



UvA-DARE (Digital Academic Repository)

Techno-legal tinkering in war

AI decision-support systems and international humanitarian law

Klonowska, K.M.

Publication date

2026

[Link to publication](#)

Citation for published version (APA):

Klonowska, K. M. (2026). *Techno-legal tinkering in war: AI decision-support systems and international humanitarian law*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

TECHNO-LEGAL TINKERING IN WAR

AI Decision-Support Systems and International Humanitarian Law

Klaudia Klonowska

ISBN/EAN: 978-94-6536-060-7

The cover illustration was designed by Sebastian Camilo Rico Ruiz using generative AI tools.

This thesis was developed in the context of the DILEMA Project on Designing International Law and Ethics into Military Artificial Intelligence, funded by the Dutch Research Council Platform for Responsible Innovation (Nederlandse Organisatie voor Wetenschappelijk Onderzoek Maatschappelijk Verantwoord Innoveren (NWO-MVI)) under the project number MVI.19.017.

© 2026 Klaudia Klonowska. All rights reserved. The work may not be reproduced or distributed, in whole or in part, without explicit permission.

Techno-Legal Tinkering in War: AI Decision-Support Systems and
International Humanitarian Law

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor

aan de Universiteit van Amsterdam

op gezag van de Rector Magnificus

prof. dr. ir. P.P.C.C. Verbeek

ten overstaan van een door het College voor Promoties ingestelde commissie,

in het openbaar te verdedigen in de Agnietenkapel

op woensdag 18 maart 2026, te 10.00 uur

door Klaudia Maria Klonowska

geboren te Toruń

Promotiecommissie

<i>Promotores:</i>	prof. dr. M. C. Zwanenburg prof. dr. T. Marauhn	Universiteit van Amsterdam Universiteit van Amsterdam
<i>Copromotores:</i>	dr. B. L. Boutin	Universiteit van Amsterdam
<i>Overige leden:</i>	prof. dr. M. Kanetake prof. dr. M. de Goede dr. B.M.J. Pijpers dr. A.B. Leiter-Bockley dr. M. Arvidsson prof. dr. M.N. Schmitt prof. dr. T.M. van Engers	Universiteit van Amsterdam Universiteit van Amsterdam Universiteit van Amsterdam Universiteit van Amsterdam University of Gothenburg University of Reading Universiteit van Amsterdam

Faculteit der Rechtsgeleerdheid

Acknowledgements

The doctoral years were anything but a lonely journey. I have been fortunate to make numerous friends and to meet many inspiring individuals who have shaped my thinking, guided me at crucial junctures, and shared their professional and personal experiences. Many of them also offered me opportunities to present my work-in-progress at conferences, workshops, and departmental meetings, and provided their generous feedback. Without them, this manuscript would not have been the same!

I am first and foremost grateful for the unwavering support offered by my supervisors. Marten, thank you for your indispensable attention to detail in all your revisions, for introducing me to military jargon and ranks before conducting interviews, and for your continued support, even when my research temporarily took me in non-doctrinal directions. Thilo, your enthusiasm for this project from the very beginning gave me the necessary confidence to continue writing. I greatly enjoyed our exchanges, as you provided a big-picture perspective throughout the doctoral years; thank you. Berenice, thank you for setting the project up for success. You identified the research gap in AI decision-support systems, which has only grown in importance since then, and you recruited a spectacular research team that has enabled me to explore this topic beyond disciplinary boundaries. The influence of the research environment you curated is evident throughout this manuscript.

Taylor, you were my officemate, co-author, peer reviewer, tea companion, accountability partner, and the best friend a doctoral candidate could wish for! It was spectacular sharing this doctoral journey with you, from reading exchanges, PhD methodology workshops, to writing sessions, and adopting your habit of sending notes and articles to myself by email along the way. Though this manuscript might not be formatted in Times New Roman, I hope you consider it complete.

There are also those with whom I've worked closest on papers, conference presentations, and "side" (though not insignificant) projects. These collaborations shaped me as an academic. I want to thank Rebecca for the mentorship that has transformed my understanding of law, law-making, and their encounters with technology, which laid the groundwork for the research approaches on which this manuscript is based. Tasniem, although our paper suffered from longer-than-pregnancy publication cycles, I am very grateful for all the time we've spent together writing, refining, and presenting it. I've learned a great deal from you; thank you. Sofie, it is incredible how much we have been able to organize together over the years – from a workshop at EISA to a collective discussion piece with the wonderful Jasper, Ruben, and Natalie in the *International Political Sociology Journal*. My energy always doubles when brainstorming with you! Jonathan, I am grateful for all our fascinating discussions and for your generosity in sharing the depths of your knowledge with me. Each conversation with you resulted in a long list of follow-up ideas, which I save in the hope of future collaboration opportunities.

This journey would not have been possible without the initial support of Janne Nijman, a former academic director of Asser, who believed in my ability to pursue a doctoral degree. Likewise, Gert Griff, a managing director at Asser at the time, provided me with crucial institutional support to organize various PhD workshops, numerous writing boot camps, and, most importantly, the research stay at West Point. Thank you both for making this possible! Many others at Asser enriched my academic life. I am grateful to Sadjad (a member of “the originals”), Magda, and Tomasz – our postdoctoral researchers on the DILEMA project – who shared valuable insights, support, and laughs in those years. Sofia and Geoff, thank you for always finding the time, whether for a coffee or a long walk, to listen to my half-baked ideas and for offering repeated inspirations! I’ve learned more from you both than I can begin to tell. Geoff, I still have the first post-it note with your suggestions to read the works of Jasanoff, Haraway, and Suchman, which has essentially sent me into a large rabbit hole (thank you!). I also want to thank James, Marta, Thea, Romy, and Christophe for creating a supportive environment and for helping me navigate academic life on various occasions.

The US Military Academy at West Point became my second academic home. Special thanks go to Prof. Sean Watts, who arranged my research stay from March to May of 2024, and Prof. Hitoshi Nasu, who, despite my almost daily visits to his office with numerous questions, kept his doors open. All of the West Point community was uniquely welcoming, but I especially want to thank Col Winston Williams, Col Joshua Berry, Maj Stijn Henken, Jenny Maddocks, Cdr Prema MacIntosh, and Jose Malave. Discussions with you and your cadets were not only always joyous but also deepened my appreciation for the practical challenges many of you faced during military operations. I also want to thank Prof. Michael Schmitt, with whom I collaborated most during the last months of my PhD, for the doctrinal (re)disciplining, practical advice, and meetings that always begin and end with a smile.

I am fortunate to stand on the shoulders of many more friends and colleagues whom I crossed paths with at conferences and workshops. Although they may not realize it, their comments on earlier drafts of this manuscript inspired the final product. I must thank Wouter Werner and his colleagues at the VU, Lianne de Boer, Tanja Aalberts, Renske Vos, and Keri van Douwen, for the online Covid research sessions. I also want to thank Jessica Dorsey, Rupert Barrett-Taylor, Jasper van der Waa, Jeroen van den Boogard, Ingvild Bode, Marijn Hoijsink, Jutta Weber, Dustin Lewis, Fleur Johns, Emma Nyhan, Pauline Wernotte, Matilda Arvidsson, Peter Pijpers, Claudia Aradau, Dimitri van der Meersche, Aurel Sari, and Mitt Regan. Though academic institutions are often described as competitive environments, I am grateful to have encountered your collegiality, generosity, and kindness. I hope to pay it forward.

This list would not be complete without a heartfelt thanks to all my interviewees, who offered their time, shared their valuable experiences, and opened doors to their institutions. Special thanks to Michael Street (NATO NCIA), Michael Brasseur (formerly at Task Force 59), and colleagues at the Army Futures Command.

I am forever grateful for the support of my parents, who were often the first (sometimes even before me) to know when my publications appeared online and eagerly followed my LinkedIn updates. Finally, Camilo, mi amor, you not only supported the effort but took every step of this journey with me! You inspired me to embark on this intellectual path, persistently challenged my assumptions, and repeatedly got lost with me in the depths of the military and algorithmic worlds, even during our road trips. Your thoughts and pointed analysis are found in between every line of this manuscript. Thank you!

Executive Summary

This book explores the growing reliance on Artificial Intelligence (AI) decision-support systems in military operations and their profound impact on how targeting decisions are made during armed conflict. AI systems today analyze vast amounts of data – from satellite imagery to social media activity – to monitor emerging threats, detect and classify objects, identify and develop targets, and recommend courses of action. Their use in recent conflicts, such as in Gaza and Ukraine, illustrates how AI influences critical battlefield decisions, affecting who, what, how, and when is targeted in warfare.

Contrary to the popular promise that AI will make military decisions better and faster, this research reveals that integrating AI into military practice is slow, complex, and fraught with challenges. Through in-depth interviews with military practitioners from the United States and the Netherlands, this study demonstrates that the effective use of AI requires high-quality data, specialized training, and an ability to wrestle with new uncertainties.

This book introduces the concept of “techno-legal tinkering,” an ongoing, experimental process of modifying, re-training, and fine-tuning AI models that simultaneously shapes legal norms. The practice of tinkering allows the promises of gradual improvement of AI to obscure its inherent shortcomings, flaws, and inaccuracies. While aspiring to rectify issues in an ad-hoc manner, tinkering with AI decision-support systems concurrently introduces and normalizes uncertainties and risks associated with their use, and shifts the boundaries of what commanders consider reasonable in attack decisions.

The tinkering approach, as detailed in Chapter 2, arises from the difficulty of aligning AI training data with real-world battlefield conditions. The continuous work of developers is essential for maintaining AI system functionality, adapting them to evolving contexts of use, and mitigating risks of adversarial interference. As observed in the deployment of prototypical AI systems to contemporary battlefields, forward-deployed software engineers already actively adjust AI systems on the front lines to rectify performance fluctuations and sustain their use.

Despite the growing role of AI decision-support systems in armed conflicts, no binding treaty explicitly governs their development, acquisition, or use. Their lawfulness rests on existing international legal obligations, as addressed in Chapter 3. A doctrinal analysis of the core targeting rules of distinction, precautions, and proportionality in attack reveals that these systems have mixed effects. For example, statistical inaccuracies of AI systems can lead to target misidentification and increased civilian casualties. At the same time, an AI analysis of the indirect consequences of an attack can expand what is considered reasonable in a proportionality assessment. The impact of AI decision-support systems on compliance with the law is, therefore, highly contextual. This is why Chapter 3 calls for a deeper understanding of how AI reshapes decision-making and legal norms in practice than is possible through doctrinal approaches alone.

A crucial legal norm underlying the interpretation of open-textured targeting rules in international humanitarian law, especially under conditions of uncertainty, is

reasonableness. Chapter 4 critically examines attempts to embed this norm in AI systems and concludes that such processes do not lead to objective or reasonable decision outputs. Instead, AI systems introduce quantification logics and subjective choices that are built in by developers through data selection and model design. Rather than removing legal ambiguities, AI shapes the scope of what is accepted as reasonable military conduct, redistributing authority toward those who design and program these systems.

To further explore the realities of AI integration in military contexts, Chapter 5 focuses on the role of dashboards in shaping end users' perceptions of the battlefield. AI dashboards are portrayed as a "single pane of glass," promising a frictionless compression of diverse data streams into a comprehensive common operational picture. Using the example of United States Task Force 59, the chapter demonstrates the complexity hidden beneath promises of transparency. The single pane of glass draws a multitude of private actors and their systems-of-systems into military practices. This multiplicity contributes to the epistemological opacity of AI systems, making the ambition to "open the black box" of such interdependent, layered dashboards a formidable challenge. The chapter concludes that AI dashboards shape not only vision but also cognition, creating regimes of in/visibility that prefigure what merits attention and is acted upon.

The book also explores AI confidence scores – indicators that measure the likelihood of accuracy of AI outputs – and their effects on doubt in targeting decisions. Contrary to prevailing narratives, Chapter 6 suggests that AI integration often adds uncertainty and risk to military decision-making. Militaries struggle to implement and interpret confidence scores, frequently experiencing user frustration, and experiment with pre-setting confidence score thresholds to reduce system rejection. The chapter cautions that confidence scores are not directly translatable to legal concepts of doubt. Without an understanding of the system in its context of use, these scores cannot provide a reliable representation of decision-makers' degrees of doubt/certainty or the lawfulness of an attack decision.

Importantly, this book addresses how AI systems may constrain – rather than support – the exercise of deliberate judgement during attacks. Given the difficulty of contesting and verifying every AI output, users' perceptions of the battlefield and the dynamism of targeting operations are shaped by AI systems. At the same time, the uncertainty and context-dependence of AI performance must be acknowledged, and commanders should accordingly be expected to actively question and understand whether reliance on AI systems is reasonable in a specific situation. The notion of reasonable reliance on AI systems in military decision-making, as advanced in Chapter 7, emphasizes that managing the uncertainties of AI systems is a vital part of the contemporary responsibilities of commanders.

In sum, *"Techno-Legal Tinkering in War: AI Decision-Support Systems and International Humanitarian Law"* offers a nuanced and critical examination of AI's role in military targeting, emphasizing the entanglement of technology, law, and human judgment. It reveals how AI's subjective design choices and ongoing tinkering practices reshape legal

norms, operational procedures, and the boundaries of permissible military conduct. By bringing to light AI's inherent uncertainties and legal implications, the book calls for greater attention to the challenges and risks of integrating AI into high-stakes decisions in warfare.

This book will interest scholars, policymakers, and practitioners concerned with the impact of AI on warfare. It provides essential insights for those seeking to understand not only what AI does on the battlefield, but also how it changes the very nature of military decision-making and the interpretation of legal norms under international humanitarian law.

Nederlandse Samenvatting

Dit boek verkent de groeiende afhankelijkheid van besluitvormingsondersteuningssystemen die gebruik maken van kunstmatige intelligentie ('Artificial Intelligence (AI) *decision-support systems*') in militaire operaties en hun aanmerkelijke invloed op hoe beslissingen in het kader van doelbestrijding worden genomen tijdens gewapende conflicten. Tegenwoordig analyseren AI-systemen grote hoeveelheden data – van satellietbeelden tot sociale media-activiteit – om opkomende dreigingen te monitoren, objecten te detecteren en te classificeren, doelwitten te identificeren en ontwikkelen, en aanbevelingen te doen voor acties. Het gebruik ervan in recente conflicten, zoals in Gaza en Oekraïne, laat zien hoe AI beslissingen op het slagveld beïnvloedt, en daarmee mede bepaalt wie, wat, hoe en wanneer in een oorlog wordt aangrepen.

In tegenstelling tot de populaire belofte dat AI militaire beslissingen beter en sneller zal maken, toont dit onderzoek aan dat de integratie van AI in de militaire praktijk langzaam, complex en vol uitdagingen verloopt. Uit diepte-interviews met militaire professionals uit de Verenigde Staten en Nederland blijkt dat effectief gebruik van AI hoogwaardige gegevens, gespecialiseerde training en het vermogen vereist om om te gaan met nieuwe onzekerheden.

Dit boek introduceert het concept van '*techno-legal tinkering*', een doorlopend, experimenteel proces van het aanpassen, opnieuw trainen en verfijnen van AI-modellen. Tegelijkertijd vormt dit proces ook juridische normen. Deze praktijk van sleutelen ('*tinkering*') zorgt ervoor dat de belofte van geleidelijke verbetering de inherente tekortkomingen, fouten en onnauwkeurigheden van AI verhullen. Ondanks het streven naar ad-hoc oplossingen voor problemen, brengt het sleutelen aan AI *decision-support systems* nieuwe onzekerheden en risico's met zich mee die genormaliseerd worden. Ook verschuift daarmee wat commandanten als "redelijk" beschouwen in het kader van beslissingen over doelwitten en aanvallen.

De noodzaak voor deze tinkering-aanpak, zoals beschreven in hoofdstuk 2, vloeit voort uit de moeilijkheid om AI-trainingsdata goed af te stemmen op de realiteit van het slagveld. Het continue werk van ontwikkelaars wordt zo een essentieel onderdeel van het functionerend houden van AI-systemen, het aanpassen aan nieuwe gebruikssituaties en het beperken van risico's als gevolg van vijandige beïnvloeding. Zoals blijkt uit de inzet van prototype AI-systemen op hedendaagse slagvelden, zijn software-ingenieurs vaak in het veld aanwezig om AI-systemen actief te corrigeren wanneer de prestaties fluctueren en het gebruik ervan te borgen.

Ondanks de toenemende rol van AI *decision-support systems* in gewapende conflicten, bestaat er geen bindend verdrag dat hun ontwikkeling, aanschaf of gebruik expliciet reguleert. De rechtmatigheid van deze systemen berust op bestaande internationale juridische verplichtingen, zoals besproken in hoofdstuk 3. Een doctrinaire analyse van kernregels voor de selectie van doelwitten – onderscheid, voorzorg en proportionaliteit –

toont aan dat AI-systemen gemengde effecten hebben. Statistische onnauwkeurigheden kunnen bijvoorbeeld leiden tot een verkeerde identificatie van doelen en toename van burgerslachtoffers, terwijl een AI analyse van indirecte gevolgen van een aanval juist kan bijdragen aan een breder begrip van proportionaliteit. De impact van AI op de naleving van het recht hangt dus sterk af van de context. Hoofdstuk 3 roept op tot een dieper inzicht van hoe AI besluitvorming en normen in de praktijk hervormt dan louter op basis van doctrinaire benaderingen mogelijk is.

Een cruciale juridische norm die ten grondslag ligt aan de interpretatie van de open normen voor doelwitselectie in het internationaal humanitair recht, vooral in situaties van onzekerheid, is redelijkheid. Hoofdstuk 4 bevat een kritische analyse van pogingen om deze norm te verankeren in AI-systemen en concludeert dat dit niet leidt tot objectieve of redelijke beslissingen. In plaats daarvan introduceert AI kwantificeringslogica's en subjectieve keuzes die door ontwikkelaars worden ingebouwd via dataselectie en modelontwerp. In tegenstelling tot het wegnemen van juridische onzekerheid, bepaalt AI juist wat als redelijk wordt beschouwd in militaire gedragsnormen, waarbij de machtspositie verschuift naar degenen die deze systemen ontwerpen en programmeren.

Om de realiteit van AI-integratie in militaire contexten verder te verkennen, richt hoofdstuk 5 zich op de rol van 'dashboards' bij het vormen van percepties van eindgebruikers op het slagveld. AI-dashboards worden gepresenteerd als een 'single pane of glass', die diverse gegevensstromen zonder frictie samenbrengt tot een overzichtelijk gemeenschappelijk operationeel beeld. Aan de hand van het voorbeeld van de Amerikaanse Task Force 59 toont het hoofdstuk aan dat deze schijnbare transparantie gelaagd is: meerdere private partijen en hun complexe systemen zijn verweven in militaire processen. Die complexiteit draagt bij aan de epistemologische ondoorzichtigheid van AI-systemen, waardoor het openen van de 'black box' van zulke gelaagde dashboards een enorme uitdaging is. Het hoofdstuk concludeert dat AI-dashboards niet alleen het gezichtsveld, maar ook het cognitieve proces uitbreiden, en daarmee bepalen wat zichtbaar is, aandacht verdient en waarop gehandeld wordt.

Het boek gaat ook in op AI betrouwbaarheidsscores ('confidence scores') – indicatoren die de waarschijnlijkheid van de juistheid van AI-uitkomsten meten – en hun effect op twijfel in besluiten over doelwitten. In tegenstelling tot gangbare narratieven stelt hoofdstuk 6 dat AI-integratie vaak extra onzekerheid en risico's toevoegt aan militaire besluitvorming. Militairen worstelen met de implementatie en interpretatie van betrouwbaarheidsscores, ervaren frustratie bij het gebruik van de systemen en experimenteren met het instellen van drempelwaarden om afwijzing van het systeem te verminderen. Het hoofdstuk waarschuwt dat betrouwbaarheidsscores fundamenteel niet te vertalen zijn naar juridische concepten van twijfel. Zonder begrip van het systeem in de gebruikscontext bieden deze scores geen betrouwbare weergave van de mate van twijfel of zekerheid van besluitvormers of de rechtmatigheid van een aanvalsbeslissing.

Belangrijk is dat het boek onderzoekt hoe AI-systemen de mogelijkheid om weloverwogen beslissingen te maken tijdens aanvallen juist kan beperken in plaats van ondersteunen.

Vanwege de moeilijkheid om elke AI-uitkomst te weerleggen of te verifiëren, worden de waarnemingen van gebruikers op het slagveld en de dynamiek van operaties mede gevormd door AI. Tegelijkertijd moet de kwetsbaarheid en contextafhankelijkheid van AI-toepassingen erkend worden, en dienen commandanten actief te bevragen en te begrijpen of het redelijk is om in een specifieke situatie op AI-systemen te vertrouwen. Het idee van redelijke afhankelijkheid van AI in militaire besluitvorming, besproken in hoofdstuk 7, benadrukt dat het omgaan met AI-onzekerheden een essentieel onderdeel is van de hedendaagse verantwoordelijkheden van commandanten.

Samenvattend biedt *“Techno-Legal Tinkering in War: AI Decision-Support Systems and International Humanitarian Law”* een genuanceerde en kritische blik op de rol van AI in militaire besluitvorming, met nadruk op de verwevenheid van technologie, recht en het menselijk oordeelsvermogen. Het onthult hoe subjectieve ontwerpkeuzes en voortdurende aanpassingen de juridische normen, operationele procedures en grenzen van toegestaan militair handelen veranderen. Door de inherente onzekerheden en juridische gevolgen van AI bloot te leggen, roept het boek op tot meer aandacht voor de uitdagingen en risico's van AI in *high-stake* beslissingen in oorlogvoering.

Dit boek is gericht aan wetenschappers, beleidsmakers en praktijkbeoefenaars die geïnteresseerd zijn in de impact van AI op oorlogvoering. Het biedt essentiële inzichten voor iedereen die wil begrijpen wat AI precies doet op het slagveld, maar vooral ook hoe het de aard van militaire besluitvorming en juridische normen onder het internationaal humanitair recht verandert.

Table of Contents

Acknowledgements	5
Executive Summary.....	8
Nederlandse Samenvatting	11
List of Tables and Figures.....	17
List of Interviews	18
Abbreviations and Acronyms	19
CHAPTER 1: INTRODUCTION	21
1.1. <i>Scope: AI Decision-Support Systems and Military Targeting</i>	23
1.2. <i>Defining Techno-Legal Tinkering</i>	26
1.3. <i>Research Question</i>	28
1.4. <i>Theoretical Framework and Methodological Choices</i>	30
1.4.1. Posthumanism and the Fiction of the Human-In-The-Loop	30
1.4.2. Science and Technology Studies and Cognitive Assemblages	31
1.4.3. Critical AI Studies and the Task of Demystifying AI	33
1.5. <i>Methods and Study Limitations</i>	34
1.5.1. Doctrinal Analysis	34
1.5.2. Empirical Research Methods: Qualitative Expert Interviews	35
1.5.3. Reflections on Secrecy and "Black Boxes"	37
1.6. <i>Chapters Overview</i>	39
CHAPTER 2: AI DECISION-SUPPORT SYSTEMS: DEFINITION, MILITARY APPLICATIONS, AND UNIQUE CHALLENGES	43
2.1. <i>About AI: a Socio-Technical Apparatus</i>	44
2.2. <i>Military Applications of AI-enabled Systems</i>	50
2.2.1. Critical Versus Non-Critical AI Applications	51
2.2.2. Decision-Support Systems	52
2.2.3. AI Decision-Support Systems versus Autonomous Weapons Systems: Distinct or Related Issues?	55
2.3. <i>Unique Characteristics of AI Systems</i>	59
2.3.1. The Constant Work-in-Progress	59
2.3.2. Tinkering at the Front Lines	61
2.3.3. Fluctuations in AI Performance.....	64
2.3.4. Risk of Error Propagation	66
2.3.5. Opacity.....	67
2.4. <i>Chapter Conclusion</i>	70

CHAPTER 3: REGULATION AND TARGETING RULES APPLICABLE TO MILITARY AI DECISION-SUPPORT SYSTEMS71

3.1. *Regulatory Efforts and Their Normative Effect on the Lawful Use of AI Decision-Support Systems in Targeting* 73

3.1.1. Group of Governmental Experts on Lethal Autonomous Weapon Systems 73

3.1.2. First Committee of the United Nations General Assembly..... 75

3.1.3. Responsible AI in the Military Domain Summits..... 76

3.1.4. Political Declaration on Responsible Military Use of Artificial Intelligence and Autonomy 77

3.1.5. Other Regional Initiatives 77

3.2. *Application of Core Targeting Rules in the Use of AI Decision-Support Systems* 78

3.2.1. Distinction in Attack..... 80

3.2.2. Precautions in Attack 89

3.2.3. Proportionality in Attack 93

3.3. *Chapter Conclusion* 98

CHAPTER 4: THE REASONABLE COMMANDER STANDARD AND THE LOGICS OF QUANTIFICATION 101

4.1. *Reasonable Commander Standard and its Interpretive Grey Zones*..... 103

4.1.1. Common Ambiguities 103

4.1.2. Reasonableness and Other Provisions of International Humanitarian Law 106

4.1.3. Margin of Discretion: Risks or Benefits? 107

4.2. *The Impossibility of Programming Objective Reasonableness*..... 108

4.2.1. Public Surveys and the Illusion of “Common” Sense..... 108

4.2.2. Methodological Challenges of Programming Reasonableness 111

4.2.3. Situated Storytelling..... 112

4.3. *Technological Mediation of Commander’s Subjective Assessment of Reasonableness* 113

4.4. *Reasonableness as Norm Emergence* 116

4.5. *Procedural Approach and the Zone of Reasonableness*..... 118

4.6. *Chapter Conclusion* 119

CHAPTER 5: FROM SEEING TO KNOWING: THE IN/VISIBILITIES OF THE COMMON OPERATIONAL PICTURE 121

5.1. *Sensors, Data Infrastructures, and the Limitations of Seeing Everything* 125

5.2. *Algorithmic Filtering of Relevant Information*..... 126

5.3. *The Multiplicity of the Single Pane of Glass*..... 128

5.4. *Developers, Experimentation, and Common Complications*..... 131

5.5. *Regimes of In/Visibilities* 133

5.6. *Chapter Conclusion* 135

CHAPTER 6: ASSEMBLING CERTAINTY: THE QUEST FOR ACCEPTABLE THRESHOLDS OF DOUBT.....	137
6.1. <i>Legal Threshold of Doubt/Certainty in Targeting.....</i>	139
6.2. <i>A Primer on Confidence Scores.....</i>	142
6.3. <i>“Just Tell Me if it’s Right”: User Reactions to Confidence Scores.....</i>	144
6.4. <i>The Skeptical Soldier and the Practices of Setting Confidence Thresholds</i>	147
6.5. <i>The Techno-Legal Assembling of Knowledge Production.....</i>	152
6.6. <i>Chapter Conclusion</i>	154
CHAPTER 7: TOWARDS REASONABLE RELIANCE ON AI DECISION-SUPPORT SYSTEMS... 	157
7.1. <i>Constraints to Deliberation in the Use of AI Decision-Support Systems</i>	159
7.1.1. <i>Operational Tempo</i>	159
7.1.2. <i>Reduction of “Legal Breaks”</i>	160
7.1.3. <i>Limitations of Explainability.....</i>	162
7.1.4. <i>Trust as a Double-Edge Sword</i>	166
7.2. <i>When is Constrained Deliberation Unlawful? Considerations on Dynamic Targeting</i>	169
7.3. <i>Honest Mistakes? Legal Disincentives to Reject AI Outputs</i>	173
7.4. <i>Measures to Support Reasonable Reliance on AI Decision-Support Systems</i>	175
7.5. <i>Chapter Conclusion</i>	178
CHAPTER 8: CONCLUSIONS: LASTING DILEMMAS	181
8.1. <i>Main Findings</i>	182
8.2. <i>Main Contributions</i>	185
8.3. <i>Practical Suggestions</i>	186
8.4. <i>Future Research Avenues</i>	188
Bibliography.....	191
About the Author.....	218

List of Tables and Figures

FIGURE 1:

An illustration of the envisioned role of an AI system in assessing expected collateral damage. ...
..... 97

FIGURE 2:

A captured videoframe from a report published by the Wall Street Journal on 17 January 2023 representing the interface of a computer vision software used by TF59. 127

FIGURE 3:

A sample demonstration of BigBear.AI's Ursa Minor dashboard. 130

FIGURE 4:

A screenshot of an interactive tool illustrating the predictive certainty of words in a sentence in the BERT model. 143

FIGURE 5:

An illustration of a saliency map..... 163

TABLE 1:

Example of measures shaping reasonable reliance on AI DSS. 177

List of Interviews

All interviews utilized in this study were conducted by the author, either in person or via online platforms. References to these interviews within the text are identified by the respective date of each interview. All ranks and institutional affiliations mentioned below reflect the status at the time the interview was conducted. Interviews were accompanied by field observations.

Date	Location	Participant
28.01.2022	The Netherlands	Military practitioner 1 (NL)
31.03.2023	The Netherlands	Col Mietta Groeneveld, NATO Command & Control Center of Excellence
03.04.2023	Online	LtCol Ton van Heusden, Dutch Envoys to the U.S. Navy Task Force 59
09.05.2023	Online	Cdr Michael Brasseur, Commander, U.S. Navy Task Force 59
13.06.2023	The Netherlands	LtCol Wim A. Verloop, Head of Data Science, Artificial Intelligence and Robotics (NL)
27.07.2023	The Netherlands	Michael D. Street, Chief of Data Science & AI, NATO Communications and Information Agency
12.12.2023	The Netherlands	Military practitioner 2 (NL)
19.03.2024	The United States	Col Joshua F. Berry, Deputy Head of the Department of Law, the United States Military Academy at West Point
19.03.2024	The United States	Col Winston Williams, Head of the Department of Law, the U.S. Military Academy at West Point
02.05.2024	The United States	Military practitioner 3 (U.S.)
06.05.2024	The United States	Military practitioner 4 (U.S.)
23.05.2024	Online	Dr Jasper van der Waa, Scientific Researcher, TNO (NL)
30.05.2024	Online	Military practitioner 5 (U.S.)
06.06.2024	The Netherlands	Military practitioner 6 (NL)
22.07.2024	Online	Military practitioner 7 (U.S.)
12.08.2024	The Netherlands	Military practitioner 8 (NL)

Abbreviations and Acronyms

AI	Artificial Intelligence
AI DSS	Artificial Intelligence Decision-Support System(s)
AIP	Artificial Intelligence Platform
AIS	Automatic Identification Systems
AP I	Additional Protocol I
AWS	Autonomous Weapons Systems
C2COE	Command and Control Centre of Excellence (NATO)
C4ISR	Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance
CCW	Convention on Certain Conventional Weapons
CDE	Collateral Damage Estimate
COA	Course of Action
DARPA	Defense Advanced Research Projects Agency
GC	Geneva Conventions
GGE LAWS	Group of Governmental Experts on Lethal Autonomous Weapons Systems
ICL	International Criminal Law
ICRC	International Committee of the Red Cross
ICTY	International Criminal Tribunal for the Former Yugoslavia
IDF	Israeli Defense Forces
IHL	International Humanitarian Law
ISR	Intelligence, Surveillance, and Reconnaissance
IVAS	Integrated Visual Augmentation System
LAWS	Lethal Autonomous Weapons Systems
LLM	Large Language Model
NATO	North Atlantic Treaty Organization
NCIA	Communications and Information Agency (NATO)
NCV	Non-combatant and Civilian Casualty Cutoff Value
NGCV-S	Next Generation Combat Vehicle Suite
PTG	Perceptually-enabled Task Guidance
RAS	Robotics & Autonomous Systems
REAIM	Responsible AI in the Military Domain
ROE	Rules of Engagement
ROC	Robotics Operations Center
SIPRI	Stockholm International Peace Research Institute
STS	Science and Technology Studies
TF59	Task Force 59 (U.S. Navy)
U.K.	United Kingdom
UN	United Nations
UNIDR	United Nations Institute for Disarmament Research
UNGA	United Nations General Assembly
U.S.	United States
XAI	Explainable AI