfy each requirement to maintain operational legitimacy under international humanitarian law (Guo, 2025).

Each *jus in bello* principle faces specific challenges under AWS deployment. Algorithmic target classification may fail to achieve reliable distinction in complex environments where combatant-civilian boundaries blur (Guo, 2025; Sparrow, 2007). Proportionality judgments require contextual moral reasoning about anticipated harm and military advantage that current AI cannot replicate (Taddeo and Blanchard, 2023). Precaution demands adaptive human oversight that autonomous systems may lack when operating beyond real-time intervention (Horowitz, 2016). These tensions suggest that AWS acceptance depends not only on perceived operational effectiveness but critically on stakeholder beliefs about whether autonomous systems can satisfy fundamental ethical requirements without delegating irreducible moral decisions to algorithms (Sagan and Valentino, 2019; Verdiesen et al., 2019).

Accordingly, this study examines whether perceptions of *jus in bello* compliance, particularly distinction and proportionality, predict acceptance patterns across development and deployment stages.

### 2.3. Acceptance of AWS

The development and deployment of AWS represent not merely technological advancement but strategic and institutional transformation in the conduct of warfare (Amoroso and Tamburrini, 2021; Arai and Matsumoto, 2024; Sparrow, 2007). As emerging military AI technologies become increasingly embedded in national security agendas and defense innovation ecosystems, decisions surrounding AWS extend beyond technical adoption and constitute broader questions of governance, legitimacy, and societal feasibility (Altmann & Sauer, 2017; Horowitz, 2018). Technological capability alone is insufficient to determine AWS diffusion trajectories; “acceptance” among key stakeholder groups fundamentally conditions regulatory approval, deployment timing, and political legitimacy (Amoroso & Tamburrini, 2021; Bode & Huelss, 2018).

In democratic contexts, the introduction of novel weapon technologies must therefore be evaluated not only through operational utility but also through public perceptions of moral acceptability, accountability, and human control (Horowitz, 2016; Rosendorf et al., 2024). Furthermore, acceptance must also be understood as a multi-level phenomenon that includes organizational acceptance within the armed forces. Military personnel represent primary users whose operational practices, responsibility structures, and ethical judgments may be fundamentally reshaped by autonomous systems. Without sufficient trust and willingness to employ AWS in practice, the expected effectiveness of such technologies may remain unrealized, regardless of technical performance (Macdonald & Schneider, 2019). Research on drones and AI-enabled battlefield systems suggests that soldiers often prefer manned platforms in high-risk environments, reflecting persistent psychological and normative barriers that can produce operational friction regardless of technical performance (Macdonald and Schneider, 2019; Verdiesen et al., 2019).

Studies consistently find that the prospect of “machines making life-and-death decisions” triggers strong intuitive resistance, regardless of operational performance or situational considerations (Arai & Matsumoto, 2024; Horowitz, 2016; Ikkatai et al., 2022; Rosendorf et al., 2024; Verdiesen et al., 2019). This opposition is often grounded in normative concerns about human dignity, moral agency, and the perceived ontological distinction between humans and algorithmic systems (Asaro, 2012; Matsumoto & Arai, 2025; Sparrow, 2016). Large-scale surveys similarly indicate widespread rejection of delegating lethal authority to machines (Ipsos, 2021).

However, recent research suggests that such opposition is not fixed but context dependent. When respondents are informed that AWS may exhibit lower error rates than human soldiers, dignity-based objections become less salient (Verdiesen et al., 2019), and support increases when comparative risk information highlights harm reduction or improved reliability (Rosendorf et al., 2024). By contrast, appeals to command responsibility alone appear insufficient to increase support, suggesting that legitimacy concerns extend beyond formal accountability arrangements. More broadly, surveys

that measure attitudes in abstract terms may overestimate opposition shaped by cultural imaginaries rather than realistic operational scenarios (Horowitz, 2016). When scenarios emphasize troop protection, military utility, or competitive pressures from adversary development, support for AWS development tends to rise, underscoring the conditional and strategic nature of acceptance judgments.

Beyond situational framing, acceptance varies systematically across individual-level characteristics and attitudes toward AI. Gender and age differences are consistently observed, with women and older respondents often expressing stronger concern, while men tend to be relatively more accepting (Arai & Matsumoto, 2024; Ikkatai et al., 2022). Greater AI knowledge does not necessarily translate into technological optimism; instead, higher familiarity may heighten sensitivity to ethical risks and uncertainty, producing more skeptical attitudes, whereas general interest in science and technology may correlate with more positive evaluations (Arai & Matsumoto, 2024; Ikkatai et al., 2022). These findings highlight the importance of distinguishing between technological literacy and technological enthusiasm as separate explanatory dimensions.

Acceptance further differs by stakeholders' proximity to military decision-making and by mission context. Military-experienced respondents and policymakers tend to evaluate robotic deployment as more acceptable than the general public, suggesting that attitudes are shaped by role-based identities and experiential familiarity with practices of force (Moshkina & Arkin, 2007). Moreover, support varies across mission profiles: acceptance is lower in scenarios involving direct coercion or likely civilian contact, but higher for non-offensive missions such as reconnaissance and surveillance (Moshkina & Arkin, 2007). Autonomy as a technical attribute alone does not determine acceptance; rather, its normative implications shift across operational contexts.

Despite these insights, three critical gaps remain. First, most studies treat AWS acceptance as monolithic, failing to distinguish between support for development versus deployment and use—stages that invoke distinct risk perceptions and ethical considerations (Arai and Matsumoto, 2024; Verdiesen et al., 2019). Second, systematic comparisons between military personnel and civilian publics remain rare, despite their divergent stakes, responsibilities, and evaluative criteria (Arai and Matsumoto, 2024). Third, existing research has begun to explore ethical legitimacy, AI knowledge, and value orientations as factors influencing attitudes toward AWS; however, these elements are typically examined in isolation, leaving the underlying mechanisms through which they interact to shape acceptance insufficiently theorized and empirically underexplored. Addressing these gaps, this study examines how AWS acceptance varies across development and use stages, comparing military personnel and civilians, while integrating ethical legitimacy, AI knowledge, and value priorities within a unified analytical framework. Using survey data from South Korea, we test whether perceptions of *jus in bello* compliance—particularly distinction and proportionality—predict acceptance, whether AI knowledge increases or decreases support (and for whom), and how value orientations moderate these relationships.

## 3. Method

## 3.1. Survey

This study employs a cross-sectional survey design to examine stakeholder acceptance of AWS, comparing attitudes between the general public and active-duty military personnel in South Korea. The public sample was collected through the online survey panel of Korea Gallup, a leading public opinion research organization. Data collection was conducted between December 9 and 18. A stratified random sampling approach was applied to enhance representativeness across key demographic and regional categories, including gender, age, and place of residence. Given South Korea's population concentration in metropolitan areas, regional quotas were allocated first, and a square-root proportional allocation method was used to balance sampling error across strata.

Given the practical limitations in accessing a comprehensive sampling frame for active-duty military personnel, probability-based recruitment was not feasible for this subgroup. Instead, a convenience sampling approach was employed,

leveraging professional networks to recruit voluntary participants. Data collection for this group took place between December 1 and 12, resulting in 150 completed responses. In total, responses were collected from 1,227 members of the general public and 150 active-duty military personnel. After excluding five respondents in the public sample who self-identified as active-duty personnel, the final dataset consisted of 1,222 civilian respondents and 155 military respondents.

Participants were explicitly informed of the anonymity of their responses, the absence of correct or expected answers, and their right to withdraw at any time without consequence, to mitigate potential social desirability bias within the hierarchical military organization. This study was approved by the Institutional Review Board of Ulsan National Institute of Science and Technology (UNIST) (Approval No. UNISTIRB-25-057-A) and was performed in accordance with the ethical standards and guidelines of the UNIST Institutional Review Board and the Declaration of Helsinki.

## 3.2. Survey questionnaires

The survey instrument was designed to capture acceptance of AWS in a multidimensional manner. Given that AWS represents an emerging military AI technology for which no single standardized acceptance model has yet been established, prior empirical studies have operationalized acceptance through heterogeneous and fragmented sets of explanatory factors (Arai and Matsumoto, 2024; Verdiesen et al., 2019). Building on this literature, the present study integrates key dimensions that have been treated separately in previous work and examines their combined effects within a unified analytical framework.

Acceptance of AWS was conceptualized as a multi-level outcome rather than a single undifferentiated attitude. Accordingly, the dependent variables were measured at three distinct levels: acceptance of development and adoption, acceptance of deployment and operational use, and mission-specific acceptance depending on the anticipated role of AWS. Acceptance of development was assessed using the item: “What is your overall opinion on the development of autonomous weapon systems?” Responses were measured on a five-point Likert scale (1 = strongly disagree, 5 = strongly agree). Acceptance of operational use was measured with the question: “How much do you approve of the use of fully autonomous weapons, where AI conducts the entire process from target selection to attack without human intervention?” Responses were measured on a five-point scale (1 = absolutely not allowed, 5 = strongly allowed). Recognizing that acceptance of AWS use may vary substantially across operational contexts, this study additionally measured mission-specific acceptance. Drawing on the task classifications proposed by Moshkina and Arkin (2007), and adapting them to the South Korean military’s mission-oriented weapon system framework, anticipated ground-operation roles were reconstructed. The survey focused on the ground domain in order to avoid excessive abstraction that might arise from including maritime, aerospace, or space-based missions. Respondents evaluated acceptance across roles such as surveillance and reconnaissance, maneuver, defense, attack, counter-terrorism, high-value target strikes, and cyber warfare. This approach allows for systematic assessment of how acceptance differs depending on mission profiles and normative implications.

Independent variables were grouped into four major dimensions: values orientations, perception of just war execution and associated risk, attitudes toward AI, and sociodemographic characteristics. Value orientations were operationalized through the Inglehart’s post-materialism index (Inglehart, 2020) and trust (in selected institutions and generalized trust). Post-materialist value priorities were measured following the standard approach developed by Inglehart and routinely implemented in the World Values Survey (https://www.worldvaluessurvey.org). Respondents were presented with a set of national priority items that combine materialist goals (e.g., economic stability, order, and fighting inflation) and post-materialist goals (e.g., participation, freedom of expression, and environmental protection). They were asked to select the most important and second-most important priorities, allowing classification along the materialist, mixed, and post-materialist continuum. In this study, we constructed a ‘materialist index’ by summing the number of materialist priorities selected across respondents’ choices (ranging from 0 to a maximum of 6). Trust was measured as both interpersonal trust and trust in the institutions responsible for AWS development and deployment (trust in the central government and trust in defense-related institutions), each assessed on five-point Likert scales (1 = do not trust at all, 5 = trust very much).

To capture the normative execution of AWS, perceptions of compliance with core *jus in bello* principles were measured through items assessing whether AWS would enhance distinction and proportionality, such as whether they could better distinguish combatants from civilians and reduce civilian casualties (average value of the questionnaires “AWS may improve the identification of combatants and civilians in warfare” and “The deployment of AWS has the potential to reduce civilian casualties in armed conflicts”). Furthermore, it was further assessed through respondents’ tolerance for system malfunction, measured on an 11-point scale ranging from 0% to 100% acceptable malfunction rates. Lower tolerance was interpreted as reflecting higher perceived risk and lower confidence in AWS reliability. In addition, perceived national security risk including threats posed by North Korea, neighboring military expansion, and terrorism were measured as contextual risk factors (average value).

Attitudes toward and knowledge of AI were operationalized through four sub-dimensions: subjective AI knowledge, objective AI knowledge, interest in AWS, and acceptance of algorithm-based decision-making. Objective knowledge was measured by asking respondents five general knowledge questions related to AI and analyzing their True/False/Don’t know responses to determine the number of correct answers (ranging from 0 to a maximum of 5). Examples of AI knowledge questions include “Artificial intelligence can be biased depending on training data” and “Artificial intelligence can learn patterns based on given data and perform tasks such as prediction or classification.” In contrast, subjective knowledge measured each respondent’s self-assessed level of AI knowledge (1 = don’t know at all, 5 = know very well). Interest in AWS was measured using a three-item scale adapted from prior studies (Hartwig et al., 2023; Ikkatai et al., 2022), with acceptable internal consistency (Cronbach’s $\alpha = 0.678$). The questionnaires are “I find the topic of AWS (AI weapons) itself interesting and engaging,” “I have searched for or shared with others news, articles, discussions, etc. related to AWS (AI weapons),” “I think AWS (AI weapons) are an important issue facing modern society” (1 = not at all, 5 = very much so). Acceptance of algorithmic decision-making was measured through an ethical dilemma scenario such as degree of consent for life-and-death decisions made through algorithms in accident situations adapted from the trolley-problem paradigm (Awad et al., 2018; Bonnefon et al., 2016; Christensen & Gomila, 2012; Greene & Haidt, 2002).

Sociodemographic variables included gender, age, active-duty military personnel status, and political affiliation. Military personnel were included as a key stakeholder group directly involved in the operational use of AWS, allowing systematic comparison between potential users and the broader public. Political affiliation was assessed via the question “Where would you place yourself in terms of political affiliation?” with responses recorded on a five-point scale (1 = very conservative, 5 = very progressive/liberal). Table 1 shows the demographic characteristics of respondents.

Table 1
The demographic characteristics of respondents

| | | Public | | Military Personnel | |
|---|---|---|---|---|---|
| | | **N** | **%** | **N** | **%** |
| Total | | 1,222 | 100 | 155 | 100 |
| Gender | Male | 617 | 50.5 | 148 | 95.5 |
| | Female | 605 | 49.5 | 7 | 4.5 |
| Age | 20s | 199 | 16.3 | 62 | 40.0 |
| | 30s | 226 | 18.5 | 88 | 56.8 |
| | 40s | 249 | 20.4 | 5 | 3.2 |
| | 50s | 285 | 23.3 | 0 | 0 |
| | 60s | 265 | 21.5 | 0 | 0 |
| Political affiliation | Very Conservative | 35 | 2.9 | 12 | 7.7 |
| | Conservative | 247 | 20.2 | 34 | 21.9 |
| | Moderate | 638 | 52.2 | 102 | 65.8 |
| | Liberal | 265 | 21.7 | 5 | 3.2 |
| | Very liberal | 37 | 3.0 | 2 | 1.3 |

## 4. Results

Overall, acceptance of AWS differed systematically across stakeholder groups and across stages of technological diffusion. Respondents expressed greater support for the development of AWS than for their deployment and operational use, indicating that acceptance declines when autonomy is framed as directly substituting for human lethal judgment.

# 4.1. Comparative Analysis of AWS

# 4.1.2. Comparison among other sectors

This study explored public acceptance of AI use across both high-risk domains and Korean public service sectors. The results indicate that perceptions of AI vary considerably depending on application context and risk framing. We chose high-risk domains of AI implementation, such as crime prediction and critical infrastructure management, as outlined in the recent EU AI Act. We compared the acceptance of AWS with these high-risk domains (Fig. 2). The results showed that acceptance of AWS (3.14) was the second lowest, following its application in interview and recruitment domains (2.94). Furthermore, when comparing the acceptance of AI services in public sectors with AWS acceptance (Fig. 3), the fire and emergency service sector (3.93) demonstrated the highest acceptance, while AWS showed a moderate level (3.40).

# 4.1.2. Comparison between public and military personnel

Significant differences emerged between the general public and military personnel. Military respondents expressed substantially higher support for AWS development and adoption than civilians ($\Delta M = 0.53$, $p < .001$). However, acceptance of deployment and operational use was lower among military personnel than among the public ($\Delta M = -0.28$, $p < .01$). This suggests that while military stakeholders may recognize strategic incentives for development, they remain comparatively

cautious regarding real-world operational use. Gender differences were also pronounced. Men reported significantly higher acceptance than women for both development and use. Figure 4 displays these differences.

To better contextualize these group differences, additional explanatory variables were compared. Military respondents reported significantly higher interest in AWS, consistent with their occupational proximity to the issue. They also perceived greater national security threats and evaluated the feasibility of just war compliance more positively. At the same time, military personnel exhibited lower tolerance for system malfunction, indicating heightened sensitivity to operational risk and the potential for civilian harm. Together, these findings highlight that acceptance among military stakeholders is shaped by a combination of strategic threat perception and risk-aware caution regarding battlefield reliability. Figure 5 shows these results of key variables.

## 4.1.3. Acceptance of Use by Mission

Acceptance of AWS deployment was strongly contingent on mission context. Military respondents expressed relatively higher support for non-offensive roles such as surveillance, reconnaissance, and information analysis, whereas acceptance declined sharply for missions involving direct lethal engagement, aggressive combat operations, or high-value target strikes (Fig. 6). This mission-dependent pattern underscores that acceptance is not determined by autonomy as a technical attribute alone, but by the normative implications embedded in specific operational use cases.

## 4.2 Regression model: Acceptance of AWS development and adoption

Hierarchical regression models were estimated to identify determinants of support for AWS development and adoption (Table 2). Demographics were entered first (Model 1), followed by values (Model 2), just war execution and associated risk perception (Model 3), and AI-related attitudes and knowledge (Model 4), with a full model reported in Model 5. In Model 1 (Adjusted $R^2$ = .095), men and military personnel expressed higher development support, whereas more progressive political affiliation was associated with lower acceptance. Adding values increased explanatory power modestly (Model 2; Adjusted $R^2$ = .134): materialist priorities and trust in defense institutions were positively associated with support.

A large improvement occurred when variables of just war execution and associated risk perception were introduced (Model 3; Adjusted $R^2$ = .240). Development support increased with perceived national security risk. Perceived execution with distinction and proportionality (Just War Compliance) also showed a robust positive association, whereas tolerance for malfunction was associated with reduced support. For AI-related attitudes and knowledge model (Model 4; Adjusted $R^2$ = 0.243), AWS development support showed relatively large positive associations with AWS interest, acceptance of algorithm-based decision-making, and objective AI knowledge. In the full model (Model 5; Adjusted $R^2$ = .316), a correlation was observed between AWS support and factors of AWS interest, just war compliance, algorithm acceptance and gender ranked in that order. On the other hand, the higher the tolerance for malfunction and the more progressive, the lower the support for AWS development tended to be. Table 2 summarizes these results.

Table 2
Hierarchical regression results of acceptance of AWS development and adoption (with standardized β coefficients)

| Variable | Model 1 (Demographics) | Model 2 (Demographics + Values) | Model 3 (Demographics + Just War/Risk Perceptions) | Model 4 (Demographics + AI Perceptions) | Model 5 (Full Model) | Model 6 (Include an interaction term) |
|---|---|---|---|---|---|---|
| **Demographics** | | | | | | |
| Gender (male = 1, female = 0) | 0.203*** | 0.192*** | 0.153*** | 0.101*** | 0.094*** | 0.094*** |
| Age | 0.076** | 0.037 | 0.023 | 0.027 | -0.014 | -0.012 |
| Military (military = 1, Public = 0) | 0.145*** | 0.117*** | 0.088** | 0.086** | 0.043 | 0.046 |
| Political affiliation | -0.122*** | -0.095*** | -0.073** | -0.111*** | -0.060* | -0.055* |
| **Values** | | | | | | |
| Interpersonal trust | | 0.011 | | | 0.016 | 0.016 |
| Trust in government | | 0.051 | | | 0.001 | 0.001 |
| Trust in national defense depart. | | 0.114*** | | | 0.060* | 0.057* |
| Materialism | | 0.152*** | | | 0.142*** | -0.003 |
| **Just War & Risk Perceptions** | | | | | | |
| Perceived national security threat | | | 0.115*** | | 0.034 | 0.034 |
| Tolerance for malfunction | | | -0.111*** | | -0.017*** | -0.121*** |
| Just War Compliance | | | 0.355*** | | 0.209*** | 0.208*** |
| **AI Perceptions** | | | | | | |
| Subjective AI knowledge | | | | -0.022 | -0.019 | -0.017 |
| Objective AI knowledge | | | | 0.069** | 0.046 | -0.103 |
| Interest in AWS | | | | 0.363*** | 0.271*** | 0.269*** |
| Note: * p < 0.05, ** p < 0.01, *** p < 0.001. | | | | | | |

| Variable | Model 1 (Demographics) | Model 2 (Demographics + Values) | Model 3 (Demographics + Just War/Risk Perceptions) | Model 4 (Demographics + AI Perceptions) | Model 5 (Full Model) | Model 6 (Include an interaction term) |
|---|---|---|---|---|---|---|
| **Demographics** | | | | | | |
| Algorithm acceptance | | | | 0.122*** | 0.098** | 0.098*** |
| **Interaction** | | | | | | 0.226* |
| Materialism  Objective AI knowledge | | | | | | |
| **Adjusted $R^2$** | 0.095 | 0.134 | 0.240 | 0.243 | 0.316 | 0.319 |
| **F-value** | 37.171 | 27.512 | 63.090 | 56.324 | 43.456 | 41.279 |
| Note: * p < 0.05, ** p < 0.01, *** p < 0.001. | | | | | | |

# 4.3 Regression model: Acceptance of AWS deployment and Use

In Model 1(Adjusted $R^2$ = 0.032), male reported higher acceptance of deployment and use, whereas military personnel expressed significantly lower support than the general public, indicating pronounced stakeholder divergence. Notably, the negative effect of military status remained strong and stable across all models. This suggests that proximity to operational responsibility and risk may generate greater caution toward real-world use.

A large increase in explanatory power occurred when variables of just war execution and associated risk perception were introduced (Model 3; Adjusted $R^2$ =0.287), indicating that these considerations constitute core drivers of deployment/use acceptance. In particular, perceived compliance with *jus in bello* requirements—especially distinction and proportionality exerted the strongest positive effect on acceptance and remained robust in the full model (Model 5: β = 0.303, p < .001). Willingness to tolerate malfunction was positively associated with deployment/use acceptance, indicating that respondents who are less sensitive to (or more willing to accept) the risk of civilian harm from technical failure are more supportive of operational use.

When AI-related attitudes and knowledge were added (Model 4; Adjusted $R^2$ =0.294), interest in AWS and acceptance of algorithm-based decision-making significantly increased support for deployment and use. By contrast, objective AI knowledge showed a consistent negative association with deployment/use acceptance and remained significant. This pattern suggests that greater technical understanding may heighten sensitivity to uncertainty, limitations, and ethical or risk concerns—producing informed caution or skepticism rather than technological optimism.

In the full model (Model 5; Adjusted $R^2$ =0.367), the most consistent predictors of deployment/use acceptance were perception of just war compliance, acceptance of algorithmic decision-making, interest in AWS, and malfunction tolerance. By comparison, objective AI knowledge and trust in government were negatively associated with use acceptance. Overall, these results reinforce stage-dependent acceptance dynamics: whereas development/adoption support is partly shaped by strategic, deployment and operational use are more strongly conditioned by ethical feasibility, civilian-harm risk, and stakeholder responsibility. These results are shown in Table 3.

Table 3
Hierarchical regression results of acceptance of AWS deployment and use (with standardized β coefficients)

| Variable | Model 1 (Demographics) | Model 2 (Demographics + Values) | Model 3 (Demographics + Just War/Risk Perceptions) | Model 4 (Demographics + AI Perceptions) | Model 5 (Full Model) | Model 6 (Include an interaction term) |
|---|---|---|---|---|---|---|
| **Demographics** | | | | | | |
| Gender | 0.153*** | 0.148*** | 0.108*** | 0.051* | 0.066** | 0.066** |
| Age | 0.012 | -0.006 | -0.041 | -0.024 | -0.050* | -0.049* |
| Military person | -0.128*** | -0.133*** | -0.163*** | -0.171*** | -0.185*** | -0.184*** |
| Political affiliation | -0.076** | -0.071* | -0.039 | -0.073** | -0.042 | -0.039 |
| **Values** | | | | | | |
| Interpersonal trust | | -0.034 | | | 0.018 | 0.018 |
| Trust in government | | -0.011 | | | -0.052* | -0.052* |
| Trust in national defense depart. | | 0.154*** | | | 0.048 | 0.047 |
| Materialism | | -0.052 | | | 0.010 | -0.063 |
| **Just War & Risk Perceptions** | | | | | | |
| Perceived national security threat | | | 0.017 | | -0.012 | -0.012 |
| Tolerance for malfunction | | | 0.178*** | | 0.099*** | 0.097*** |
| Just War Compliance | | | 0.461*** | | 0.303*** | 0.303*** |
| **AI Perceptions** | | | | | | |
| Subjective AI knowledge | | | | -0.006 | -0.021 | -0.020 |
| Objective AI knowledge | | | | -0.186*** | -0.132*** | -0.207** |
| Interest in AWS | | | | 0.307*** | 0.174*** | 0.173*** |
| Algorithm acceptance | | | | 0.280*** | 0.207*** | 0.207*** |
| Note: * p < 0.05 ** p < 0.01 *** p < 0.001 | | | | | | |

| Variable | Model 1 (Demographics) | Model 2 (Demographics + Values) | Model 3 (Demographics + Just War/Risk Perceptions) | Model 4 (Demographics + AI Perceptions) | Model 5 (Full Model) | Model 6 (Include an interaction term) |
|---|---|---|---|---|---|---|
| **Demographics** | | | | | | |
| **Interaction**<br><br>Materialism □<br>Objective AI knowledge | | | | | | 0.114 |
| Adjusted $R^2$ | 0.032 | 0.057 | 0.287 | 0.294 | 0.367 | 0.367 |
| F-value | 12.404 | 10.387 | 79.993 | 72.601 | 54.206 | 50.949 |
| Note: * p < 0.05 ** p < 0.01 *** p < 0.001 | | | | | | |

## 5. Discussion

This study examined acceptance of AWS by distinguishing between support for technological development and support for deployment and operational use, comparing military personnel and civilians in South Korea. The findings demonstrate that AWS acceptance cannot be treated as a uniform attitude but reflects a multidimensional governance judgment conditioned by ethical legitimacy perceptions, contextual risk, and stakeholder-specific responsibility structures.

## 5.1 Stage-Dependent Acceptance and the Development-Deployment Paradox

The most theoretically consequential finding is the systematic divergence between acceptance of AWS development and operational use—and the reversal of this pattern across stakeholder groups. Military personnel expressed significantly higher support for AWS development than the public, yet significantly lower support for battlefield deployment and use (RQ1). This inversion suggests that proximity to operational responsibility does not straightforwardly produce technological enthusiasm; rather, military stakeholders appear to apply qualitatively different evaluative criteria to abstract capability development versus concrete autonomous lethal engagement.

From a diffusion and governance perspective, this pattern implies that strategic necessity arguments—competitive pressure from neighboring states, capability enhancement imperatives—may generate sufficient legitimacy for development programs while failing to sustain acceptance at the point of operational use. Acceptance of use appears to function as a legitimacy gatekeeper in the AWS diffusion trajectory, meaning that governance constraints may emerge even where developmental momentum is strong (Bode et al., 2023; Tamburrini and Amoroso, 2020). This finding extends prior work by Arai and Matsumoto and Verdiesen et al. (Arai and Matsumoto, 2024; Verdiesen et al., 2019) by demonstrating empirically that these two stages activate distinct evaluative logics rather than representing a single underlying attitude. In other words, acceptance is shaped not by general technological preference, but by role-based exposure to the practical consequences of autonomy in warfare (Gao et al., 2025; Sparrow, 2007).

## 5.2 Ethical Legitimacy as the Central Determinant of Use Acceptance

For the RQ2, perceptions of ethical execution captured through confidence in compliance with *jus in bello* principles such as distinction and proportionality emerged as the strongest predictor of AWS acceptance, particularly at the deployment/use stage. Furthermore, individuals with stronger materialist value orientations were more likely to support the

development of AWS, but no clear relationship was observed with respect to their use. This indicates that societal and organizational judgments about AWS are fundamentally grounded in concerns about whether meaningful human control and civilian protection can be maintained under autonomous engagement conditions. Notably, the effect of perceived just war compliance was substantially stronger for acceptance of use than for acceptance of development. This suggests that ethical legitimacy concerns become more salient when respondents evaluate concrete battlefield employment rather than abstract technological advancement. Operational acceptance therefore functions as a legitimacy gatekeeper in the diffusion of AWS, highlighting that governance constraints may arise even when strategic incentives for development are present.

## 5.3 Knowledge, Interest, and Informed Skepticism Toward AWS

A further contribution of this study concerns the relationship between AI knowledge and acceptance. Objective AI knowledge exhibited a consistent negative association with acceptance of AWS use, challenging the conventional "knowledge deficit" assumption that opposition stems primarily from ignorance (RQ3). Instead, greater technical understanding may increase awareness of uncertainty, limitations, and ethical risk, thereby producing informed skepticism toward high-stakes autonomous applications. This finding is consistent with the results reported by Arai and Matsumoto (2024) and Ikkatai et al. (2022). This relationship often appears in other technological systems that involve potentially fatal risks. For example, Chung (2020)'s study on knowledge and acceptance of nuclear power similarly showed that as knowledge about nuclear power increased, acceptance of nuclear power plants tended to decrease. For AWS, this negative knowledge effect was most pronounced for operational use rather than development support, indicating that knowledge may foster optimism toward abstract innovation while simultaneously heightening critical evaluation when autonomy is linked to real-world lethal decision-making (Said et al., 2025; Tully et al., 2025). At the same time, interest in AWS remained a strong positive predictor across both stages, suggesting that cognitive engagement and technical literacy operate through distinct mechanisms in shaping acceptance.

## 5.4 Value Orientations as a Moderator of AWS Acceptance

An additional contribution of this study lies in identifying the moderating role of underlying value orientations in shaping acceptance of AWS. As illustrated in Fig. 7, the relationship between contextual conditions and acceptance differs systematically across materialist (materialist index = 5 or 6) and postmaterialist (materialist index = 0 or 1) groups. For respondents with stronger materialist priorities, acceptance of AWS development remains relatively stable or even increases under higher level of knowledge conditions, suggesting that strategic necessity and national security considerations may outweigh ethical reservations. By contrast, respondents with postmaterialist orientations exhibit a pronounced decline in acceptance, particularly at the operational use stage (RQ3), when their knowledge increase. This pattern indicates that when autonomy is linked to concrete battlefield deployment, ethical concerns regarding civilian protection, human dignity, and meaningful control become substantially more salient among postmaterialist groups.

These findings suggest that societal acceptance of emerging military AI technologies is not only driven by perceptions of risk and legitimacy but is also fundamentally conditioned by deeper value-based worldviews. In this sense, AWS diffusion represents a socio-technical governance challenge in which public support depends on how technological innovation interacts with competing normative priorities in society.

# 6. Conclusion

This study reconceptualized acceptance of AWS not as a single attitudinal outcome, but as a staged ethical evaluation that differentiates between technological development and operational deployment. Drawing on survey data from both the South Korean public and military personnel, the findings demonstrate that support for AWS shifts systematically across diffusion stages and is structured by ethical legitimacy judgments, technical knowledge, and value orientations.

Most notably, perceptions of compliance with core *jus in bello* principles emerge as the decisive determinant of operational acceptance. Respondents do not evaluate AWS primarily through abstract innovation frames; they assess whether delegating lethal decision-making to autonomous systems can remain normatively defensible. Ethical legitimacy therefore functions as a threshold condition for deployment: development may be tolerated under strategic or institutional logics, but operational use requires moral justification grounded in civilian protection and accountable control.

The divergence between military and civilian responses further underscores the ethical salience of responsibility. Although military personnel express relatively stronger support for technological development, they exhibit greater caution toward battlefield use. Proximity to operational accountability appears to heighten sensitivity to moral risk, suggesting that experiential awareness of consequences intensifies ethical scrutiny rather than technological endorsement.

The negative association between objective AI knowledge and operational acceptance challenges the knowledge-deficit model often invoked in technology governance. Greater technical literacy does not automatically generate confidence. Instead, it appears to sharpen awareness of algorithmic uncertainty, system limitations, and the moral hazards of delegating life-and-death decisions to probabilistic systems. Knowledge, in this context, produces reflective restraint.

Finally, value orientations moderate these dynamics. Post-materialist respondents—those prioritizing civil participation, freedom of expression, and environmental protection —exhibit pronounced declines in acceptance as knowledge increases. Acceptance of AWS thus unfolds within deeper normative conflicts over the moral boundaries of automation in warfare. Governance strategies that rely solely on performance claims or informational outreach will therefore remain insufficient unless they address these underlying ethical commitments.

Taken together, the findings suggest that the societal feasibility of AWS depends less on technological capability than on sustaining ethically credible conditions of control, accountability, and harm mitigation. Governance frameworks must explicitly articulate how meaningful human control will be operationalized, how responsibility gaps will be prevented, and how proportionality assessments will be preserved under algorithmic mediation. Without such safeguards, diffusion will encounter structurally differentiated resistance across both civilian and military constituencies.

This study is not without limitations. Constraints on probability sampling among military personnel and potential social desirability pressures may influence reported attitudes. Future research should incorporate qualitative inquiry, experimental designs, and cross-national comparisons to further unpack how institutional cultures and strategic contexts shape ethical evaluation. Nevertheless, by demonstrating that acceptance operates as a morally conditioned, stage-specific, and value-mediated construct, this study contributes to AI ethics scholarship by situating autonomous weapons not merely as technological artifacts, but as contested sites of normative governance.

# Declarations

### Competing Interests

The authors declare that they have no competing interests.

### Data Availability Statement

The survey instrument (Korean original with English translation) and codebook used in this study are available via the Open Science Framework at https://osf.io/zcvd2/?view_only=2e1f33ad5c8a412abe4bcd79c2410a33. Due to the confidentiality assurances provided to active-duty military respondents under IRB approval (UNISTIRB-25-057-A), individual-level raw response data are not publicly available. De-identified data and the analysis scripts used to generate the reported results are being finalized and will be added to the same repository; both may also be made available to qualified researchers upon reasonable request, subject to approval by the corresponding author's institution.

**Ethical Approval**

This study was approved by the Institutional Review Board of Ulsan National Institute of Science and Technology (UNIST) (Approval No. UNISTIRB-25-057-A). This study was performed in accordance with the ethical standards and guidelines of the UNIST Institutional Review Board and the Declaration of Helsinki.

**Informed Consent**

Participants were explicitly informed of the anonymity of their responses, the absence of correct or expected answers, and their right to withdraw at any time without consequence, to mitigate potential social desirability bias within the hierarchical military organization.

**Funding**

This work was supported by the National Research Foundation of Korea (NRF) [grant number RS-2025-00517091].

**Acknowledgement**

The present research was originally conducted as part of the first author's Master thesis (Risk Perception and Acceptance of Autonomous Weapon Systems: A Comparison Between Public and Military Personnel) at the Ulsan National Institute of Science and Technology.

# References

1. Abraham, Y. (2024). *'Lavender': The AI machine directing Israel's bombing spree in Gaza*. 972mag. Retrieved 2025 from https://www.972mag.com/lavender-ai-israeli-army-gaza/
2. Alston, P. (2010). *Interim Report of the Special Rapporteur on Extrajudicial, Summary or Arbitrary Executions* (A/65/321). U. Nations.
3. Altmann, J., & Sauer, F. (2017). Autonomous weapon systems and strategic stability. *Survival*, *59*(5), 117–142.
4. Amoroso, D., & Tamburrini, G. (2021). Toward a normative model of meaningful human control over weapons systems. *Ethics & International Affairs*, *35*(2), 245–272.
5. Arai, K., & Matsumoto, M. (2024). Public perceptions of autonomous lethal weapons systems. *AI and Ethics*, *4*(2), 451–462.
6. Arkin, R. (2017). Lethal autonomous systems and the plight of the non-combatant. In *The Political Economy of Robots: Prospects for Prosperity and Peace in the Automated 21st Century* (pp. 317–326). Springer.
7. Asaro, P. (2012). On banning autonomous weapon systems: human rights, automation, and the dehumanization of lethal decision-making. *International review of the Red Cross*, *94*(886), 687–709.
8. Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J.-F., & Rahwan, I. (2018). The moral machine experiment. *Nature*, *563*(7729), 59–64.
9. Blanchard, A., & Bruun, L. (2025). *Autonomous Weapon Systems and AI-enabled Decision Support Systems in Military Targeting: A Comparison and Recommended Policy Responses*.
10. Bode, I., & Huelss, H. (2018). Autonomous weapons systems and changing norms in international relations. *Review of International Studies*, *44*(3), 393–413.
11. Bonnefon, J.-F., Shariff, A., & Rahwan, I. (2016). The social dilemma of autonomous vehicles. *Science*, *352*(6293), 1573–1576.
12. Boyd, J. R. (1996). The essence of winning and losing. *Unpublished lecture notes*, *12*(23), 123–125.
13. Bruun, L., & Bo, M. (2025). *Bias in Military Artificial Intelligence and Compliance with International Humanitarian Law*. S. I. P. R. I. (SIPRI). https://doi.org/10.55163/NLWV5347

14. C.S.K.R. (2013). *Urgent Action Needed to Ban Fully Autonomous Weapons – Non-governmental organizations convene to launch Campaign to Stop Killer Robots*. Campaign to Stop Killer Robots. Retrieved 2025 from https://www.stopkillerrobots.org/news/press-release-urgent-action-needed-to-ban-fully-autonomous-weapons-non-governmental-organizations-convene-to-launch-campaign-to-stop-killer-robots/
15. Choudhury, M. (2021). Letter dated 8 March 2021 from the Panel of Experts on Libya established pursuant to resolution 1973 (2011) addressed to the President of the Security Council. In: UN Doc. No. S/2021/229.
16. Christensen, J. F., & Gomila, A. (2012). Moral dilemmas in cognitive neuroscience of moral decision-making: A principled review. *Neuroscience & Biobehavioral Reviews*, *36*(4), 1249–1264.
17. Chung, J.-B. (2020). Public deliberation on the national nuclear energy policy in Korea – Small successes but bigger challenges. *Energy Policy*, *145*, 111724. https://doi.org/https://doi.org/10.1016/j.enpol.2020.111724
18. de Laat, P. B. (2021). Companies committed to responsible AI: From principles towards implementation and regulation? *Philosophy & technology*, *34*(4), 1135–1193.
19. DiBenedetto, C. (2024). *OpenAI no longer bans military uses ChatGPT*. https://mashable.com/article/open-ai-no-longer-bans-military-uses-chatgpt
20. Elias, J. (2025, February 4, 2025). *Google removes pledge to not use AI for weapons, surveillance*. CNBC. Retrieved June 11 from https://www.cnbc.com/2025/02/04/google-removes-pledge-to-not-use-ai-for-weapons-surveillance.html
21. Fleming, N. (2025). Why we should limit the autonomy of AI-enabled weapons. *Nature*.
22. Gao, S., Wang, L., Zhang, N., Bai, Y., Tian, Y., Chen, S., Zhang, Y., & Zhou, H. (2025). Pilot flying versus pilot monitoring: The effect of role assignment on pilots' perceived risk in flight simulation experiment. *Risk analysis*.
23. Greene, J., & Haidt, J. (2002). How (and where) does moral judgment work? *Trends in cognitive sciences*, *6*(12), 517–523.
24. Hartwig, T., Ikkatai, Y., Takanashi, N., & Yokoyama, H. M. (2023). Artificial intelligence ELSI score for science and technology: a comparison between Japan and the US. *AI & SOCIETY*, *38*(4), 1609–1626.
25. Harwood, S. (2024). A cybersystemic view of autonomous weapon systems (AWS). *Technological Forecasting and Social Change*, *205*, 123514.
26. Heyns, C. (2013). *Report of the Special Rapporteur on extrajudicial, summary or arbitrary executions* (A/HRC/23/47). U. Nations.
27. Horowitz, M. C. (2016). Public opinion and the politics of the killer robots debate. *Research & Politics*, *3*(1), 2053168015627183.
28. Horowitz, M. C. (2018). Artificial intelligence, international competition, and the balance of power (May 2018).
29. Ikkatai, Y., Hartwig, T., Takanashi, N., & Yokoyama, H. M. (2022). Octagon measurement: Public attitudes toward AI ethics. *International Journal of Human–Computer Interaction*, *38*(17), 1589–1606.
30. Inglehart, R. (2020). *Modernization and postmodernization: Cultural, economic, and political change in 43 societies*. Princeton university press.
31. Ipsos. (2021). Global Survey Highlights Continued Opposition to Fully Autonomous Weapons.
32. Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, *1*(9), 389–399.
33. Lucas, G. (2018). The automated battlefield. In: American Association for the Advancement of Science.
34. Macdonald, J., & Schneider, J. (2019). Battlefield responses to new technologies: Views from the ground on unmanned aircraft. *Security Studies*, *28*(2), 216–249.
35. Matsumoto, M., & Arai, K. (2025). What's wrong with "Death by Algorithm"? Classifying dignity-based objections to LAWS. *AI & SOCIETY*, *40*(5), 3519–3530.
36. Moshkina, L., & Arkin, R. C. (2007). Lethality and autonomous systems: Survey design and results.

37. Rosendorf, O., Smetana, M., & Vranka, M. (2024). Algorithmic aversion? Experimental evidence on the elasticity of public attitudes to “killer robots”. *Security Studies*, *33*(1), 115–145.
38. Russell, S. (2023). AI weapons: Russia’s war in Ukraine shows why the world must enact a ban. *Nature*, *614*(7949), 620–623.
39. Said, N., Schumacher, L., & Huff, M. (2025). Artificial intelligence in medicine: the influence of medical expertise and perceived causability on medical AI risk and benefit perception. *Journal of Risk Research*, 1–15.
40. Sparrow, R. (2007). Killer robots. *Journal of applied philosophy*, *24*(1), 62–77.
41. Sparrow, R. (2016). Robots and respect: Assessing the case against autonomous weapon systems. *Ethics & International Affairs*, *30*(1), 93–116.
42. Tully, S. M., Longoni, C., & Appel, G. (2025). Lower artificial intelligence literacy predicts greater AI receptivity. *Journal of Marketing*, 00222429251314491.
43. Turchin, P., Hoyer, D., Korotayev, A., Kradin, N., Nefedov, S., Feinman, G., Levine, J., Reddish, J., Cioni, E., & Thorpe, C. (2021). Rise of the war machines: Charting the evolution of military technologies from the Neolithic to the Industrial Revolution. *Plos one*, *16*(10), e0258161.
44. Verdiesen, I., De Sio, F. S., & Dignum, V. (2019). Moral values related to autonomous weapon systems: an empirical survey that reveals common ground for the ethical debate. *IEEE Technology and Society Magazine*, *38*(4), 34–44.

## Figures

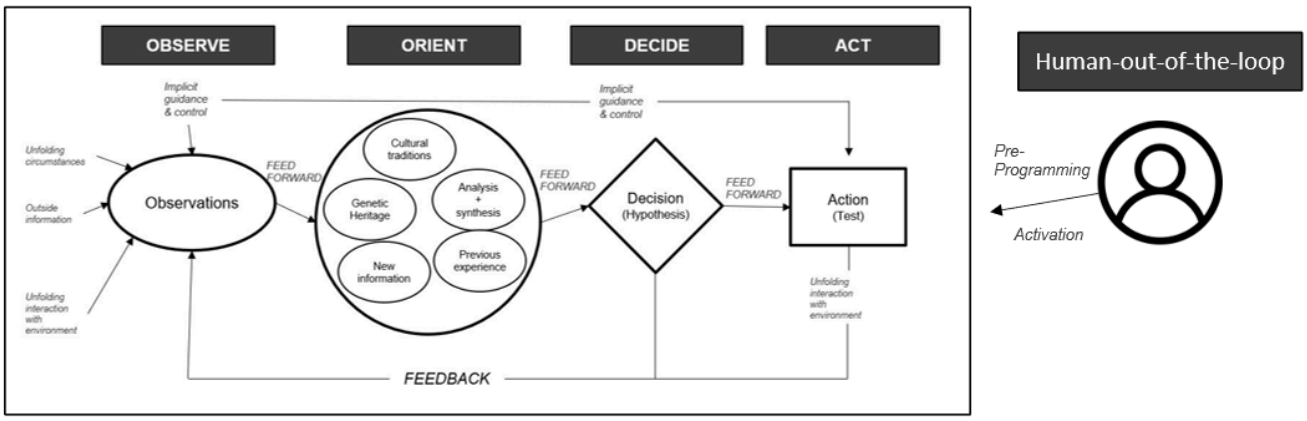



**Figure 1**

Human-out-of-the-loop

* Source: Boyd (1996). Adapted by the authors.

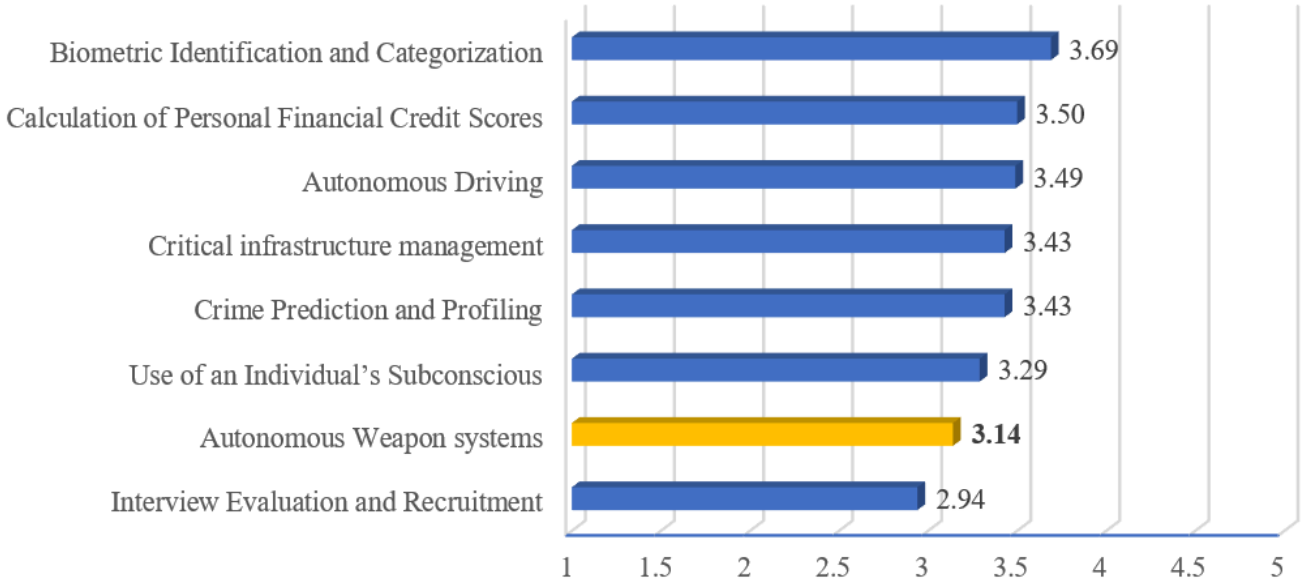



**Figure 2**

Mean Levels of Acceptance of High-Risk AI Application

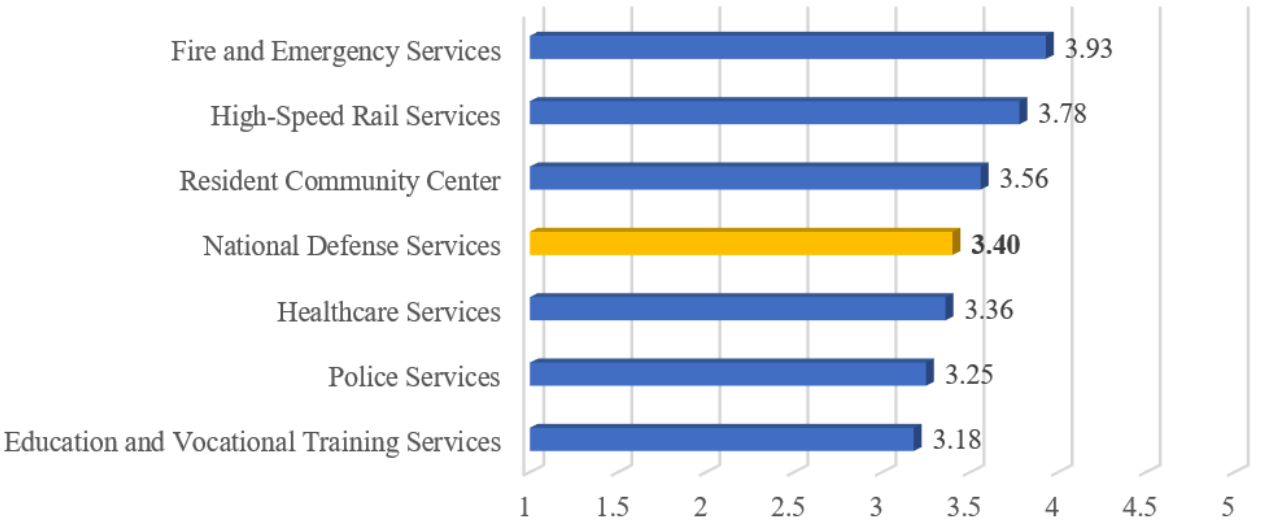



**Figure 3**

Mean Levels of Acceptance of Public Service Domains

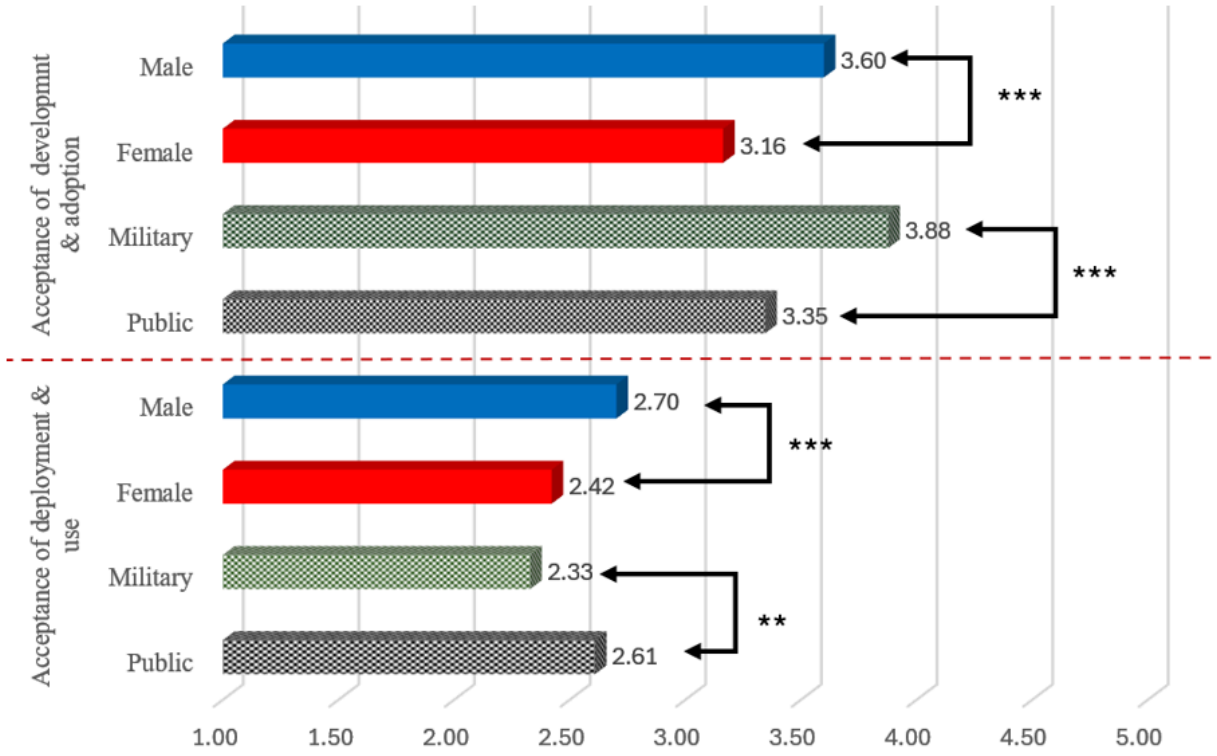



**Figure 4**

Group differences in acceptance of AWS (use, development)

Note: * p<0.05, ** p<0.01, *** p<0.001.

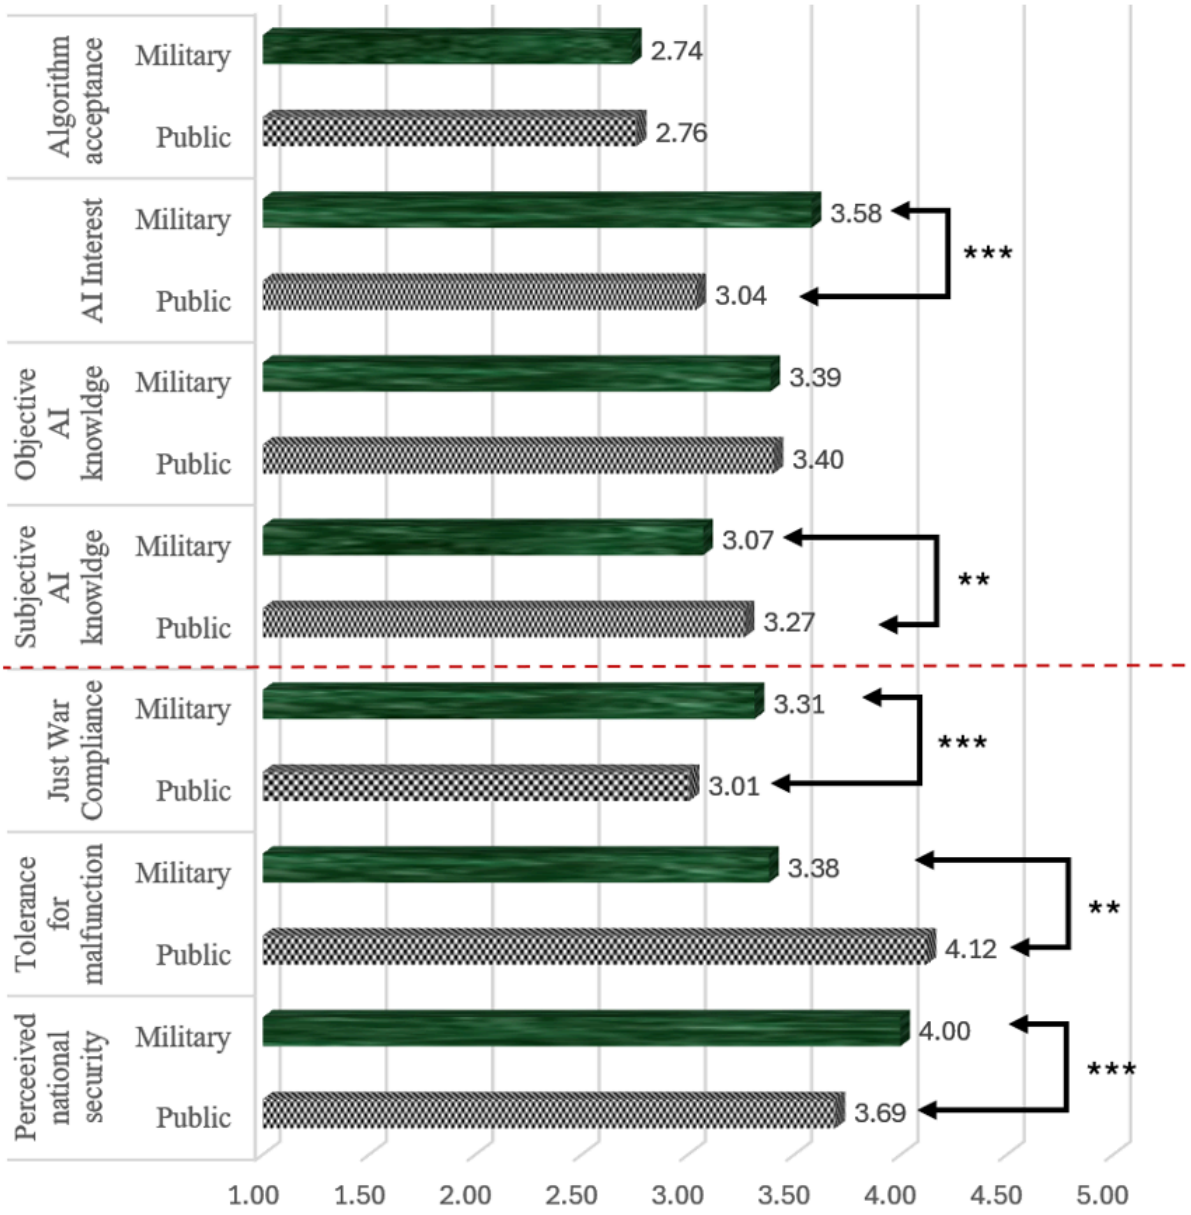



**Figure 5**

Group difference in the perception of AI and Just War principles

Note: * p<0.05, ** p<0.01, *** p<0.001.

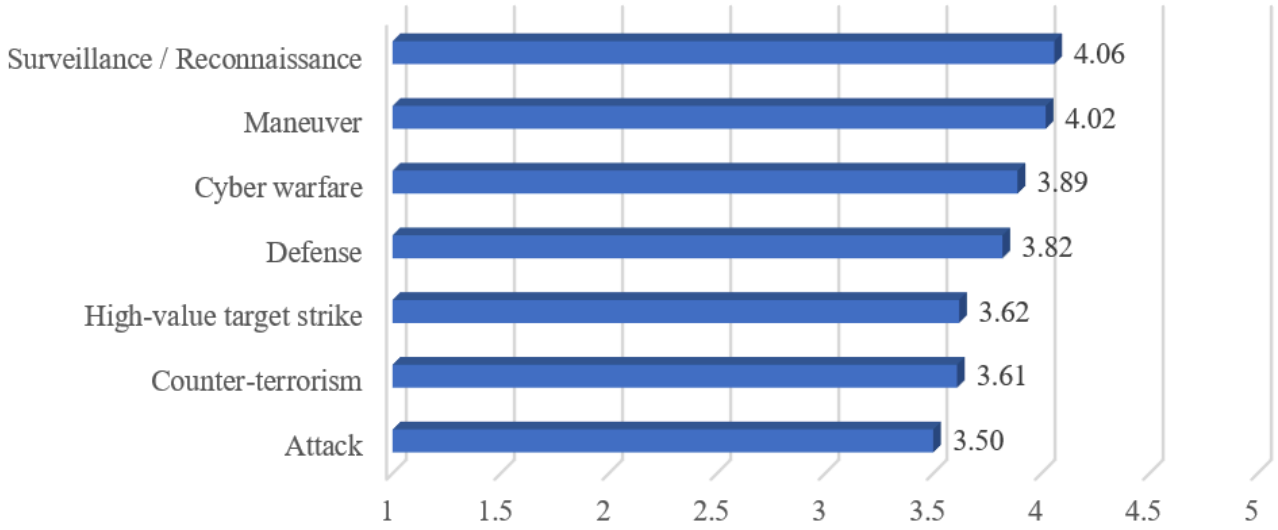



**Figure 6**

Military respondents' Acceptance of Use by Mission

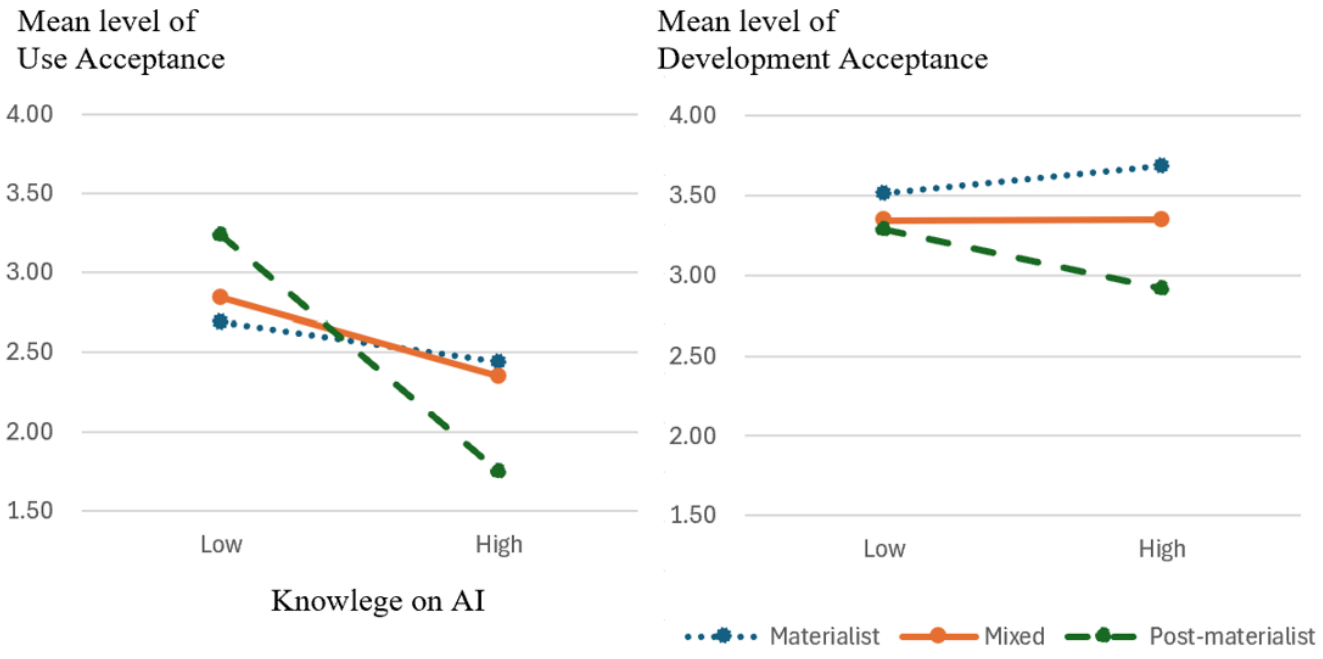



**Figure 7**

Interaction Effects of Value and AI knowledge