



**TURUN
YLIOPISTO**
UNIVERSITY
OF TURKU

WAR MACHINE LEARNING

AI in Defence

Lauri Vasankari

University of Turku

Faculty of Technology
Department of Computing
Information and Communication Technology
Doctoral Programme in Technology

Supervised by

Professor, Jukka Heikkonen
University of Turku

PhD, Luca Zelioli
University of Turku

PhD, Paavo Nevalainen
University of Turku

Reviewed by

Professor emeritus, Mika Hyttiäinen,
National Defence University

Professor, Miklós Krész, University of
Szeged, Hungary

Opponent

Professor, Juha Röning
University of Oulu

The originality of this publication has been checked in accordance with the University of Turku quality assurance system using the Turnitin OriginalityCheck service.

ISBN 978-952-02-0645-1 (PRINT)
ISBN 978-952-02-0646-8 (PDF)
ISSN 2736-9390 (PRINT)
ISSN 2736-9684 (ONLINE)
Painosalama, Turku, Finland, 2026

*I should dedicate this dissertation to many, but it would not repay any of them.
Instead, I dedicate this work to pigheadedness, relentlessness, and perseverance.*

UNIVERSITY OF TURKU
Faculty of Technology
Department of Computing
Information and Communication Technology
VASANKARI, LAURI: War Machine Learning
Doctoral dissertation, 312 pp.
Doctoral Programme in Technology
April 2026

ABSTRACT

This dissertation examines the development and integration of machine learning within the military domain, arguing that the primary constraint and greatest opportunity for advancing military Artificial Intelligence (AI) is the data ecosystem. Across research in computer vision (CV), reinforcement learning (RL), federated learning (FL), and generative AI (GenAI), the analyses consistently show that progress is limited by systemic issues related to data availability, quality, and infrastructure.

The work synthesizes findings from six original publications to demonstrate that practical military AI requires a shift from an algorithm-centric view to a holistic, system-focused perspective that treats data as a first-class operational capability. To bridge the gap between high-level strategy and granular technical research, this dissertation adapts the Cross-Industry Standard Process for Data Mining (CRISP-DM) as a framework for assessing military AI applications.

Key findings from the studies validate this thesis. A CV study on sonar imagery highlighted model failure due to poor-quality sensor data, underscoring the need for integrated data pipelines. RL research revealed that a lack of high-fidelity simulators and operational data hampers real-world transfer. The investigation into GenAI identified a dependency on proprietary models misaligned with military needs, proposing FL as a secure, collaborative paradigm for developing military-specific foundation models. Finally, an ethical analysis addresses the "reliability-oversight paradox" in autonomous systems, proposing a new human-machine teaming model of human support rather than simple oversight.

In conclusion, this dissertation claims that the effective integration of AI into military forces depends on building a robust data ecosystem that includes expertise and understanding on doctrinal and policy-making levels, data and algorithm understanding on the technical level as well as governance, operator-in-the-loop feedback and annotation mechanisms, and interoperable infrastructure.

KEYWORDS: artificial intelligence, defence, military, machine learning, deep learning

TURUN YLIOPISTO
Teknillinen tiedekunta
Tietotekniikan laitos
Tietotekniikka
VASANKARI, LAURI: War Machine Learning
Väitöskirja, 312 s.
Teknologian tohtorihjelma
Huhtikuu 2026

TIIVISTELMÄ

Tämä väitöskirja tarkastelee koneoppimisen hyödyntämistä ja integrointia toimintaan asevoimissa ja sotilaallisessa toimintaympäristössä. Työn keskeinen väite on, että datakosysteemi on sekä merkittävin rajoite että suurin mahdollisuus sotilaallisen tekoälyn (AI) kehitykselle. Konenäön (CV), vahvistusoppimisen (RL), hajautetun oppimisen (FL) ja generatiivisen tekoälyn (GenAI) tutkimusalueita koskevat julkaisut osoittavat johdonmukaisesti, että edistystä rajoittavat systeemiset ongelmat liittyen datan saatavuuteen, laatuun ja ympäröivään tietotekniseen infrastruktuuriin.

Työ syntetisoi kuuden alkuperäisjulkaisun tulokset osoittaakseen, että käytännönläheinen sotilaallinen tekoäly vaatii siirtymää algoritmi- ja tekoälymallikeskeisestä näkökulmasta kokonaisvaltaiseen, systeemikeskeiseen lähestymistapaan, jossa dataa käsitellään keskeisenä operatiivisena kyvykkyytinä. Kaventaakseen kuilua korkean tason strategian ja käytännön teknisen tutkimuksen välillä tämä väitöskirja soveltaa CRISP-DM-viitekehystä (Cross-Industry Standard Process for Data Mining) sotilaallisten tekoälysovellusten arviointiin.

Tutkimusten keskeiset tulokset vahvistavat tämän teesin. Konenäköön keskittynyt tutkimus kaikuluotainkuvista osoitti mallien epäonnistuvan heikkolaatuisen sensoridatan vuoksi, mikä korostaa integroitujen dataputkien tarvetta. Vahvistusoppimisen tutkimus paljasti, että korkealaatuisten simulaattoreiden ja operatiivisen datan puute haittaa menetelmien siirtämistä todelliseen käyttöympäristöön. Generatiivisen tekoälyn tutkimuksessa tunnistettiin riippuvuus sotilaallisiin tarpeisiin soveltumattomista kaupallisista malleista ja ehdotettiin hajautettua oppimista turvallisena ja yhteistoiminnallisena mallina sotilaskäyttöön tarkoitettujen perusmallien kehittämiseksi. Eettinen analyysi käsittelee luotettavuuden ja valvonnan välistä paradoksia autonomisissa ja älykkäissä järjestelmissä ja ehdottaa uutta ihmisen ja koneen yhteistoimintamallia, joka perustuu ihmisen tukeen pelkän valvonnan sijaan.

Lopuksi tämä väitöskirja esittää, että tekoälyn tehokas integrointi asevoimiin on riippuvainen vankan datakosysteemin rakentamisesta. Tämä ekosysteemi edellyttää asiantuntemusta ja ymmärrystä doktriinien ja politiikan tasolla, teknisen tason data- ja algoritmiymmärrystä sekä hallintamalleja, operaattorin palautteen ja toiminnan huomioivia mekanismeja ja kokoavaa infrastruktuuria.

ASIASANAT: tekoäly, puolustus, asevoimat, koneoppiminen, syväoppiminen

Foreword and Acknowledgements

Apart from my supervisors, reviewers, opponent and custos, I am quite sure other people I ought to thank for aiding me in this endeavor will not eventually read this dissertation. I do not blame them.

Instead, as a published internal monologue, I wish to thank my supervisor, professor Jukka Heikkonen, for his unyielding assistance and support in my complete academic journey, through two Master's Degrees to this Doctor of Technology degree. We are still some years away from the Star Wars moment where the circle is complete, but without Jukka, this circulation might not have ever started. Fellow supervisors, Luca Zelioli and Paavo Nevalainen deserve acknowledgment due to their support in this endeavor, and Luca also as a fellow researcher making some of the original publications possible.

I also owe my interest in AI to my father, who gave me the initial push, now more than seven years ago, that has thus served as the pivot point in my career, from a naval officer to an AI professional. My life would not be on this track without him, even if we discount the initial onset of life provided. Same applies to my mother, albeit the scope of her influence is less poignant on subject matter expertise.

I owe thanks to my superiors and colleagues within the military as within my current company, the most impactful being Petteri Hemminki, Christian Andersson and my collaborating researchers Aapo Koski, Kalle Saastamoinen, and Adrian Borzyszkowski, as well as Mark Rempel, Marten Schaad and Maximilian Moll. Special thanks is owed to my friend, colleague, and research partner Jan Joutsu. Support from Matti Ristimäki and Heikki Härkönen has also been invaluable.

Finally, I owe my apologies to my loved ones, for being unavailable in this never ending pursuit of something still out of reach. Twisting the words of Elaine Rich, this is a journey towards things that, at the moment, remain unreached. I hope none of you hold a grudge against me.

Just like Uncle Scrooge I find resemblance in Robert W. Service's poem The Spell of the Yukon:

*"Yet it isn't the gold that I'm wanting so much as just **finding** the gold."*

27.3.2026
Lauri Vasankari

Table of Contents

Table of Contents	iv
Abbreviations	vii
List of Original Publications	x
1 Introduction	1
1.1 Background and definitions	2
1.2 Motivation, objectives, and methodology	6
1.3 Organization of the thesis	9
2 Literature review	11
2.1 Academic publications	11
2.2 Governmental policies and strategies	16
2.3 Think tanks	18
2.4 Other literature	19
2.5 Summary	20
3 AI and Military	22
3.1 The Intelligence Artifice	22
3.2 Military Domain	28
3.2.1 Tasks	29
3.2.2 Capabilities	30
3.2.3 Organization	30
3.2.4 Decision-making processes	31
3.2.5 Military Information Systems	33
3.2.6 On Complexity	37
3.3 Domain features of data	38
3.4 Main application areas	39
4 Machine Learning research areas	43
4.1 Computer Vision background	48
4.1.1 CV in military domain	53
4.2 Reinforcement Learning	55
4.2.1 RL in military domain	59

4.3	Federated Learning	64
4.3.1	FL in military domain	66
4.4	Generative Artificial Intelligence	67
4.4.1	Natural Language Processing	67
4.4.2	Generative Models	68
4.4.3	GenAI in military domain	71
4.5	Ethical considerations regarding AI systems	74
4.6	Testing, evaluation, validation, and verification	76
4.7	Field observations from Ukraine	78
4.8	Summary of Findings	79
5	Contribution of this thesis	83
5.1	Publication I: <i>Deep Mix: AI in Littoral Sonar Operations</i>	83
5.1.1	Summary	83
5.1.2	Methods and Data	83
5.1.3	Results and contribution	84
5.1.4	Author's contribution	84
5.2	Publication II: <i>Strategizing the Shallows: Leveraging Multi-Agent Reinforcement Learning for Enhanced Tactical Decision-Making in Littoral Naval Warfare</i>	85
5.2.1	Summary	85
5.2.2	Methods and Data	85
5.2.3	Results and contribution	86
5.2.4	Author's contribution	87
5.3	Publication III: <i>Reinforcement Learning for decision support in defense and security: A systematic review</i>	87
5.3.1	Summary	87
5.3.2	Methods and Data	87
5.3.3	Results and contribution	88
5.3.4	Author's contribution	89
5.4	Publication IV: <i>Emerging trends in federated learning: from model fusion to federated X learning</i>	89
5.4.1	Summary	89
5.4.2	Methods and Data	89
5.4.3	Results and contribution	90
5.4.4	Author's contribution	90
5.5	Publication V: <i>GenAI in Military: Trends and Opportunities</i>	91
5.5.1	Summary	91
5.5.2	Methods and Data	91
5.5.3	Results and contribution	91
5.5.4	Author's contribution	92

5.6	Publication VI: <i>The dilemma of AI reliability</i>	92
5.6.1	Summary	92
5.6.2	Methods and Data	93
5.6.3	Results and contribution	93
5.7	Conceptual framework	94
5.8	Methodological Framework for Synthesis	96
6	Conclusion	97
6.1	Summary of Key Findings	98
6.2	Implications	99
6.3	Limitations and Future Research	101
	Declarations	102
	Bibliography	103
	Original Publications	143

Abbreviations

Abbreviation	Meaning
AGI	Artificial General Intelligence
AI	Artificial Intelligence
API	Application Programming Interface
AR	Augmented Reality
AUC	Area Under Curve
BERT	Bidirectional Encoder Representations from Transformers
C2	Command and Control
CNN	Convolutional Neural Networks
COA	Course of Action
COP	Common Operating Picture
CONOPS	Concept of Operations
CoT	Chain-of-Thought
CRISP-DM	Cross-Industry Standard Process for Data Mining
CV	Computer Vision
DARPA	Defense Advanced Research Projects Agency
DDQN	Double Deep Q-Networks
DIANA	Defence Innovation Accelerator for the North Atlantic
DL	Deep Learning
DoD	Department of Defence
DP	Differential Privacy
DSS	Decision Support System
EW	Electronic Warfare
FedRL	Federated Reinforcement Learning
FL	Federated Learning
FMTL	Federated Multi-Task Learning
FPCA	Federated Principal Component Analysis
FTL	Federated Transfer Learning
GAN	Generative Adversarial Network
GenAI	Generative Artificial Intelligence
GNSS	Global Navigation Satellite System
GOFAI	Good Old-Fashioned AI

IID	Independent and Identically Distributed
IRL	Inverse Reinforcement Learning
IS	Information System
ISR	Intelligence, Surveillance and Reconnaissance
IT	Information Technology
KD	Knowledge Distillation
KL	Kullback-Leibler
LAWS	Lethal Autonomous Weapon Systems
LBP	Local Binary Patterns
LLM	Large Language Model
LRM	Large Reasoning Model
MADDQN	Multi-Agent Double Deep Q Network
MAPPO	Multi-Agent Proximal Policy Optimization
MARL	Multi-Agent Reinforcement Learning
MCM	Mine Countermeasure
MCTS	Monte Carlo Tree Search
MDP	Markov Decision Process
MDMP	Military Decision-Making Process
MILCO	Mine-like Contact
ML	Machine Learning
MLP	Multilayer Perceptron
MoE	Mixture of Experts
NLM	Neural Language Model
NLP	Natural Language Processing
NN	Neural Network
OODA	Observe-Orient-Decide-Act
OR	Operations Research
OT&E	Operational Test and Evaluation
PCA	Principal Component Analysis
PLA	People's Liberation Army
POMDP	Partially Observable Markov Decision Process
POSG	Partially Observable Stochastic Game
R-CNN	Regions with CNN features
RAG	Retrieval Augmented Generation
RF	Random Forest
RL	Reinforcement Learning
RLHF	Reinforcement Learning from Human Feedback
RNN	Recurrent Neural Network
ROC	Receiver Operating Characteristic
SOP	Standard Operating Procedure
SoR	System-of-Record

SSM	Soft Systems Methodology
SSS	Side Scan Sonar
STO	Science and Technology Organization
SVM	Support Vector Machine
T&E	Test and Evaluation
TEMP	Test and Evaluation Master Plan
TEVV	Test, Evaluation, Validation and Verification
UAV	Unmanned Aerial Vehicle
UGV	Unmanned Ground Vehicle
UMAP	Uniform Manifold Approximation and Projection
USV	Unmanned Surface Vehicle
UUV	Unmanned Underwater Vehicle
VAE	Variational Autoencoder
VGG	Visual Geometry Group
ViT	Vision Transformer
VR	Virtual Reality
WCSS	Within-Cluster Sum of Squares
XAI	Explainable Artificial Intelligence
YOLO	You Only Look Once

List of Original Publications

This dissertation is based on the following original publications, reproduced with the permission of the copyright holders, which are referred to in the text by their Roman numerals:

- I. Lauri Vasankari, Adrian Borzyszkowski, Luca Zelioli, Jukka Heikkonen. Deep Mix: AI in Littoral Sonar Operations. *J. Marine. Sci. Appl.* (2025).
<https://doi.org/10.1007/s11804-025-00695-4>
- II. Lauri Vasankari, Kalle Saastamoinen. "Strategizing the Shallows: Leveraging Multi-Agent Reinforcement Learning for Enhanced Tactical Decision-Making in Littoral Naval Warfare," In: Maglogiannis, I., Iliadis, L., Macintyre, J., Avlonitis, M., Papaleonidas, A. (eds) *Artificial Intelligence Applications and Innovations. AIAI 2025. IFIP Advances in Information and Communication Technology*, vol 712, Springer, Cham, pp 129–141, 2024.
https://doi.org/10.1007/978-3-031-63215-0_10
- III. Maarten Schadd, David S. Berman, Carolyn Chen, Mika Cohen, John Dorsch, Alexander Gegov, Maximilian Moll, Oliver Rose, Anna Rösner, Kalle Saastamoinen, Thomas Schiller, Andreas Strand, Lauri Vasankari, Mark Rempel. "Reinforcement Learning for decision support in defense and security: A systematic review," *Annals of Operations Research*, Springer, 2025. In publication.
- IV. Shaoxiong Ji, Yue Tan, Teemu Saravirta, Zhiqin Yang, Yixin Liu, Lauri Vasankari, Shirui Pan, Guodong Long, Anwar Walid, "Emerging trends in federated learning: from model fusion to federated X learning", *International Journal of Machine Learning and Cybernetics*, Springer, pages 3769–3790, 2024.
<https://doi.org/10.1007/s13042-024-02119-1>
- V. Lauri Vasankari, Aapo Koski, "GenAI in Military: Trends and Opportunities", *Scandinavian Journal of Military Studies*, 8(1), pages 416–434. 2025.
<https://doi.org/10.31374/sjms.415>
- VI. Lauri Vasankari, "The Dilemma of AI Reliability," in *Research Papers on Artificial Intelligence in the Military Operational Environment and Wargaming*, Saulius Keturakis, Arto Mutanen, Antti Rissanen & Jouko Vankka (eds.), *National Defence University*, Series 2: Research Reports No. 9, Helsinki, 2026, pp. 38–48. ISBN 978-951-25-3585-9.

1 Introduction

Artificial Intelligence (AI) is widely regarded as the next revolution in warfare, primarily as the enabler of autonomous weapon systems [1]. Beyond the role in autonomous systems AI is a multi-use, general-purpose technology with broad applications in warfighting, on different scales and tasks. Research institutions like RAND have conducted in-depth analyses of AI military impact, highlighting that the current development is commercially driven rather than state or defence-industry led [2]. Consequently, AI influence is pervasive, affecting all military functions beyond generic battlefield operations.

The application of AI in the military domain is not a new phenomenon. It can be argued that its first use in a military context occurred during World War II. At that time, Alan Turing and the Hut 8 team in Bletchley Park employed the Banburismus procedure, a method involving sequential Bayesian probability, electromagnetic Bombe computers [3], and manual analysis, to improve on the work of Polish cryptologists led by Marian Rejewski to decrypt the Enigma machine used by the Nazi armed forces [4]. Whether this constitutes a true application of AI remains a subject of debate and hinges on a precise definition, of which no universal standard has been agreed upon. This matter will be elaborated upon in Section 1.1.

The current, continuing trend of AI is built on its subfield known as Machine Learning (ML). In essence, ML is predictive modeling that utilizes an iterative *training* loop which is used to approximate a function that maps inputs to outputs. Hence, it is described as learning from data. The aim is to create a model that can perform well on an unforeseen data. Essentially, there exists a hypothesis space \mathcal{H} of applicable functions or features, namely hypothesis maps h that projects input $x \in X$ into output $y \in Y$ [5]. Therefore, if the annotations, i.e., actual outputs are known, a supervised learning ML algorithm can be described as

$$h(x) \rightarrow \hat{y} \sim y, \quad (1)$$

where the distance between the predicted output \hat{y} and actual, known output y is calculated with a cost or loss function \mathcal{L} , which is then used to update the function $h \in \mathcal{H}$ to minimize the error. A classical loss function is the Euclidean distance, $\mathcal{L}(y, \hat{y}) = \sqrt{\sum_{i=1}^n (y_i - \hat{y}_i)^2}$. For an unsupervised learning task, where the annotated y does not exist, the error is usually calculated as distance between inputs x , and then used to combine similar inputs into groups, a technique also known as clus-

tering. A hybrid solution known as semi-supervised learning has annotations for some of the data, while some or most of the annotations have to be deduced for the rest by the applied model or algorithm in the hypothesis space.

When the hypothesis map employed is a neural network with multiple layers, the technique is referred to as Deep Learning (DL), a term referencing the network's depth [5; 6]. Fundamentally, DL neural networks, of which Multilayer Perceptrons (MLPs) are the simplest and most straightforward approach, function as universal approximators applicable to a vast range of tasks. Most, if not all, modern breakthroughs in AI stem from the use of very deep Neural Networks (NNs). The large number of neurons, or network parameters, enables them to learn highly complex patterns from massive datasets and to generalize effectively across diverse application domains, from image analysis to natural language processing.

Fueled by advancements over the past three decades and now largely driven by the private sector [7], the use of AI has also proliferated within armed forces worldwide. As mentioned afore, in past decades the initial military interest in AI focused primarily on machine autonomy [8]. However, since the introduction of the transformer architecture with its self-attention mechanism [9] and the following launch of general-use natural language interfaces [10; 11], the focus has expanded to treat AI as a transformative capability in its own right. The potential to process ever-growing volumes of data is another significant driver, offering the ability to enhance situational awareness and facilitate faster, more accurate decision-making for a strategic advantage. AI is now recognized as a fundamental technology that may alter all aspects of human life, including warfare. NATO, for instance, has classified AI as a key disruptive technology, with an aim to maintain the technological edge by advancing AI [12]. Other strategic initiatives and public statements portray AI as a critical game-changer, potentially the proverbial "silver bullet" for achieving or maintaining military supremacy.

Reflecting the growing interest in and the emergence of novel military applications for AI, this thesis investigates the practical application of AI, specifically ML, across various military and defence domains and contexts. The investigation focuses on specific applications within the subfields of Reinforcement Learning (RL), Federated Learning (FL), Computer Vision (CV), and Generative Artificial Intelligence (GenAI).

1.1 Background and definitions

Common terminology and agreed-upon definitions are crucial for communicating complex ideas, a principle articulated by thinkers such as Francis Bacon [13] and conceptually mirrored by René Descartes' pursuit of foundational certainty [14; 15]. Bacon cautioned that imprecise words hinder the exchange of information and the advancement of knowledge, while Descartes advocated for clear and distinct ideas

instead of solely relying on words that can easily become disconnected from the subject itself. For centuries, this pursuit of precision and shared understanding has been essential to science and society. The absence of a universally accepted definition for AI, coupled with the variety of terms in use and the apparent disconnect between the word and the distinct, original idea in common discourse, present a significant obstacle to its systematic adoption.

The field of military AI is particularly affected by this ambiguity, as the act of defining terms is a strategic decision in itself. The Defence Acquisition University under the U.S. Department of Defense (DoD), recently renamed as Department of War, has highlighted a critical need to "align around logical AI definitions and terminology" [16], publishing its own analyses and glossaries. In academia, Russell and Norvig [17] provide a comprehensive overview of various definitions for AI, discussing the merits and drawbacks of each. For the purposes of this thesis, AI is defined as a computational system that applies logic and probabilities to solve problems that traditionally require human intelligence or are beyond human capabilities. This definition builds upon Elaine Rich's assertion that "AI is the study of how to make computers do things at which, at the moment, people are better" [18]. By this standard, the efforts of Hut 8 can be classified as an early AI application; despite the absence of digital computers, the computational processes employed exceeded the human capabilities of the era. Since then, the field has advanced rapidly, with modern AI solutions demonstrating near-ubiquitous applicability, as documented in the AI Index Report 2025 [7]. The report indicates that AI has matched or surpassed human baseline performance in many narrow fields, leaving only the most complex reasoning tasks as the domain of human intelligence.

In contemporary discourse, the term AI is often used to describe an active, human-like entity or system, a tendency known as anthropomorphism, which means the attribution of human characteristics to non-human entities [19]. One would not, for example, use the term mathematics in such sense. To maintain terminological consistency, this thesis avoids treating AI as an anthropomorphic agent that interacts with its environment in a human-like manner. McDermott [20] and Mitchell [21] have discussed this issue in terms of wishful mnemonics: terms that falsely suggest human-like properties in AI applications. A prominent current example is the term hallucination, used to describe the phenomenon where Large Language Models (LLMs) produce fallacious or nonsensical outputs due to their non-deterministic architecture [22; 23]. While the term evokes a human-like subjective experience, the underlying computational causes are well understood. A more precise term might be confabulation, though it could also be seen as anthropomorphic. Therefore, a neutral, non-anthropomorphic term such as erroneous generation is preferable.

At the highest level, AI can be divided into symbolic and connectionist approaches. The symbolic field, which relies on programmed rules and knowledge bases, is often referred to as Good Old-Fashioned AI (GOFAI) and has seen di-

minished focus in the modern era, though its principles remain crucial to software development [17]. In contrast, the connectionist field, which focuses on learning patterns from data, has dominated recent decades. The most prominent connectionist approach, ML, is a subfield of AI that augments the parent definition with the concept of learning from data without explicitly programmed rules [17]. In this context, "learning" is the iterative process of tuning a model's parameters to minimize error on a given dataset, thereby improving the accuracy of its hypothesis map h as shown in Equation 1.

Instead of referring to AI as a monolithic entity, this dissertation considers its applicable forms to be AI models. Building on this framework, an AI system is created when an AI model is integrated into a broader Information System (IS). Alternatively, in the symbolic GOF AI paradigm, an AI system can be built by algorithmically replicating the knowledge and decision-making processes of human experts, a leading approach from the 1950s to the 1980s [17]. Although some modern regulatory frameworks, such as the EU AI Act [24], explicitly exclude such rule-based systems from their scope, this dissertation considers them as AI systems because they fit the thesis's definition of a computational system applying logic to solve problems that originally require humans.

The term model itself carries multiple meanings. In Operations Research (OR), it refers to a mathematical formulation of a problem to be optimized [25; 26]. In symbolic AI, a logical model is a set of rules expressed in formal logic [17]. In ML, a model is the mathematical function produced by a training algorithm, defined by its architecture and a set of learned parameters. In RL, however, the term model typically refers to a representation of the environment, while the AI artifact is called an agent [27]. An RL agent learns a policy through trial-and-error interaction with the environment. To avoid ambiguity in this dissertation, AI model will refer to an ML artifact or a system's core logic, while environment model will be used exclusively in the context of RL. This convention ensures that "AI model" consistently refers to the mechanism by which an AI system maps inputs to outputs.

Integral to the development of an AI model are the processes of *training*, *validation* and *evaluation*. Training is the iterative process where model parameters are tuned to minimize error on a training dataset. A separate validation set is used to monitor the training process. This dataset allows for the assessment of the model's performance on unseen data during training. While providing insight to the training accuracy, it also helps to assess *overfitting*, a situation where the model memorizes the training data at the expense of generalization. Conversely, *underfitting* occurs when a model is too simple to capture the underlying patterns in the training data. Finally, an evaluation set, also known as a test set, is used to provide an unbiased assessment of the final model's expected real-world performance. The need for larger datasets for more complex models is explained by concepts such as Rademacher complexity, which measures a model's capacity to fit random noise [28].

To properly scope this research, the terms *national security*, *defence* and *military* must also be clarified. National security is the broadest of these, encompassing the safeguarding of a nation-state against all existential threats to its core values, territory, and population through the coordinated use of diplomatic, informational, military, and economic power [29]. It addresses a wide array of challenges, including but not limited to military threats.

Defence is the specific subset of national security concerned with countering external military threats. As a concept, it encompasses the full spectrum of state measures to resist a military attack. Its primary political objective is the preservation of the state, distinguishing it from offense, which aims at conquest [30]. While the distinction can blur at the tactical or operational level, where offensive actions may serve a strategic defensive goal, the overarching purpose remains preservation. For this thesis, defence is defined as the comprehensive framework of national capabilities, including military forces, strategic doctrines, industrial resources, and technological systems—oriented toward deterring aggression and protecting the nation’s sovereignty and interests from external threats. The military, in turn, is the principal instrument of the defence framework. It refers to the state-sanctioned armed forces that constitute the state’s monopoly on the legitimate use of physical force [31]. The military serves as an instrument of political will, capable of applying force, or the threat of force, to achieve political objectives [30]. This definition excludes non-state actors such as private military contractors (PMCs), insurgents, and terrorist organizations.

A doctrine describes the fundamental principles that guide how military forces conduct their operations and actions to reach their objectives, of which the U.S. doctrine for the armed forces [32] acts as an example. A Concept of Operations (CONOPS), as described in Joint Operation Planning [33], is a more concise expression of a commanders intent and the path to execution, usually developed for a specific task, operation or mission. A Standard Operating Procedure (SOP) is, as described in ADP 5-0 The Operations Process [34], a detailed, step-by-step, instruction that describes how to perform a recurring, preplanned or routine task or action, for example a river crossing for maneuverable ground forces.

Regarding the premise of AI in the military domain, the emphasis is often on the speed and cognitive intelligence that can be acquired through deploying novel AI based technology. These features are often framed as decision advantage [35]; making *better* decisions *faster*, and being more resilient against adversary’s development and actions. Decision advantage and resilience are, however, not new or revolutionary ideas, as these attributes have been highlighted by decorated strategists since Sun Tzu [36]. He highlighted foreknowledge, deception, speed, momentum and formlessness, which correlate with intelligent actions, intelligence, speed and resilience. As an another example, Alexander Suvorov highlighted three principles: speed, ”eye judgment”, and onslaught [37]. The eye judgment or ”eye measure”

refers to commanders battlefield intuition, the ability to assess a situation instantly and accurately, constituting of the whole context that encompasses the terrain, the enemy, own troops, decision points and timing. From a chosen perspective, the adaptation of AI in military domain can be seen as the latest frontier at which these recognized factors are taken to the next level in applying the same ancient principles of successful warfighting.

1.2 Motivation, objectives, and methodology

To address the research gap described in detail in Chapter 2, this research investigates the practical applicability of ML into military domain from multiple perspectives and application levels to identify solutions, challenges and future trends that will define the battlefields of the upcoming decades. There has been a considerable push in military AI research in the past years, as denoted in Chapter 2, of which the AI vision of the United States Department of Defense serves as an example [38]. Simultaneously, the perception of AI and its implications in warfare have indoctrinated a lot of variance, when experts and recognized spokespeople for AI claim more and more extravagant promises for the near future.

The underlying, system-level hypothesis is that AI can and will act as a core capability that will largely contribute to the success of the party that better utilizes it in operations, as predicted by NATO [12; 39] and other major military stakeholders [40; 41]. Hence, the purpose of this study is to provide an expert insight on the imminent implications of AI in the military context and what future applications and impacts can be deemed most likely according to the current state of research, development and deployment. The research contribution lies in rooting the current scope into the scientific background and quantifying the claims and future visions into respective truth anchors. Quantification, in this case, refers to projecting the theories into real-life use cases that provide concrete indications of the actual applicability and capabilities beyond hypotheses. The underlying motivation and required expertise is inherited from the author's military background and experience, which is combined to the field of view on the implications of AI and ML. Failure to understand the premise, requirements and limitations of AI as a technology can and arguably will have a tremendous impact on future conflicts that may shape the future of, for example, Western democracies. Therefore, the primary motivation for this dissertation is to serve as an exploratory research to aid in estimating the maturity and the depth of reach of the use and applicability of AI in Western military forces. While the research is general in its methodology and aims to produce generally applicable results, the research perspective inherits an European point of view from to the author's background and motivation to partake and enhance the development of European capabilities in this regard.

The objectives of this thesis are

- Performing a system-level investigation into the applicability, impact and needs of military AI capabilities.
- Providing empirical evidence on the performance and shortcomings of a selection of state-of-the-art ML methods in narrow military problems.
- Providing insight into future directions and challenges of novel ML research areas within the military domain.
- Integrating ethical considerations as a system-level factor into military AI capability development.
- Adapting Cross-Industry Standard Process for Data Mining (CRISP-DM) framework as a methodological tool to bridge technical AI research with capability development and strategic level decision-making.

The objectives are divided into two groups, technology assessment and applicability assessment, that build understanding on different levels.

Technology assessment utilizes a systematic review on ML paradigms and methods, namely RL, FL and GenAI, to provide background information on current state of the art as well as challenges related to each technology.

In applicability assessment, the mathematical premise, computational considerations and the possibilities and challenges related to these technical factors are examined through applicatory research. ML methods, in this case from the fields of CV and RL, are applied to solve constrained and narrow military problems to provide the concrete evidence and generalizable insights.

Together, these points of view are combined into the high-level understanding of the possibilities and limitations of AI in the military, enhanced by practical issues such as data availability, security constraints, problem complexity, explainability of outcomes and ethical considerations. The applicability of the CRISP-DM process from industry to defence is evaluated as a methodological approach and an objective.

The justification for this displayed objective selection lies in the system-level point of view of this research, which itself stems from the hypothesis that narrow applicability studies nor systematic reviews on a certain subtopic can tackle the fundamental bottlenecks of military AI development. This hypothesis has been worded in author's previous thesis on military sciences, which explored the use and development of AI from the perspective of the Finnish Navy, and stated in 2022 that "the degree to which the data could be utilized was very low, and collecting the data was challenging. This was not due to organizational resistance: the reception to all information requests was very positive, and there was expressed interest in the study's results. The difficulties stem from data capture and archiving practices, the distributed locations of databases, fragmentation of information, and security classifications. This is a similar challenge to that faced by the U.S. Navy, for which assembling sufficiently large data volumes into a usable form is identified as a critical need to enable research and development work" [42]. These observations gave

grounds to formulate the hypothesis that over different paradigms and research areas, the system-level issues are largely similar, and the bottleneck is less in the AI methods and algorithms and more in the overall digital readiness to develop and deploy such capabilities.

In summary, this research does not dive particularly deep into a particular technology within the field of AI, but instead examines the broad scope of methods and solutions in the military context from a system level perspective, providing grounded and thought-through insight on the military domain, and an assessment of applicable methodologies to create insight and establish functioning frameworks and policies to enable and exploit novel AI solutions. The overarching scientific contribution is the definition and validation of the 'Military Data Ecosystem' as a primary capability. This thesis demonstrates that the effective integration of AI is not primarily an algorithmic or AI model capability challenge, but a systemic one, providing a new theoretical framework that redefines data from an ephemeral byproduct into a managed operational asset.

To achieve these objectives, this thesis is built upon a series of targeted studies. The objective of technology assessment is primarily addressed through a systematic reviews of Reinforcement Learning (Publication III) and Federated Learning (Publication IV). The applicability assessment is realized through hands-on research in applying Computer Vision to sonar imagery (Publication I) , using Multi-Agent Reinforcement Learning for tactical decision-making (Publication II) , exploring opportunities in Generative AI (Publication V) , and examining AI deployment ethics (Publication VI).

The research presented in this dissertation is conducted as a cumulative work based on six original, peer-reviewed publications. This compilation thesis allows for an in-depth exploration of multifaceted research questions through a series of focused studies. The individual methodologies employed in each study, ranging from systematic literature reviews to the empirical application of machine learning models, are detailed within the respective publications (I-VI).

The overarching methodology for this dissertation is synthesis. It follows a structured approach to build a comprehensive understanding of the application of ML in the military domain. The research strategy was designed to align with the objectives outlined above, progressing from foundational technology assessments to practical applicability assessments.

The process involved:

1. Identifying Core Research Areas: The primary research areas of CV, RL, FL, Natural Language Processing (NLP), and GenAI were selected based on their emerging prominence and disruptive potential within the defence sector, as well as the fundamental differences between the paradigms. Essentially, CV and NLP can be viewed as ML paradigms on different modalities that can be

turned into generative methods as GenAI, while RL is a different approach to learning altogether, and FL functions as a possible umbrella for distributed application of any of these paradigms.

2. **Systematic Investigation:** Each core area was investigated through one or more original publications. This involved both systematic reviews of existing literature to establish the state-of-the-art and challenges (as in Publication III and Publication IV) and applicatory research where ML models were developed and tested against specific military problems (as in Publication I, Publication II, and Publication V) and an ethical, epistemic analysis of the non-technical challenges with AI.
3. **Synthesizing Findings:** The final step, which is the primary work of this thesis's introductory and concluding chapters, is to synthesize the findings from these individual publications. This synthesis aims to construct a holistic view, connecting the low-level technical insights from narrow problem-solving to the high-level strategic and operational implications for military forces.

The theoretical framework through which the original publications are synthesized is CRISP-DM. While not exactly academic, it shares resemblance to well-established theoretical methodologies such as Soft Systems Methodology (SSM) [43] and OR [26]. While SSM and OR are more profound scientific methods, CRISP-DM brings the theoretical background into a concrete and applicable process that is meant to provide concrete results for business and enterprises. Hence, for the scope of this research, CRISP-DM is transformed into a military-compatible format to assess the findings of the original papers. The layout of the CRISP-DM framework is displayed and described in Chapter 3.

The analysis was supported with a research visit that was conducted to Ukraine, an unfortunate but prime example of modern nationwide warfare in effect. The author visited several sites in Kyiv region between July 15 and July 25, 2025, while participating in a defence-focused venue to meet with local and international startups and military personnel to gather insights into the current state of technology and innovation, and the applicability of AI in the contemporary forms of tactical warfare.

1.3 Organization of the thesis

This thesis is structured to guide the reader from fundamental concepts of AI and the framework of military context towards specific, tangible research contributions, and finally to a high-level synthesis of the findings. An in-depth literature review in Chapter 2 provides an overview of the field from multiple publication perspectives including academia, governmental publications, think tanks, and non-scientific expert literature. Chapter 3 provides the necessary background on AI, defence and military organizations. Chapter 4 introduces the key ML research areas that form the basis of the original publications. Chapter 5 then details the core contribution of

this thesis by summarizing and synthesizing the findings of the six included publications, demonstrating a bottom-up approach where practical, low-level insights from specific applications are used to inform high-level considerations about the future of AI in defence. Finally, Chapter 6 concludes the dissertation by summarizing the key outcomes and discussing their broader implications and future research directions.

2 Literature review

This section reviews some of the most influential work on AI in military context, focusing on expert publications, dedicated think tanks as well as recognized scientific papers. This literature review creates the background for this dissertation and highlights the research gap that is addressed, as despite the breadth of research, there is a lack of operational-level bridging to fully integrate AI capabilities into the warfighting reality.

2.1 Academic publications

The AI research field is expanding at an accelerating pace, also in the defence and military fields. For example, a ScienceDirect database query "(artificial intelligence" OR "machine learning") AND (military OR warfare OR "armed forces" OR navy OR "air force" OR army)" for title, abstract or author-specified keywords, filtered to include engineering and computer science papers, returned 408 results from 1992 to halfway through 2025 when queried on first of August in 2025. The number of research papers per year grows exponentially, as shown in Figure 1. The words "defence/defense" and "security" are excluded as they induce hits on cyber defence and security, which is often unrelated to stark military context, albeit being just as applicable to the military as well. The sheer volume and rapid growth of this technical literature make a comprehensive review intractable, but also highlights a critical challenge: the potential for a widening gap between the highly specialized academic research and the strategic-level policy discussions reviewed later.

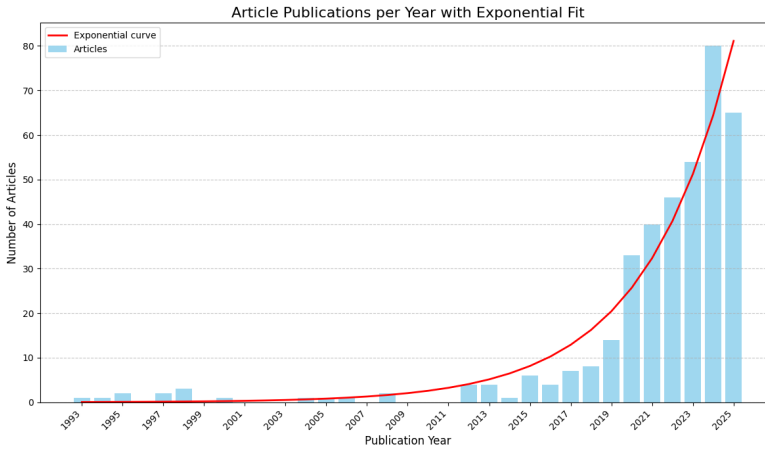


Figure 1. ScienceDirect publications per year for military AI indicating a growing trend.

To narrow down the results, being a dissertation on ML and AI, the bibtex information with abstracts was downloaded, combined and processed with three small, locally run language models, Llama 3.1 7B [44], Mistral 7B [45] and Gemma3 4B [46]. The task of the language models was to determine, based on meta information and abstracts, if the paper is actually focused in the military field and not just mentioning it in some context, according to the following prompts:

```
PROMPT_SCHEMA_EXAMPLE = {
    "name": "<string: article title>",
    "is_military_ai_ml": "<boolean: true if the abstract is about
AI/ML in a military context>",
    "topic": "<one of TOPIC_CHOICES>",
    "method_type": "<one of METHOD_CHOICES>",
    "key_findings": "<string: 1-3 concise points summarizing the
key findings>",
}

SYSTEM_INSTRUCTIONS = (
    "You are a meticulous research analyst. Read the abstract and
metadata. "
    "Decide if the work is about artificial intelligence or machine
learning in a MILITARY context. "
    "Choose ONE topic and ONE method_type from the provided
choices. "
    "Answer STRICTLY as minified JSON matching the schema. Do not
include explanations, Markdown, or backticks."
)

USER_TEMPLATE = (
    "Paper metadata as JSON follows. Return a single JSON object
with keys: name, is_military_ai_ml, topic, method_type,
key_findings.\n\n"
```

```

"ALLOWED TOPICS: {topics}\n"
"ALLOWED METHOD TYPES: {methods}\n\n"
"JSON SCHEMA EXAMPLE (values are placeholders):\n{schema}\n\n"
"PAPER: {paper_json}"
)

```

After the processing, the results were examined for all 408 papers, and if the paper got marked as relevant (`is_military_aiml = True`) by at least two out of three small models, it was examined more closely. This approach narrowed down the search to 72 papers.

This analysis of 72 scholarly articles reveals a clear and concentrated focus within the domain of computational intelligence in military applications. The literature is predominantly characterized by applied research aimed at developing tangible solutions, particularly in the areas of autonomous systems, intelligence gathering, and decision support. The distribution of research topics underscores a significant academic and practical interest in three primary areas, which together account for nearly 78% of the reviewed literature.

Command and Control (C2) and decision support is the most dominant theme, with 20 articles. The research explores a wide scope of topics from systems that simulate operational procedures [47] to providing early warnings to improve international stability [48].

Autonomous systems & Robotics are the second-largest category with 10 articles followed by Intelligence, Surveillance and Reconnaissance (ISR) with 9 articles. ISR research highlights the critical role of data processing and analysis in modern military operations. The focus is on leveraging computational methods to extract actionable intelligence from vast amounts of sensor data. Uncrewed systems research is heavily focused on practical applications, such as using machine learning for real-time object recognition for unmanned vehicles, e.g., Buluswar and Draper [49], and predicting structural responses to blast loads [50], indicating a drive towards creating more resilient and intelligent unmanned platforms.

Other topics such as Logistics and maintenance (8 articles), Cybersecurity (6 articles), Medical (3 articles), and Personnel (3 articles) represent smaller but notable areas of research. In contrast, foundational domains like Electronic Warfare (1 article) and Communications (1 article) appear significantly underrepresented in this body of literature. For Communications, it has to be stated that three papers concern networks, but under another topic, such as Unmanned Aerial Vehicles (UAVs) or cyber security [51; 52; 53]. The analyzed papers are represented in as a topic summary in Table 1.

The methodological landscape of this particular sample is overwhelmingly skewed towards practical implementation, reinforcing the applied nature of the research field. Application and implementation was the methodological approach for a vast majority of the papers in Table 1, covering 51.39% of the papers. This indicates that the

Topic	Count	Citations
Command & control / decision support	20	Ângelo Lellis Moreira et al. [54]; de Araújo Costa et al. [55]; James and Herget [56]; Jiang et al. [57]; hsien Liao [47]; Zabala-López et al. [58]; Sánchez-Ruiz and Miranda [59]; Liebowitz and Davis [60]; Liu et al. [61]; Masud et al. [62]; Nõmm and Venables [63]; Oh et al. [64]; Perry et al. [65]; Mechergui and Jayakumar [66]; Mendonça et al. [67]; Xia et al. [68]; Yadav and Kim [69]; Aha [70]; Scrimgeour [48]; Zhao et al. [71].
Autonomous systems & robotics	10	Altinors et al. [72]; Batista et al. [73]; Buluswar and Draper [49]; Fualdes and Barrouil [74]; Gilmore [75]; Hosseinzadeh et al. [51]; Kaur et al. [53]; Sutton and Roberts [76]; Rahmani et al. [77]; Cai et al. [78].
Intelligence, surveillance & reconnaissance (ISR)	9	Akbal et al. [79]; Wei et al. [80]; Guo et al. [81]; Zhao and Morikawa [82]; Hashemi and Hall [83]; Kılıç et al. [84]; Kwon and Lee [85]; Mehta and Shah [86]; Petrov et al. [87].
Logistics & maintenance	8	Baker et al. [88]; Mohril et al. [89]; Candelieri et al. [90]; Bortolan Neto et al. [50]; Li et al. [91]; Malkoff [92]; Vasilikis et al. [93]; Boutselis and McNaught [94].
Cybersecurity & electronic warfare	6	Akbani et al. [52]; Whelan et al. [95]; Sojitra et al. [96]; Maathuis and Cools [97]; Almaslakh [98]; Shamshirband et al. [99].
Personnel	3	Hoecherl et al. [100]; Wasilefsky et al. [101]; Zhang et al. [102].
Medical	3	Satava [103]; Ahamed et al. [104]; Gondalia et al. [105].
Human-machine teaming / HCI	2	Canan et al. [106]; Sheridan [107].
Air operations	2	Wittig and Onken [108]; De Giorgi and Quarta [109].
Small arms	2	Chandan et al. [110]; Yang et al. [111].
Wargaming & simulation	2	Klahr [112]; Knapp et al. [113].
Targeting & fire control	2	Govindarajan et al. [114]; Li et al. [115].
Communications	1	Aloqaily et al. [116].
Electronic Warfare	1	Wang et al. [117].
Innovation & adaption	1	Kagiwada [118].

Table 1. Topics with representative citations.

field is primarily concerned with building, testing, and implementing computational solutions to specific, narrow military problems rather than more exploring foundational concepts. 16.67% of the papers introduced methodological novelty in AI or ML, although the novelty and methodological impact is open for reinterpretation especially in comparison with some application papers. 9.72% papers proposed framework or architecture solutions, with or without experimentation, some akin to methodological novelty papers. Review and survey papers constituted for 20.83% of the papers, and the innovation paper by Kagiwada [118] is in fact an essay from personal experience in the field.

For a short look-through, in the top 25 ScienceDirect query results, sorted by relevance, there are 16 open access papers with direct military relevance. These can be classified into C2 & DSS, small arms, autonomous systems & robotics, personnel-related topics, Decision Support System (DSS), and cyber defence. The reviewed technical literature reveals a strong focus on enhancing existing military functions rather than creating revolutionary new ones. In warfighting, concrete low-level problems such as weapon target assignment [115], small arm firing skill evaluation [110], chemical weapon detection [80], and assessment of urban destruction [82] have been examined. On a higher level decision support, interest towards AI or ML based military DSSs is trending [58; 55; 59]. Another trend, uncrewed (previously known as unmanned) systems related research includes non-Global Navigation Satellite System (GNSS) based navigation for UAVs [73], intelligent communication solutions [116], swarming [78], and UAV related cyber capabilities [95]. The digitally cross-sectional cyber security perspective has been researched from, e.g., Explainable Artificial Intelligence (XAI) point of view on malicious data classification [96] and overall evaluation of AI-based cyber security solutions [97]. Applicability of ML techniques to personnel-related tasks, such as retention [100], candidate selection [101] and harm prevention [102], can be hypothesized to be related to the availability of structural data and close resemblance to civilian problems in similar areas. The focus on discrete, solvable problems within the technical literature exemplifies one side of the research gap, showcasing deep but narrow progress that is often difficult to translate into broad strategic advantage.

In summary, the scientific literature on ML and AI in military contexts is heavily focused on the practical implementation of algorithms to solve problems in C2 & DSS, autonomy, and ISR. The primary goal of engineering and computer science efforts appears to be in creation of functional, real-world applications. The field seems to prioritize the development of intelligent systems that can perceive the environment, process information, and support or supplant human decision-making. The relative scarcity of theoretical and foundational research shows the focus on certain field, i.e, the military: foundational research happens in more general scope, and military applications follow those innovations. This focus on applied, narrow problems becomes even more apparent when examining the most prominent papers from this

cohort, which reveal a clear pattern of enhancing existing military functions rather than creating revolutionary new ones.

2.2 Governmental policies and strategies

The policy documents from NATO and the US DoD operate at a high level of abstraction, articulating principles and goals that often lack a clear connection to the granular technical capabilities currently being developed, thus representing the other side of this gap. From a governmental perspective, 2021 NATO Artificial Intelligence Strategy and 2024 revised strategy aim to provide the Alliance with a aims and outcomes, underpinning responsibility principles such as lawfulness, traceability, reliability, governability and bias mitigation AI [39; 119]. In the United States, DOD Ethical Principles for AI provide similar ethical principles and guidelines for acquisition [120]. Data, analytics and AI Adoption Strategy [40], superseding 2018 AI strategy, aims to improve the organizational environment to enable achieving decision advantage with AI. The nation-wide AI competitiveness and defence innovation priorities have been outlined in National Security Commission on AI final report [121]. Within NATO Defence Innovation Accelerator for the North Atlantic (DIANA), accelerator hubs and test centers have been established across the alliance, focusing on AI, autonomy, quantum, and other revolutionary technologies [122]. U.S. Department of Defense [123] demonstrates rapid tri-lateral deployment of AI and autonomous systems. Regarding Lethal Autonomous Weapon Systems (LAWS), DoD Directive 3000.09 [124] establishes legal and technical safeguards and framework, for example, the responsibilities and requirements regarding use of lethal force. For United Nations, International Committee of the Red Cross [125] expands the LAWS discussion to humanitarian-law concerns and recommendations for UN. United Kingdom's Defence AI Strategy [126] along with [127] outlines UK-specific governance and assurance mechanisms. Actual, executable regulation and laws around AI are somewhat non-existent. The European Union's landmark AI Act [24], in a deliberate policy choice, excludes military applications from its scope. This allows the EU to advance a market-focused regulatory policies while pursuing defence AI separately. Hence, defence related AI is mainly unregulated and while there is considerable research on the subject, regulatory and legislative framework is in its infancy.

From a non-Western point of view, People's Liberation Army (PLA) incorporates AI to its modernization strategy, labeled "intelligentization", which aims to develop a world-class military that leverages AI for new forms of warfare and transforming key areas such as situational awareness, decision-making, unmanned systems and cognitive warfare domain [41]. PLA operational concepts [128], inferred from strategic guidelines and recent examples of PLA in combat, show that the desired modernized status is to create information dominance, deal new realities between combat and war space, and be able to defeat adversary's operational system through

target-centric warfare. In Chinese military writings, combat space is the geographic area where physical conflict occurs, while war space encompasses all domains of war from physical to non-physical, including political, economic, diplomatic and information spheres. In the operational concept perception, the war space is expanding while combat space is shrinking. Target-centric warfare denotes the use of precision-strike capabilities and intelligent munitions to surgically impact the combat space.

These advances underline the ongoing race towards military supremacy between competing superpowers, namely the United States and China. Borchert et al. [129] have published a very thorough case study into the defence AI that covers 25 countries, highlighting that defence AI is at the center of geopolitical and geoeconomic competition. The introductory chapter notes that the case study aims to fill the research gap of how countries think about defence AI, how they prepare for its adoption, and how they develop existing concepts and processes and related capabilities. The study identifies three strategic motives for defence AI: threat-based, fear or falling behind, and AI as a capability multiplier. Most of the countries reside in the third category, while the countries at the bleeding edge such as China and United States are in the first, accompanied by smaller countries that aim to maintain their strategic edge against prominent opponents. These countries include Greece, South Korea, India and Ukraine. From capability categories, the most sought-after is the combination of AI with uncrewed systems, followed by predictive maintenance, C2 combined to data analytics and data management, Electronic Warfare (EW), wargaming and, in minority, mission planning and tactics development. The research highlights that Russia and China seem to considerably prioritize capabilities that aim towards autonomous reconnaissance-strike complexes. The research indicates that almost all nations are focused on data-driven and correlational learning, which can be effectively translated as ML. The approach is denoted the second wave of AI, in accordance with Defense Advanced Research Projects Agency's (DARPA's) three waves of AI technology [130], while the United States is the only nation exploring the third wave by focusing on contextual reasoning and self-learning under uncertainty. The tension between sovereignty and cooperation is also highlighted as national interests compete with collaborative resources. Despite the recognized impact, it is prevalent that most countries focus on training military personnel to handle specific AI systems instead of advancing their general AI talents and competence. The insights highlight the fact that current military innovation in AI is emulation, mainly mimicking the US, while it usually aims to enhance existing practices and systems. Human-in-the-loop and human-on-the-loop solutions are emphasized globally, and the case study points out that the "valley of death" problem, where adopting promising AI technologies to practice is not a straightforward process that often times fails, is a common struggle for most nations.

2.3 Think tanks

RAND corporation, a non-profit American think tank that conducts research and analysis on a wide range of public policy issues, including defence and national security, has published research reports regarding AI since 1960s. In their database, 211 research papers have been published, spanning from 1962 until today. It is conceivable that the "AI winter" and the "dotcom" bubble seem to correlate with the publication frequency, as there are gaps from 1973 to 1989, from 1993 to 2001 and from 2002 to 2012. The vast majority of the research, 193 reports in total, have been published since 2017, while only 18 papers cover the interval from 1962 to 2012. This also correlates with the progress of digitalization, AI and ML, including the AI winters when AI failed to meet expectations. In a similar manner, SIPRI (Stockholm International Peace Research Institute) that focuses on research into conflict, armaments, arms control and disarmament have published reports on military AI since 1987. Their database returns 21 publications when querying for "artificial intelligence", of which 20 are linked to armed forces and military.

RAND usually employs a scenario-based and policy oriented methodology while SIPRI leans towards arms control and international law. For both institutions, early research addresses issues such as problem-solving programs [131], NLP and symbolic AI [132; 133], as well as early neural networks [134], advanced computing [135] and strategies [136]. There is a major shift in later research, where the conceptual and theoretical approach turns to more concrete applications and implications, as well as the introduction of specific military topics from strategic to tactical levels, akin to the categories represented afore. The identified categories include DSS and C2, logistics [137], cyber defence [138], space technologies [139], human resources, uncrewed and autonomous systems [140; 141], responsible use [142; 143] and governance [144] as well as strategic principles [145] and bias mitigation [146]. The DSS and C2, in this case, include wargaming [147; 148], as wargames can be defined as "representations of conflict or competition in a safe-to-fail environment, in which people make decisions and respond to the consequences of those decisions" [149]. This can be seen as an exercise or test-time environment prior to decision-making itself. In addition, a conceptual review on the future of C2 systems [150] has been conducted. C2s can be perceived as systems that produce situational awareness that allow decision making while also allowing one to act on the decision. In human resource management, ML is also related to decision-making [151] and its fairness [152]. Notably, the work of Schulker et al. [151] aligns with that of Wasilefsky et al. [101], as both focus specifically on Air Force applications, creating a targeted body of research in this area. Due to the scale and gravity of nuclear capabilities and the nature of, e.g., SIPRI's methodology, the impact on stability, deterrence, and nuclear risks are well presented in the related reports [153; 154; 155; 156; 157; 158; 159].

Both institutions have researched similar and supplemental topics on military AI,

which accumulates into a comprehensive high-level understanding of the impact and reach of this cross-sectional technology. The low-level impact is less researched, and often the results are policy suggestions to better enable, govern, and manage the agreed upon change that AI, as a capability, brings to different warfighting and military functions.

2.4 Other literature

In addition to academic literature, governmental policies and strategies as well as think tank reports, there is a small number of expert books that have been especially influential in shaping how policymakers, defence practitioners and the wider audience discuss AI, autonomous weapons and the future of warfare. These sources have substantial insider subject-matter expertise and synthesis value to them, although they simultaneously display a subjective point of view that can be deemed likely to include author bias, agenda-setting and a selective framing. Despite these limitations, they remain useful for capturing the dominant practitioner narratives that frequently guide real-world discourse. This thesis selects a sample of books from Brose, Scharre and a professional anthology edited by Tangredi due to wide engagement in defence policy and professional military discourse, jointly covering complementary analyses on the subject.

In this category, Colby award winning "Army of None" [160] by Paul Scharre was one of the most widely cited early syntheses to describe the advance of AI and especially autonomous systems in military forces, focusing on United States [161]. With field experience and a long career in military affairs, Scharre manages to describe the history, level of development, and trends in a manner suitable for wide audiences. Highlighting the premise of early drones and the underlying AI technology, the book has become a prescient of the military reality of the 2020s.

Scharre's book was followed by a practitioner anthology "AI at War: How Big Data, Artificial Intelligence, and Machine Learning Are Changing Naval Warfare", edited by Tangredi and Galdorisi [162]. It serves as an effort to provide a balanced and practical overview that seeks to demystify the technology for a non-technical audience of national security professionals, policymakers, and concerned citizens, examining both the promising applications and the inherent limitations of AI in a defence context. Addressed topics include high-level strategy, policy, doctrine, specific weapon systems, and pressing ethical concerns. Although its subtitle specifies a focus on naval warfare, its themes and findings are abstract and thus relevant across all military services and the broader defence community. The book also aims to identify the significant real-world barriers to AI adoption, including entrenched institutional cultures, inter-service rivalries, and the political realities of defence acquisition. It explicitly cautions against over-reliance on AI, particularly in contested environments where systems could be vulnerable to cyberattacks or electronic de-

ception. This skepticism is balanced with a sense of urgency, as the book frames AI development within the context of great power competition.

Brose [163], a former Senate Armed Services Committee staff director, argues that the current state of the US armed forces is falling behind in the critical areas of AI, autonomous systems, and networked warfare when compared to possible adversaries, mainly China. Brose claims that the victor of future conflicts will be the side with the faster, more resilient, and more intelligent kill chain, the namesake of the book, which refers to the decision-making loop from observation to action. In other words, the United States risks losing future wars, not because of insufficient defence spending, but instead due to outdated systems, slow procurement processes, and institutional inertia. In his perception, modern warfare will be defined by speed, autonomy, and decision-making advantage that are all domains where AI can be a crucial component. The book has been reviewed by RAND [164] and National Defense University Press [165].

The latest expert addition to these insights is Scharre's newest book, "Four Battlegrounds - Power in the Age of Artificial Intelligence" [166], reviewed by United States Army War College Press [167]. The book is an overall review of AI in strategic, military context, as the work focuses on the fundamental rivalry between competing nations and the technological race for advance that follows. It covers aspects from historical analogies to modern equivalents, data and hardware insights, as well as concrete and foreseeable applications, benefits, and threats, of AI.

2.5 Summary

In summary, the literature paints a clear picture: the strategic imperative for military AI is widely accepted [163; 38; 168; 126; 41; 119], and a broad consensus exists on the most promising application areas [129; 162; 40]. However, this review has also highlighted a persistent 'valley of death' [129] fueled by institutional inertia [163] and a nascent understanding of the associated risks and governance requirements [153; 24]. Consequently, a significant gap remains between the high-level strategic discourse and the granular, practical research needed to bridge concept with capability. The identified research gap is the disconnect between concept and policy research versus applications, as well as between AI research and development versus real world military requirements.

This dissertation directly addresses this disconnect. The six original articles that form the basis of this work are designed to bridge the gap between strategic concepts and operational reality. Publication I and Publication II are applicatory studies of technical, algorithmic solutions to different level military problems, providing insight not only to the problems themselves but into the challenges such adaptation of AI capabilities face in the military context. Publication III introduces similar issues on an abstract level of military DSS, and Publication V highlights the existing

gap between R&D and deployment for GenAI solutions in the military. Publication IV introduces FL as an ML paradigm that provides solutions to some of the presented challenges. Publication VI concludes the analysis with an ethical discussion on human-machine teaming and roles of responsibility. Together, these original papers collectively build the very framework that the literature shows is currently lacking. By examining ML research areas from this perspective, this research provides a synthesized model for how military AI can move from a set of disconnected applications and abstract policies to a truly integrated warfighting capability.

3 AI and Military

This chapter provides a brief background on the field of AI and a generalized glimpse into the fields of defence, warfare and the military framework.

3.1 The Intelligence Artifice

The history of AI starts from the 19th century, pioneered by Ada Lovelace [169] and lord Boole [170]. Closely related Bayesian conditional probabilities [171] were introduced a century earlier, and for example the aforementioned Euclidean distance inherits the name from Euclid [172] and his work around 300 BC. The foundation laid by these theories before digital computers exhibits the incremental scientific method, where even the current state-of-the-art leverages centuries old science to come up with novel solutions.

Figure 2 shows the subfields of AI and a way it can be understood through leading paradigms, their intersections, and the modern research areas of which most are addressed in this dissertation. The figure is based on the works of Russell and Norvig [17], Bishop and Bishop [173], and Pearl [174]. The symbolic AI era of dominance spanned from the 1950s to the 1980s [175; 176; 17; 177], during which the symbolic approach was considered the paradigm that would eventually lead to human-like intelligence through the manipulation of symbols and rules. Among the most prominent realizations of this paradigm were expert systems [178; 179]. The paradigm has not vanished but instead is a stable in common information processing and programming languages as well bridging numeric methods to symbolic domains and vice versa, creating the neuro-symbolic approach also shown in Figure 2. The probabilistic approach started its rise in 1980s, reaching maturity by 2000s, and remains central for probabilistic machine learning, Partially Observable Markov Decision Process (POMDP), and overall uncertainty modelling.

The numeric approach started out in parallel with the dominant era of symbolic AI, in 1950s, with early discoveries such as the perceptron [180] as the basic component of current linear neural networks. Due to the theoretical scrutiny [181] and limitations in computation capacity, the paradigm laid near dormant from 1970s to late 1980s, experiencing a re-emergence during the 1990s through statistical reasoning, probabilistic models, kernel methods, and classic ML algorithms such as Support Vector Machines (SVMs) [182]. The deep learning dominant era made its

breakthrough during 2010s and has been the leading paradigm since.

As explained in Section 1.1, despite lengthy and mathematical history of AI, there is no generally accepted definition for the term. Even the term AI was disputed, and it supplanted other terms in the 1950s when John McCarthy coined it [183], dwarfing competing terms such as computational intelligence. As there is no general definition, the understanding of the subject varies more than what usually happens in a living language. It should be noted that the term AI used to continuously mean something that has not yet been achieved. Elaine Rich defined AI as a field of science that aims to make computers do things that currently people are better at [18], which is a robust definition as it integrates the advance of the technology and the sliding of meaning.

However, after the release of world-renowned GenAI applications such as ChatGPT [11], Gemini [184], Claude [185], Llama [44], Grok [186], Mistral [45] and others, the term AI has now become a synonym with transformer-based generative applications. This is somewhat troublesome, as now the understanding of the broader scope of AI and its foundations is blurred by the familiarity of human-like interaction exhibited by applications that usually combine several AI methods and research areas to provide a suitable full-stack solution.

In this research, the term AI is used in an old-fashioned way as explained in Section 1.1, describing a field of science that aims to create intelligent computer systems that can execute tasks that formerly required humans. ML is understood as the most promising subfield within AI that has enabled leveraging data and computation to create models without the need to build their logic from scratch. The logical approach, also known as symbolic AI, was the most prominent paradigm from the 1950s to 1990s [17], after which ML and so called numeric solutions became the winning paradigm [17]. Deviating from statistical methods that aim to explain data, ML in general aims to create predictions based on data: this division is somewhat blurred on occasion, but generally a good distinction.

As shown in Figure 2, DL is a subfield of ML where the difference is in the use of sophisticated, multi-layer neural networks in a multitude of different network architectures to solve problems. The term deep refers to the depth created by these *layers*, although mathematically the layers are a chain of functions or a function composition. For a three layer neural network $F(x)$, each layer performs a function f_i , so that the initial input x is transformed turns into the output o_i which is then fed to the next layer function f_{i+1} , as displayed in Equation 2:

$$F(x) = f_3(f_2(f_1(x))) \rightarrow y \quad (2)$$

The approach, in ML or DL, is not different from other forms of functional optimization, as the trained models are optimized to perform with certain data and certain tasks. Before the launch of very large transformer-based models, DL solutions

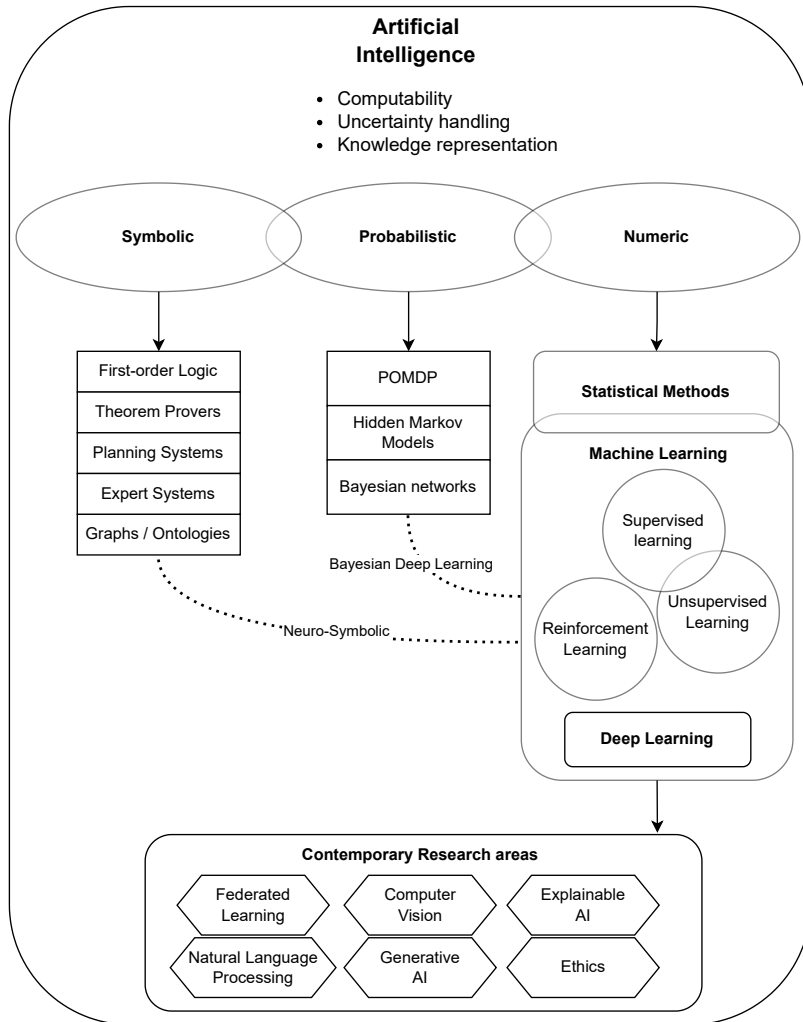


Figure 2. The Paradigms, Methods and Research areas of AI

were still considered narrow, as they were usually fitted into a certain task to solve a certain, narrow problem. After the currently available computational resources and the availability of data was exploited to train extremely large language-based transformers, also known as LLMs, it can be stated that the latest AI models are able to generalize to a multitude of tasks without further training or fine-tuning. This has had a tremendous impact on the field of AI. As a result, the current trend resembles a race toward the best model that could, hypothetically, be an Artificial General Intelligence (AGI), meaning that it could be able to solve any problem in a human-like manner. The feasibility of an AGI remains debatable, as there is no paradigm, research field or theory that shows a proven premise in reaching AGI, which indicates that we lack the architecture to reach it, as suggested by a decorated AI scientist Yann LeCun [187]. He has also stated that the future of AI will not be generative [188]. The majority of scientists in the field tend to think alike [189], while industry leaders tend to anticipate AGI within a few years at earliest [190; 191; 192]. Furthermore, the whole discussion is distorted by the fact that there is no agreed definition for an AGI either.

Despite advances, there are still clear limitations in AI models that hinder their performance. Lately, researchers have demonstrated that the LLMs and Large Reasoning Models (LRMs) merely give the illusion of thinking and understanding [193; 194]. Shojae*† et al. [193] research was criticized for placing constraints on the models like limiting context-window size and disabling code-based solutions, but as the context enables reiteration it arguable enables a form of brute forcing, and code-based solutions can be extracted from the memory of the model, these design choices can also be defended. Mancoridis et al. [194] showcase the illusion of thinking by examining the way LLMs fail to understand concepts in a human manner which leads to memorizing instead of deeper knowledge and renders certain benchmarks invalid in measuring the models so-called cognitive performance. The ARC-AGI leaderboard [195] showcases that the size of the models and the performance has indeed increased, leaving a decreasing gap between human-level performance and the current state-of-the-art. It can be argued that the last mile is the hardest to beat, but it can also be stated that novel innovations may take the field by surprise.

As presented in the Chapter 1, the ML methods aim to fit a model to the data by decreasing the error that results from the predictions and the chosen evaluation metric. This applies to the state-of-the-art LLM and LRM as well. The key is to have adequate high-quality data that enable training the model, a sophisticated model architecture that is able to exhibit intelligent properties when used in inference, and a validation and evaluation methodology to assess whether the model performance is good to begin with. As a common framework, the CRISP-DM displayed in Figure 3 [196], is as valid as ever.

CRISP-DM necessitates existence of relevant data. The data is leveraged through business understanding that translates into data understanding. In other words, the




Figure 3. CRISP-DM

business understanding precedes interpreting the data and creates the insight to leverage some new performance or capability from the data. After a data understanding has been reached, in collaboration with the business understanding, the data is prepared which usually includes cleaning and transforming it into a suitable format. Then, the modeling occurs. The modeling is iterative and it can be also be incremental, consisting of training and tuning phases. Resulting models have to be evaluated with regard to the original business understanding, by comparing the model performance to the actual, recognized need for example. After evaluation has been completed, the model can be deployed.

On an important note, the CRISP-DM process may halt before deployment. For example, if there is no relevant data available, or that the quality of data is insufficient, the process will not proceed to deployment. Likewise, if the evaluation fails, the process redirects back to start and may necessitate an alternative approach. Data itself is the information prerequisite for knowledge, but it is not knowledge by itself. Also, data can be viewed from a multidimensional perspective, as it can possess different qualities, such as temporal, granular, and structure, displayed as general trade-offs in Figure 4. In addition to the qualities shown in the figure, for example veracity and completeness are factors that greatly affect the usability of data: how reliable is the source or the data itself, is it complete or does it require fusion with other sources in order to be useful.

In essence all ML development and deployment projects and endeavors follow a similar process. There are variations, for example, the data may incite a new need that has not been recognized by the business understanding, which may still prove

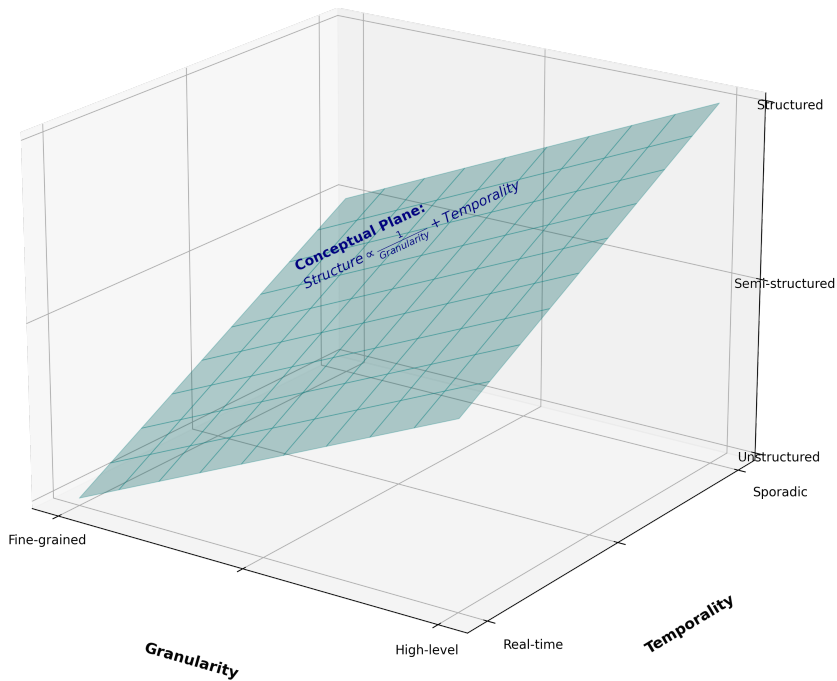


Figure 4. General data quality trade-offs

to be valuable and deployable. Likewise, the business understanding should not be understood in a narrow sense. In the scope of this research, business understanding can be translated into doctrines, military knowledge and standard operating procedures. As such, CRISP-DM functions as the ML framework through which the military problems are viewed. To transform the framework into a suitable assessment methodology it can be translated into

- Domain understanding: Identifying the problem areas within military operations and warfighting suitable for AI applications
- Data Understanding: Identifying the data, how it is accumulated and stored, what it enables and what is lacking
- Data preparation: What is required to formulate the data into usable, high-value format
- Modeling: What models from the hypothesis space seem applicable and what are the limitations
- Evaluation: How is the model evaluated, in which environments and how is the performance measured
- Deployment: What are the implications of deployment, including regulatory and ethical perspectives as well as hardware and personnel resource perspectives

This translated version is displayed in Figure 5, where the main difference is understanding the domain, i.e., warfighting, through doctrines, concepts, procedures, rules and regulations, capabilities, resources, caveats, and crucially, all the aforementioned aspects also from the perspective of the adversary and the adversarial impact.

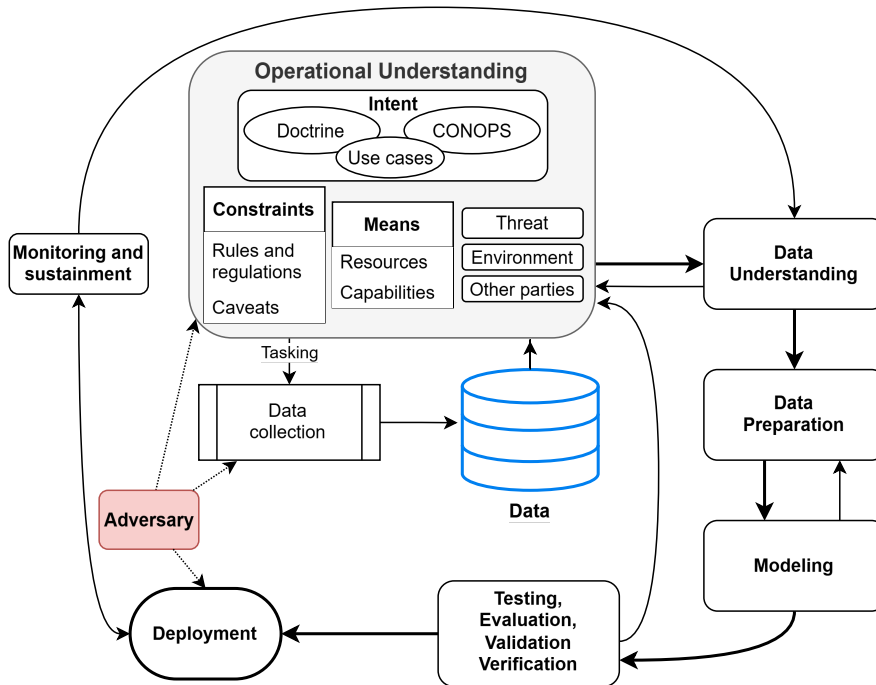


Figure 5. Military adaptation of the CRISP-DM

The technical details of AI and ML will be examined in further depth under respective sections for original publications. Now that the AI premise has been established, the following sections introduce the military domain and its specialties regarding the subject.

3.2 Military Domain

This section introduces the military domain in general, providing context, background information and "business understanding" for the synthesis, through which the arising problems and suitability of AI methods is assessed.

As a demonstrative use case it is established that there is a military group that operates in a designated operations area, tasked to produce intelligence and compile recognized operational picture from the area of responsibility with available assets, to maintain readiness to deter or repel adversary aggression, and to sustain operational

capability over the period of months in liaison with logistics support. By definition, a group consists of units, which have troops and platforms, such as infantry and armoured vehicles. This demonstrative example does not consider the particular military branch, but the group consists of headquarters and subordinate units that have a selection of sensors, effectors, and personnel that together create the capabilities available to the group. This use case is reflected in following subsections.

3.2.1 Tasks

The classic Roman maxim *si vis pacem, para bellum*, translated as “if you want peace, prepare for war”, captures the universal rationale for developing and maintaining a military force and capability [197]. Machiavelli advances the same logic regarding the political necessity of military readiness in *The Prince* and *The Art of War* [198; 199]. Another often cited theorist of war, Carl von Clausewitz, describes war as merely an extension of politics used to compel the enemy to do ones will [30], not neglecting the idea that if the will is to have peace, military power must be maintained to achieve or maintain it. These notions justify the existence and national function of military forces, where specific operational tasks are a concrete realization of the ways of executing this primary function of serving these national interests.

As stated by Wilén and Strömbom [200], contemporary roles and tasks of military forces comprise warfighting and irregular warfare, military assistance and international crisis management as well as aid to the nation which includes disaster relief, military support to internal security forces and, for example, epidemics support. While these tasks differ greatly in their purpose and execution, the military capabilities do not change drastically according to tasks. Instead, the orders and constraints regarding the use of said capabilities differ. Military operations and exercise do not exclusively exhibit warfighting functions, as peacetime operations may, for example, consist of patrols, training, and maintaining or building readiness in order to activate warfighting functions when necessary. Even in this case, warfighting capability is the main output, even if not used in full effect.

For the sake of this research, it is not applicable to assess the differences between tasks, as AI solutions can be technically suitable for any task if they serve a purpose for a given capability within the military. Therefore, the main focus of this research is in warfighting, but the findings and conclusions are not limited to this primary class of tasks for a military organization.

In the aforementioned demonstrative use case, the exemplary task is to monitor an area and maintain readiness to use force should the situation evolve negatively through adversarial actions. To accomplish this task, the headquarters must plan the use of subordinate units and their rotation in the operation, as well as ensure logistic support to replenish the units. Likewise, the group must compile intelligence information and situational awareness of the theatre to enable such planning and

proactive tasking of units and deployment of capabilities.

3.2.2 Capabilities

The term capability is not self-explanatory in a military context. As described by Finnish Admiral Anteroinen [201], the term can have several definitions depending on the context: it can refer to an effect, a function to execute tasks, physical weapon systems, platforms, fighting power, or a system model. In this research, the term capability reflects the ability to execute tasks with given systems and platforms, which ultimately generates the fighting power of a military force, holistically encompassing all auxiliary functions that contribute to that power.

While systems and platforms provide the physical means to execute a task, making them a fundamental component of a capability, they are not, by this definition, capabilities themselves. For example, a warship is a platform equipped with multiple systems, effectively operating as a system of systems. The platform itself is the asset; the capabilities it produces include, for example, the ability to compile a tactical picture from the air, sea surface, and subsurface domains, effectively creating a multi-domain surveillance capability. Likewise, the effectors on the platform enable the execution of air defence or surface warfare tasks, thereby providing air defence and surface warfare capabilities. As stated in the demonstrative use case definition, the subunits have and operate the assets, which create the capabilities available to the headquarters and the group as a whole. These capabilities can be extended beyond the hierarchical structure, for example in an effort to support another group, or vice versa in receiving support.

3.2.3 Organization

The military organizations consist of a hierarchical structure which differs from nation to nation and coalition. Fundamentally, there are high level headquarters under which lower echelons are established. The highest command level is usually the joint command that has service branches and other high level establishments like military intelligence under it. Under branches and their respective headquarters are the hierarchy trees of subordinate units which go all the way down to unit and platoon level.

The organization composition enables the military to have a rather pure rational-legal bureaucratic leadership structure described originally by Max Weber [202]. This bureaucratic structure creates a clear chain of command and division of effort, resulting in an effective, commander-led structure that is, in theory, nimble and adaptive as the decision-making is concentrated to the officer in command. While other leadership systems have shown promise that argues towards distributed leadership, the clear chain of command is straightforward and, when effective, the most decisive

way to lead and manage military forces.

Each level of the military organization is characterized by the scale it operates in. The levels are usually described as tactical, operational and strategic [203; 204]. The concept of dividing war into these levels has historical roots in the Napoleonic Wars and the American Civil War, was formally developed by Prussian and Soviet military theorists [205], and was formally adopted into U.S. doctrine in 1982 via Field Manual 100-5 [206]. The tactical can be underlined with "technical" level, meaning individual unit maneuvers and actions. In the lower echelons of the hierarchy, i.e., tactical level, the tempo is higher and the time span of decision-making shorter. Essentially, for example in a platoon or unit level, the operational decisions are done in a time frame from seconds to hours, while the unit headquarters plan and execute in a time span of days or weeks. Likewise, the scope of impact for the actions that result from the decisions have near immediate effect. In higher echelons, the actions may have effects over the course of years to come.

As an example, modern U.S. Army doctrine refines this traditional three-tier model into four distinct levels of warfare: the national strategic, theater strategic, operational, and tactical levels [207]. This updated framework serves to clarify the relationship between broad national objectives, the operational approach, and the execution of tactical tasks. At the highest echelon, the national strategic level involves the government formulating policy goals and global strategy using all instruments of national power. Subordinate to this, the theater strategic level focuses on combatant commanders synchronizing activities to fulfill those policy aims within an assigned region. The operational level acts as the vital link between these strategic goals and tactical force employment, where campaigns and major operations are planned, conducted, and sustained over broader aspects of time and space. Finally, the tactical level is where forces directly plan and execute battles and engagements to achieve assigned military objectives. While tactical actions at the corps or division level might span days or months in the form of battles, lower-echelon engagements executed by brigades and below are typically resolved in minutes or hours, reflecting the immediate, high-tempo impact characteristic of the lowest levels of the hierarchy [207].

Referring back to the demonstrative use case, the headquarters needs to plan days ahead, while the operational units act in real-time while executing the tasks that span from hours to days.

3.2.4 Decision-making processes

The organizational decision-making processes are denoted as Military Decision-Making Process (MDMP)s, which often follow a similar structure. MDMP has been defined by the US Army Colonel Mueller, director of the Center of Army Lessons Learned, as "a systematic process that enables commanders and their staffs to apply

critical and creative thinking and doctrine to solve problems and establish the framework and conditions for commanders to make effective decisions” [208]. Similar process structure as in [208] has been highlighted in NATO APP-28 [209], which also references several other MDMPs that are very much alike. The key points to highlight are:

- Receipt of Mission: What is to be done, what is the task
- Analysis of the situation: What are the resources, what is the adversary, what is the environmental impact
- Development of Courses of Action: Creating alternatives for decision-making
- Evaluating the alternatives: War gaming to decide for the best course of action identified
- Orders and Execution

When compared to CRISP-DM, the MDMP is, on abstract level, very similar. The business understanding evokes a need to understand the data. Then, the data is preprocessed, a model is fitted, alternatives are evaluated and finally the end result is deployed. However, when considering AI solutions, CRISP-DM can be considered for each of these MDMP phases separately: Can AI aid in analyzing the mission, the situation, to develop courses of action, evaluate alternatives and enhance execution?

A decision-making framework that is widely adopted in the military and scales from individual fighters to staff headquarters has been proposed by US Colonel Boyd and goes by the name of Observe-Orient-Decide-Act (OODA) loop [210; 211]. Colonel Boyd used OODA to describe the decision-making process of an individual or a group of individuals. This framework provides a more ground-level approach to the application of AI when individual operators and warfighters are examined. Just as with the MDMP, the same questions can be placed upon OODA: Can AI be used to aid in observing the environment? Can it aid the orientation, where the individual synthesizes the observed data against his or hers knowledge base and experience? Can the decision be better informed, faster, more concise? The Chapter 4 answers these question from the point of view of each original research paper.

It is to be noted that the OODA loop has given ground for a simplistic interpretation that the speed is the primary characteristic to achieve military superiority, but this point of view is narrows the focus of larger scale warfighting and has garnered criticism [212]. While speed and closing the kill chain are key factors in tactical success, the overall success of an operation or military campaign relies on the quality and scope of the decision made.

In the use case example, headquarters employ a process like APP-28 or MDMP, while the units in the area of operation execute their own OODA loop in real-time, under the constraints and autonomy issued by headquarters. For example, the headquarters have activated certain rules of engagement, which dictate the way the units are to respond to different situations. Usually in low-risk phases of an operation the

use of force is allowed only in self-defence for all units, and the headquarters retain the authorization in all other scenarios. When the situation evolves, the headquarters may delegate this authority to the units as well, with new limitations such as a list of accepted targets that can be engaged without further notice. Assessment of the situation and tasking the units is done via the decision-making process, which collects the information from the units to develop and evaluate possible courses of action to choose the best execution from.

3.2.5 Military Information Systems

Despite the proliferation of information technology, information systems do not have a succinct, singular definition or a concept, although there are widely recognized and utilized definitions. Checkland and Holwell [213] have proposed general concepts of information systems, of which a combined interpretation is displayed in Figure 6. The key distinction between an IS and an Information Technology (IT) system is that IT systems consist of hardware, software and networks, while IS includes humans in the system view.

Essentially, applying the systems methodology of Checkland, there are elements that lead to actions and an IS. The IS, in this case, is the system which serves, or supports, the system that executes actions. The actors that execute actions have information needs that need to be met to perform *purposeful* actions and inflict changes in the elements. This process can be viewed as a generalized, schematic interpretation of any IS, encompassing also IT systems in it. Military information systems can be examined from the perspective of this particular concept.

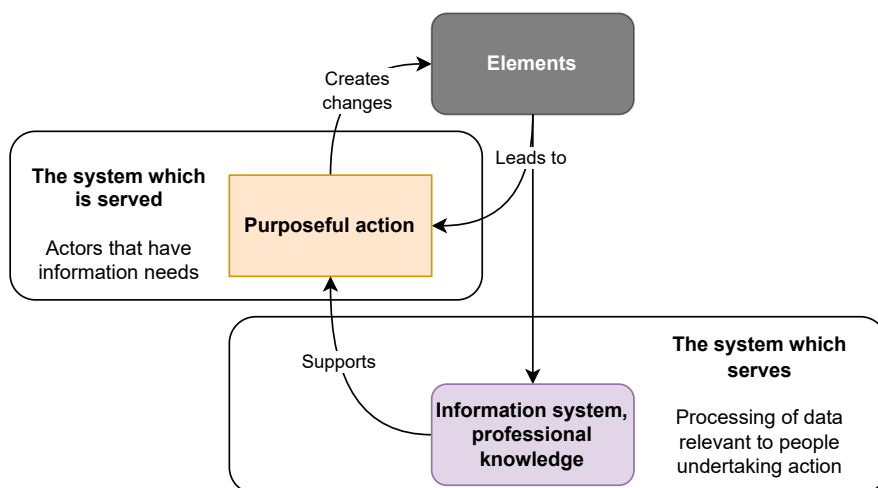


Figure 6. Information system concept [213]

The fundamental nature of military operations and related IS have undergone a profound transformation over the past several decades, shifting from platform-centric warfare, where individual tanks, ships, or aircraft operated as largely independent and isolated entities, to network-centric warfare which is defined, e.g., as an "information superiority-enabled concept of operations that generates increased combat power by networking sensors, decision makers, and shooters to achieve shared awareness, increased speed of command, higher tempo of operations, greater lethality, increased survivability, and a degree of self-synchronization" [214]. In this modern operational paradigm, the decisive advantage on the battlefield is no longer derived solely from kinetic mass, armor thickness, or raw firepower. Instead, superiority is linked to information dominance.

As stated above, one of the most widely recognized concepts underpinning military information systems is the OODA loop [210]. It has seemingly transcended its origins as a cognitive model for individual pilots and been scaled up to encompass entire military organizations, automated defensive grids, and global sensor networks. In essence, it can be stated that

- Observation involves the collection of System-of-Source (SoS) data, such as raw or lightly processed, mainly unfiltered data from the operational environment via sensors, as well as System-of-Record (SoR) data such as human-generated or automated reports and messages that provide an insight into the environment and its actors.
- Orientation requires the synthesis of the newly observed data with prior knowledge, context, historical intelligence, and strategic objectives to form a coherent, accurate operational picture of the battlespace.¹
- Decision involves selecting a specific course of action based on this orientation.
- Action is the execution of that choice, usually physically or electronically.

This view of the military information system is depicted in Figure 7, where the domains exist within the contested environment. Objects, elements and phenomena denote artifacts within the environment that can be observed, thus creating the information flow through the observation layer. The observers and data sources consist of units that encompass sensors and troops, denoting humans, as well as internet sources, fixed sensors from radars to weather stations, and third parties such as partners and co-operators. The observations and resulting reporting that is based on initial observations creates SoS and SoR data that exists in a variety of types and qualities, from structured to unstructured, real-time to sporadic, fine-grained to high-level in Figure 4. This data is then processed, aggregated and displayed for C2 purposes such as maintaining the operational picture to monitor and guide the operation.

¹ A combined, favorably domain-crossing picture is known as Common Operating Picture (COP).

Operating requires that the C2 is linked back to the units in order to react to changes within the environment. Overall, the aggregation and processing of data enables orientation, where the surrounding context is combined with the observations to create situational awareness and understanding, which gives grounds for further analysis, automated or manual, and decision-making that results in courses of action, plans and orders that are promulgated again to the units to have a desired effect on the environment and the objects or elements within. The actions ought to create changes that then result in novel observations, and the loop is reiterated.

On the right side of Figure 7 is a depiction of hardware regarding the IT components of the system, although technology resides in the IS part as well. In the hardware, there are certain capabilities at the level which executes actions, either kinetic or otherwise, requiring those capabilities and usually some form of sensors, mobility, local computing and a power source. On the supporting level, which serves the actors, the emphasis is on information processing in contrast to mechanical equipment and real-world effects. The connectivity requirement, consisting of radios, cables, satellites, and fiber optics as well as all the other related hardware, stretches from the acting front to the supporting layer to ensure the information flow.

Despite the simplified view in Figure 7, military information systems do not operate in a flat, decentralized hierarchy. Instead, they are strictly organized around the aforementioned three echelons of warfare: Strategic, Operational, and Tactical [204]. Each echelon requires fundamentally different types of information, processed at different speeds, and presented at vastly different levels of granularity. Additionally, the more granular hierarchy of units, formations, and commands dictates how this information physically and structurally flows. To execute operations across these levels, forces are organized, for example, into a standardized unit echelon hierarchy as displayed in US Army Field Manual 3-0 [207]. In the field manual, strategic objectives are managed by Theater Armies or Joint Commands, which pass operational directives to Field Armies and Corps. These operational echelons then translate campaigns into actionable missions for tactical formations, cascading from Divisions and Brigades down to Battalions, Companies, and individual Platoons. Consequently, an information system at an operational headquarters will aggregate, filter, and transmit long-term logistical and campaign data at different scope and scale than a system utilized by a tactical company or a platoon utilizing and requiring real-time targeting data.

Additionally, the military IS comprises multiple information security environments [215], which are not displayed in Figure 7. Namely, the environments or information classification levels are Unclassified, Restricted, Confidential, Secret, Top Secret and, in some instances, Cosmic Top Secret [215]. In practice, this means that depending on the gravity of information, it has to be stored and used within a proper environment to ensure that is not disclosed in an uncontrolled manner. In the Figure 7, the whole IS can exist within one classification, or several, depending on

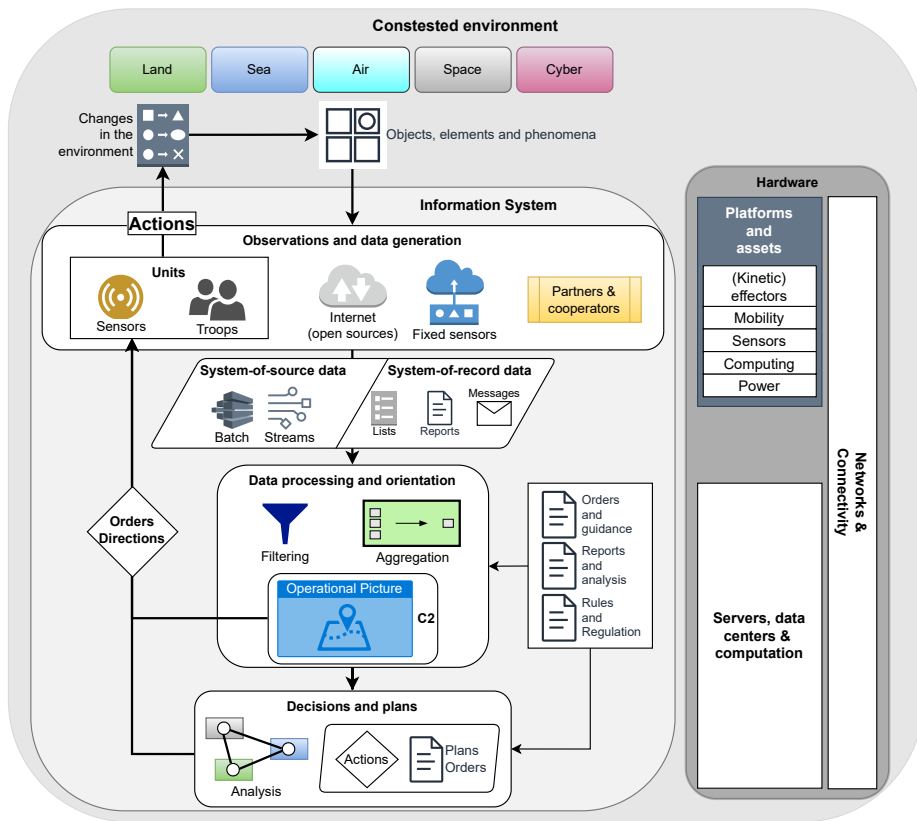


Figure 7. Military information system view

the gravity of the information.

As such, when military information systems are mapped to operational echelons and actual organizational hierarchies, the overarching architecture becomes a nested structure of codependent ISs in different physical and digital environments. As with the OODA loop, each individual entity can be viewed as its own IS, whether it is a single tactical unit, a field army, or a maritime component. At every level, there are elements that provide information, which requires processing to support purposeful actions. For example, a military group consists of subordinate units, which are themselves separate ISs that feed information upward to the group headquarters—a higher-echelon IS. In turn, the subordinate units receive the necessary command guidance and support to conduct their purposeful actions, creating localized changes within the environment in order to meet the broader goals of the operation.

3.2.6 On Complexity

It has been now established that each level of a military organization can be considered a system with inputs and outputs. At the lowest level, the inputs are real-world events that are monitored with available sensors and systems. The data sources include sensors that survey some wave length in the electromagnetic spectrum, from infrared to communication frequencies, as well as ISs that process data, nowadays mainly in the digital domain. The SoS and SoR data can be raw or processed, structured or unstructured, and it provides its primary user insight into a particular, narrow problem while contributing to the bigger picture. This information is processed into an output to higher echelons, which gather and aggregate the data, analyze it, review their mission objectives with regard to the data and either give their outputs as actions to the lower echelons or inputs to higher echelon decision-making.

The amount of data, the complexity and the uncertainties increase when moving from tactical to strategic levels. When the number of variables increases simultaneously with the time span of planning and execution, the complexity is guaranteed to increase exponentially. Additionally, the military landscape is characterized by the incompleteness of data and only partial observability into the opposing forces capabilities, composition and intent, which create uncertainty in the development of the situation towards the desired end state.

The complexity of military operations is governed by what is computationally known as the 'curse of dimensionality.' As the number of, e.g., units, weapon systems, and environmental factors increase, the possible states of the battlefield and the available actions scale exponentially. Furthermore, unexpected events that are often referred to as the 'fog of war' introduce severe stochasticity. While these factors are described conceptually here, they are formally quantified and mapped to mathematical state spaces and probability distributions later in Section 4.2, where they form the baseline for applying RL to military domain.

In summary, a military organization is created in a way that the units operating or producing military capabilities provide headquarters with information. The information functions as the input for decision-making, analyzed in a suitable decision-making process. The decision-making is an analytical workflow that results in orders for actions to be executed. The execution then serves a goal that is related to the particular task, such as warfighting or crisis management. Data plays a critical role through the military organization, but as demonstrated later in Section 4.2, the scale of complexity in decision-space calls for abstraction and heuristics, sometimes a rule-of-thumb, to be able to make decisions in due time. This has been recognized before, as the ability to grasp the essential is one of the key principles of, for example, Alexander Suvorov [37]. This underlines the next subsections that examine the data in military domain as well as identified application areas.

3.3 Domain features of data

The data is a difficult subject in the military domain. It is simultaneously both abundant and scarce. There are many factors behind this, which relate mainly to sensitivity and security, organizational culture and doctrines.

It is self-evident that a lot of military data is classified to keep, for example, capability and performance information secure from adversaries [215]. A lot of military systems are air-gapped [216; 217; 218], which means that they are neither directly or indirectly connected to the internet. This approach secures critical systems such as C2 systems from cyber attacks and data leakage. Simultaneously, it prevents collecting the data in the same manner that can be done for public, commercial and open source data, and requires considerably more complex architectures and integration to get the data from these systems into a database that can be used to train AI models. Combining the data into a large, all-encompassing data set might not be wanted at all, as distributed and fragmented data prevents potential adversaries from gaining the whole picture despite receiving some parts of it.

As a result from the fragmentation, the data is often in silos, and reaching a holistic understanding of what data is available, where, and how is very difficult. This is highlighted by Brose [163], stating that "platforms rarely cohere into one battle network that can share information effectively", quoting one US military officer saying that "The main problem is that none of my things can talk to each other". The CEO of Anduril, a US company focused on drones and autonomous systems, has stated as an industry observation in December 2024 that "Exabytes of defense data, indispensable for AI training and inferencing, are currently evaporating" [219]. These points highlight the fact that while military forces generate and process vast quantities of data, it is used ephemerally for a singular use case. If it is stored, the operator insight that occurred is usually lost, as the systems do not support recording it or that is not part of the *modus operandi*. Likewise, if analyzed data is stored with its meta data,

the raw data might be not, which again may have an impact on the usability. If there is no ground truth the supposedly iterative process between data understanding and data preparation is limited and backtracking to the source is not possible.

Additionally, as a difference when comparing with civilian domain, the military environment is contested in a different sense. In industry, data security is critical and enterprise espionage is a factor, but in the military side these aspects are amplified, as the stakeholders are governmental, and national security is at stake. Peacetime environment is contested in a different sense, as methods are more subtle than during open conflict, but may still include for example adversarial actions such as data poisoning [220; 221]. Other things such as jamming, information falsification and noise, in different mediums, may also inflict both the accumulation and exploitation of data. EW aims, in short, to maintain own C2 capabilities while breaking the connectivity of the adversary. EW deception operations aim to decrease the accuracy of adversary's intelligence gathering, surveillance, target acquisition and reconnaissance [222]. Referring back to data and its qualities, the completeness is likely to be effected both in peace and wartime. A widely known electronic warfare method to precede an aerial operation is to gradually increase the background noise with suitable jammer platforms to increase the thresholds of surveillance radars: in time, this will lead to increased detection thresholds, which enable the attacker to get closer before being detected [223]. A similar approach can be applied to other operations and data types, for example exercises and general activity. As an example, an adversary could operate in different manners, publish misleading doctrines, or use different camouflage equipment to prevent the accumulation of data for accurate ML purposes. Therefore, the data is greatly different from, for example, medical data that is collected *in vitro* and more accurately describes, e.g., the state of health of a certain populace, as both the medical institutions and the populace usually share the same goal and do not possess an adversarial stance.

Being part of a process instead being a capability or an enabler in itself underlines the cultural and doctrinal issues towards data. Data is mainly seen as an ephemeral input that results in an equally ephemeral output in a reactive system. This does not mean that insights are not drawn from the data, but that there are limited resources and limited capabilities to effectively store the data and find use for it beyond the current input-output loop. While modern technology giants have built their entire business models on collecting data, platform providers such as Amazon, Meta and Google [224; 225], modern militaries have yet to integrate the understanding and utilization of data into their operational procedures and processes.

3.4 Main application areas

Building upon the architectural foundation of military information systems established in Subsection 3.2.5 and the data constraints outlined in Section 3.3, the inte-

gration of AI can now be mapped into this ecosystem. While the previous sections illustrated the holistic, multi-echelon complexity of military operations, applying AI requires abstracting these systems into a functional pipeline.

From a dichotomic perspective, AI can be viewed to either enhance current processes and workflows or to create entirely new ones [226; 227; 228]. Enhancing current processes and workflows is usually easier, as it is simpler to analyze a process, find the issues that can be improved, apply changes, and evaluate the result, compared to innovating a completely new process or a meaningful workflow. It is also noteworthy that AI is still very much a tool instead of a source of innovation, despite advancing inference and reasoning capabilities [229]. Hence, it is up to humans to come up with the new processes suitable to be executed with AI methods, or the processes that can be disregarded due to the capabilities that can be induced with AI.

The aforementioned frameworks, MDMP and OODA-loop can be used to estimate the main application areas for AI methods within the military context. Both of these frameworks or processes can be abstracted into a process shown in Figure 8. Essentially, there is information regarding the physical world, which is observed with either biological senses or sensors. This information covers thermal, mechanical, chemical, magnetic and electromagnetic [230; 231], of which the electromagnetic is most significant in the military sense: radars, optoelectronics and signal intelligence utilize electromagnetic transmitters and receivers, and so do communication devices. The importance has been highlighted by accredited military leaders, as the control and management of electromagnetic spectrum has been highlighted as a the key component of victory in modern warfare [232; 233].

In addition to the information received directly from the physical world exists the digital world which includes technically all the rest of the available information, from documents to the sensory information. The digital information includes knowledge and models of the physical world which enable analyzing the raw sensory data to greater extent. Likewise, orders, regulations, guidelines, plans and doctrines exist in the digital domain, contributing to the understanding of the world, both digital and physical.

For both, the OODA-loop and MDMPs, there is a task that needs to be executed, and it usually requires having an effect on the physical or digital domain, or both, and in order to achieve the desired effect, the understanding of the task requires understanding from both domains. Therefore, based on Figure 8, there are five areas that can be viewed separately:

1. Collecting and accumulating observations.
2. Preprocessing and analyzing the data.
3. Analyzing and synthesizing the situation.
4. Evaluation of alternatives and finalizing decision.

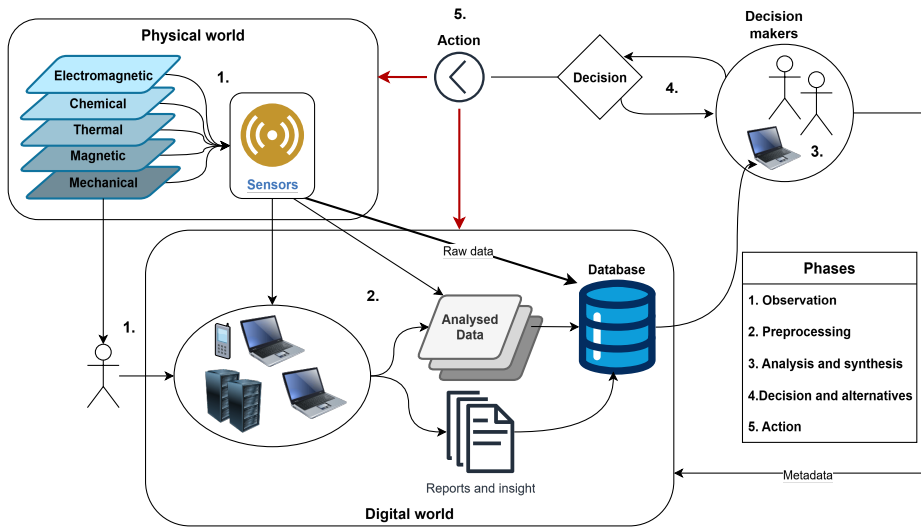


Figure 8. Information flow and decision-making

5. Executable actions that effect the desired domain.

This list is not exhaustive, but serves as a high-level abstraction of a generalizable information handling process that can be analyzed from the perspective of AI applicability.

The accumulation of information through observations happens individually both in the digital and physical worlds, and physical world observations are in general digitalized for processing and communication. It is also notable that through technology the digital and physical domains are converging, as the amount of digital data increases and for example digital twins [234] aim to bridge the gap. Simultaneously, the independent importance of digital domain has increased, due to digitalization and the amount of information. Therefore, in Figure 8, the decided actions can have effects in both domains, aggregating cyber and information warfare [235; 236] capabilities.

After and while information is being accumulated, it goes through pre-processing and post-processing before influencing decision-making. In this case, the pre-processing transforms the data into a desired, usable format that enables further processing and analysis. The post-processing, i.e., analysis, can be small-scale individual analysis executed by a sole operator with a narrow focus, or large-scale data analysis with vast amounts of data from different sources. The so-called big data analysis [237] is required to handle large quantities of data to produce impact in compressed time. Big data analysis requires multiple technologies from distributed computing to AI and ML methods to find patterns, identify anomalies and trends as well as perform predictions.

Once the data has been analyzed, conclusions can be drawn from it. These conclusions are used to induce a decision between alternatives, and it includes the details of the chosen course of action. In the military context, the decision can be a short-term individual decision by a single warfighter or fighter pilot, or a long-term strategic decision of a joint commander. The lowest level is usually denoted as technical, under tactical, and it includes the technical and immediate actions. The complexity of decisions increases with the level, as the number of effecting features, phenomena, uncertainties and variables increases both due to the level and the time-span of the decision.

From AI perspective, as partially shown in chapter 2, AI can be leveraged in each phase, from observation to action. Using the same sources as in chapter 2, the observations, from urban destruction [82] to chemical agent identification [80] can be enhanced with ML, automating the preprocessing phase and enabling faster analysis and synthesis. The analysis can be tasked to an intelligent DSS, which can then support tasks from troop deployment [58] to outcome prediction [59]. The resulting actions, from weapon assignment [115] to effect evaluation [110] are also in the scope of AI and ML applications. Hence, it can be stated that AI and specifically ML has a proven potential to enhance individual, narrow tasks, through all the phases of decision-making and information flow, on all levels of a military organization. Instead of technological hindrance, the constraints to the use of AI are the data, available computation, and expertise, both technical and military.

4 Machine Learning research areas

This chapter provides the necessary technical and theoretical background for the key ML paradigms addressed in this dissertation. By establishing the foundations of CV, RL, and other fields, it aims to provide a common framework for understanding the contributions of the six original publications, which are then synthesized in Chapter 6.

As explained in other words in Chapter 1, a basic ML model can be expressed as a mapping function f_θ ,

$$\hat{y} = f_\theta(x), \text{ e.g. } \hat{y} = w^T x + b \text{ in linear models,} \quad (3)$$

where θ denotes the parameters of function, or model, f , which in the linear case simplifies to the feature or weight vector w that maps the input x to the predicted output \hat{y} , corrected with the bias or intercept term b . Essentially, as a generalization, ML is always learning an approximation from inputs to outputs and then using the resulting model, f_θ or w in Equation 3, for inference.

However, the ability of a model to learn from a finite dataset and make accurate predictions on unseen data is not guaranteed. The theoretical framework that answers whether learning is feasible was largely established by the seminal work of Vapnik and Chervonenkis on statistical learning theory [238; 239]. As later explained by Abu-Mostafa et al. [240], the goal is to ensure that the error a model makes on the training data, the in-sample error (E_{in}), is a good proxy for the error it will make on future data, the out-of-sample error (E_{out}).

Probabilistic tools, such as Hoeffding's inequality [241], show that for a single, fixed hypothesis, E_{in} will likely be close to E_{out} under the assumption that there is enough data. Machine learning algorithms, however, do not test a single hypothesis; they search through an entire family of functions, also known as the hypothesis set \mathcal{H} , to find the one $h \in \mathcal{H}$ that best minimizes the error.

This is where the groundbreaking concept of the Vapnik-Chervonenkis (VC) dimension becomes essential [238]. The VC dimension, denoted d_{VC} , measures the capacity or expressive power of a hypothesis set. It quantifies the model's ability to split, or shatter, data points into all possible dichotomies. A model with a finite d_{VC} can be proven to generalize. The VC generalization bound, as depicted by Abu-Mostafa et al. [240], formalizes this relationship:

$$E_{\text{out}} \leq E_{\text{in}} + \sqrt{\frac{8}{N} \ln \frac{4m_{\mathcal{H}}(2N)}{\delta}} \quad (4)$$

where N is the number of data points, δ is the probability that the bound fails, and $m_{\mathcal{H}}$ is the growth function, which is bounded by a polynomial in N if the VC dimension is finite. This bound reveals the fundamental trade-off in ML:

- A more complex model with a higher d_{VC} can achieve a lower E_{in} but faces a larger penalty term, increasing the risk that E_{out} will be high. This is known as overfitting.
- A simpler model with a lower d_{VC} has a smaller penalty and generalizes better but may be too simple to capture the underlying patterns, resulting in a high E_{in} . This is known as underfitting.

This insight leads to the principle of Structural Risk Minimization (SRM) [238], a formal strategy for selecting a model that balances low training error with controlled model complexity to achieve the best possible out-of-sample performance. This theoretical foundation justifies the entire learning process and informs practical techniques like regularization, which implicitly penalize model complexity.

All ML requires three fundamental components:

1. Data,
2. Model,
3. Loss or Objective,

and the distinctions between different paradigms result from variations around these three components.

A key component that enables updating the model with regard to the received loss is the update mechanism, a rule or an optimizer. While updates can be random deviations from which the most suitable is picked for the next iteration, the most widely spread and effective method relies on gradient-based optimization. This approach calculates the gradient of the loss function with respect to each weight parameter.

Formally, if a model is parameterized by weights θ and its performance is evaluated by a loss function $L(\theta)$, the goal is to find the parameters that minimize this loss. In standard gradient descent, the weights are iteratively updated in the direction of the negative gradient, i.e., the direction of steepest descent:

$$\theta_{t+1} = \theta_t - \eta \nabla L(\theta_t), \quad (5)$$

where θ_t represents the parameters at iteration t , $\nabla L(\theta_t)$ is the gradient of the loss function with respect to those parameters, and η is the learning rate, a hyperparameter that controls the step size of the update [5; 28; 6].

This fundamental rule ensures that the weights are updated towards the gradient-indicated direction of decreasing error in the outputs. More advanced optimizers, such as Adam [242] or RMSprop [243], build upon this foundational principle by introducing adaptive learning rates or momentum to accelerate and stabilize convergence, but the underlying mathematical premise remains the same.

Supervised learning is the most straightforward approach, where the data set D has annotations, i.e., $[x_i, y_i] \in D$ and the loss calculation is straightforward between the predictions of the trained model and the known labels (annotations) in the data. A classic example is linear regression [244; 245], which can be assumed to have an input vector $x = [z_1, x_2, \dots, x_p]$ that is to be predicted into a real-valued output y . The linear regression model has the form

$$f(X) = \beta_0 + \sum_{j=1}^p X_j \beta_j, \quad (6)$$

where β_j are unknown parameters or coefficients. Unsurprisingly, linear regression closely resembles Equation 3. Supervised learning approach and more sophisticated algorithms are introduced in more detail within the field of CV in Section 4.1.

If there are no annotations and the data is simply $x_i \in D$, unsupervised methods are applicable to gain understanding of the data, such as identifying underlying patterns and doing component analysis to identify most impactful features [246]. A classic example of an unsupervised method is k -means [247], which starts with a preselected number of groups, k , that are represented by a single, random data point in each. Then, iteratively, a new data point is added to a group according to the closest proximity of the new point and the mean of previous data points within the groups, after which the mean of the group is adjusted accordingly. In the seminal work by MacQueen [247], for a given set of k centers, where $x \in (x_1, x_2, \dots, x_k)$, the region of points closest to center x_i is defined as

$$T_i(x) = [\xi : \xi \in E_N, |\xi - x_i| \leq |\xi - x_j|, j = 1, 2, \dots, k], \quad (7)$$

where $T_i(x)$ is known as the Voronoi cell [248] for the center x_i . These regions are used to build the final partition sets in a sequential manner as

$$S_k(x) = T_k(x)S_1'(x)S_2'(x) \dots S_{k-1}'(x), \quad (8)$$

so that the final partition $S(x)$ has the property that every point in a cluster $S_i(x)$ is closer to its center than the other established centers. Independently, a similar approach was proposed by Lloyd [249] at Bell labs, where the repeating steps of the iterative process were the assignment displayed in Equation 9 and update displayed in Equation 10.

$$S_i^{(t)} = \left\{ x_p : i = \arg \min_{j \in \{1, \dots, k\}} \|x_p - \mu_j^{(t-1)}\|^2 \right\} \quad (9)$$

$$\mu_i^{(t)} = \frac{1}{|S_i^{(t)}|} \sum_{x_p \in S_i^{(t)}} x_p. \quad (10)$$

Lloyd’s algorithm works by repeatedly applying these two steps until the partition and centroids no longer change. The modern formulation synthesizes MacQueen’s property and Lloyd’s procedure by framing K-means as an optimization problem. The goal is to find the partition S that minimizes the Within-Cluster Sum of Squares (WCSS), also known as inertia. The modern notation for the optimized objective function to minimize the sum of distances over points for all partitions can be stated as

$$\min_S J = \sum_{i=1}^k \sum_{x \in S_i} \|x - \mu_i\|^2. \quad (11)$$

Another widely adopted and useful unsupervised method is feature reduction through, for example, Principal Component Analysis (PCA), for which the mathematical background and formal name were introduced by Pearson [250] and Hotelling [251] approximately a century ago. PCA can be used to identify most relevant features in the data to reduce computational complexity for succeeding training tasks or to provide other analysis of the underlying implications.

While unsupervised learning is effective and useful for, e.g., large scale data analysis, this research does not cover unsupervised tasks as a primary research problem, although the methods are affected by FL, introduced in Section 4.3.

If data is not available at all, it can be in some instances generated or gathered during training. This applies to RL, elaborated in Section 4.2, as it relies on an exploratory agent that interacts with an environment, usually virtual, and by accumulating experience of interactions and having a reward mechanism instead of a loss function learns to predict most favorable choice of action from one state to the next [27]. While data may not be required to start train an RL algorithm, it still has to generate data on the go. Additionally, data is usually required to formulate the model or simulation that is used to generate the data.

As described in Section 3.1, evaluation is a critical part of the ML development. It relies on quantitative criteria such as F1-score [252], Receiver Operating Characteristic (ROC) curve [253; 254], Area Under Curve (AUC) [255], and qualitative criteria consisting of, e.g., interpretability assessment.

To evaluate the quantitative performance of classification models, the ROC and AUC are utilized as robust performance metrics that require formal definition.

The ROC curve [253; 254] illustrates the diagnostic ability of a binary classifier system as its discrimination threshold, τ , is varied. It is created by plotting the True Positive Rate (TPR) against the False Positive Rate (FPR) across all possible threshold settings. The TPR, also known as sensitivity or recall, represents the proportion of actual positives that are correctly identified, and is defined as:

$$TPR = \frac{TP}{TP + FN}, \quad (12)$$

where TP is the number of true positives and FN is the number of false negatives. The FPR, also known as the fall-out or the probability of false alarm, represents the proportion of actual negatives that are incorrectly identified as positives, and is defined as:

$$FPR = \frac{FP}{FP + TN}, \quad (13)$$

where FP is the number of false positives and TN is the number of true negatives.

While the ROC curve provides a graphical representation of the trade-off between sensitivity and specificity, the AUC [255] provides an aggregate measure of performance across all possible classification thresholds. Mathematically, it is the two-dimensional area underneath the entire ROC curve from $(0, 0)$ to $(1, 1)$, computed as the integral of the TPR with respect to the FPR:

$$AUC = \int_0^1 TPR(FPR^{-1}(x)) dx \quad (14)$$

An AUC value ranges from 0.0 to 1.0. An AUC of 1.0 represents a perfect model, while an AUC of 0.5 denotes a model performing no better than random guessing. In probabilistic terms, the AUC value represents the probability that the classifier will rank a randomly chosen positive instance higher than a randomly chosen negative one [255]. This makes it an especially valuable metric in military AI applications, such as target recognition or anomaly detection, where the cost of false positives (false alarms) and false negatives (missed threats) must be carefully weighed.

As a well known ML metric, the F1-score was first introduced as a composite E-measure by Van Rijsbergen [252], which is denoted as

$$E = 1 - \frac{1}{\alpha(\frac{1}{P}) + (1 - \alpha)(\frac{1}{R})}, \quad (15)$$

where P is precision, calculated as $\frac{|A \cap B|}{|B|}$, where A are relevant samples and B are retrieved samples. Together, the intersection $|A \cap B|$ is the set of true positives. R is recall, defined as $\frac{|A \cap B|}{|A|}$, and $\alpha \in [0, 1]$ denotes the composite weighting between precision and recall; in Van Rijsbergen [252], it is denoted as $\frac{1}{2}$ for equal weight.

In modern terms, the precision is the number of true positives divided by the sum of all positives, true and false, while recall is the true positives divided by the sum of true positives and false negatives. The precision depicts the model's accuracy in positive predictions, while recall measures how well the model is able to detect all the positive instances. The E-measure version known as F1-score combines both of these measurements into one evaluation measure, modernly formulated as

$$F1 = 2 \times \frac{P \times R}{P + R}. \quad (16)$$

The precision, recall and F1-score have originated from binary classification or prediction tasks, but can be scaled to more complex scenarios by, e.g., averaging over classes.

In the following sections, each ML research area from CV to FL with respective original publication is explored in depth. Each section starts with an introduction to the research area, after which the publication is reviewed to form an analytical insight to the research problem of this dissertation. It is noticeable that the separation between paradigms is fluid, as different model architectures and approaches afflict several paradigms, usually in a manner where the novel solutions are discovered under one paradigm and extended to others. Hence, the separation is some what arbitrary and artificial, but correlates well with the original publications as such.

4.1 Computer Vision background

The genesis of CV as a field can be traced back to the 1950s and 1960s. Early research drew inspiration from neurobiology, notably the work of David Hubel and Torsten Wiesel [256; 257] on the mammalian visual cortex, which revealed a hierarchical structure of neurons responsible for detecting edges and orientation. In 1963, Lawrence G. Roberts' Ph.D. thesis [258], often cited as a pioneering work, showed how to derive 3D information about a "blocks world" from 2D images. CV was largely based on 3-dimensional projective geometry, with hand-crafted features constructed, with for example edge detection [259] or Local Binary Patterns (LBP) [260], and then used as inputs to simple learning algorithms [261], accompanied by feature detection and extraction was dominated by handcrafted feature descriptors that were designed and tuned to be invariant to changes in scale, rotation and illumination, such as SIFT [262] and SURF [263].

The fundamental goal of computer vision is to extract meaningful information from visual data. This separates it from image processing, where the image is edited and modified. CV process is typically structured as a sequence of steps:

- Image acquisition, which happens with a sensor such as a camera, converts sensory data from physical world into a numerical representation as a grid of pixel values.

- Image processing, where low-level processing techniques are applied to the raw pixel data to prepare it for analysis
- Feature detection and extraction, a crucial step to identify salient patterns or points of interest in the image.

Core tasks in CV include image classification, object detection, image segmentation, caption generation, synthesis, inpainting, style transfer, super-resolution, depth prediction and scene reconstruction [173]. In this thesis, the focus is on classification and object detection. In caption generation a caption is automatically generated for an image, meaning that the model has been trained with data consisting of images and their captions. Synthesis refers to generation of new images, which is discussed later, while inpainting is an image editing application to, for example, remove unwanted objects. Style transfer refers to converting, for example, a photograph to an oil painting, while super-resolution improves the image resolution by increasing the number of pixels with generation. Depth prediction predicts the distance from the camera to the objects from one or more views, and scene reconstruction creates an additional dimension to an image, for example from black and white to color.

In classification, a fundamental task, whole images are assigned a single label from a set of categories, with "cat" and "dog" as classical examples. Object detection is a more complex task where the goal is to detect objects and their location within an image, which can then be classified. Image segmentation can be divided into semantic segmentation, where each pixel of the image is classified as belonging to a particular category, such as the "cat", or instance segmentation where each instance of the same object class is distinguished.

The field was one of the first to be greatly transformed by modern DL methods, predominantly using the Convolutional Neural Networks (CNN) architecture [264]. CNNs are a class of neural networks specifically designed for processing grid-like data, such as images [6; 173]. Key components of a CNN include the namesake, convolutional layers, which apply a convolution operation, a special linear operation that replaces the general matrix multiplication of linear neural networks. It is usually denoted with an asterisk, as in $s(x) = (x * w)$ [6]. The convolution is usually used over more than one axis in the grid, which means that for two-dimensional image I , a two-dimensional convolution kernel K results in formulation

$$S(i, j) = (I * K)(i, j) = \sum_m \sum_n I(m, n)K(i - m, j - n), \quad (17)$$

where I is the input image with grid coordinates (i, j) and kernel K has dimensions (m, n) . While this is the mathematical convolution, neural network programming libraries usually implement a related function called cross-correlation [6], which is similar to convolution but omits the kernel flipping:

$$S(i, j) = (K \star I)(i, j) = \sum_m \sum_n I(i + m, j + n)K(m, n). \quad (18)$$

The cross-correlation is still called convolution by convention, and the mathematical difference is mainly irrelevant for the intended purpose in CV.

Convolutional layers are followed by pooling layers. These layers reduce the spatial dimensions (width and height) of the feature maps, which helps to decrease computational complexity and control overfitting. By stacking multiple convolutional and pooling layers, CNNs can learn a hierarchy of features. Early layers learn simple features like edges and colors, while deeper layers combine these to learn more complex patterns like shapes, object parts, and eventually, entire objects. Architectures like ResNet (Residual Network) [265] later enabled the training of much deeper networks, further pushing the performance on various computer vision tasks.

The trajectory of computer vision was fundamentally altered in 2012. The introduction of a deep CNN named AlexNet [266] resulted in a dramatic reduction in error rates on the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) [267], a prominent benchmark for image classification. This event marked the beginning of the deep learning era in computer vision, rendering handcrafted feature engineering largely obsolete.

Region-based methods aim to identify regions of interest in an image for further classification task. A pioneer solution was Regions with CNN features (R-CNN) [268], which combines region proposals with CNNs. One of the most prominent region based CV methods is You Only Look Once (YOLO) [269], which frames the object detection as a regression problem and spatially separates bounding boxes, i.e., the areas where objects are detected, and provides class probabilities for each detected object, unifying these separate components into a single NN. The original model architecture comprises 24 convolutional layers and two fully connected layers, of which 20 convolutional layers, a pooling layer and a fully connected layer were used for pretraining before modifying into the final detection architecture. The design divides an input image into a $S \times S$ grid, in which each cell predicts bounding boxes as x, y, w and h , along with confidence defined as $Pr(\text{Object}) * IOU_{\text{pred}}^{\text{truth}}$. The confidence represents the Intersection over Union (IOU) between the predicted box and the ground truth. The x and y represent the center of the box in relation to the grid cell while w and h represent the relation to the whole image. During inference, the conditional class probabilities and individual box confidence predictions are multiplied as

$$Pr(\text{Class}_i | \text{Object}) * Pr(\text{Object}) * IOU_{\text{pred}}^{\text{truth}}, \quad (19)$$

which gives each box a class-specific score. Due to the combined structure, the unified YOLO architecture is fast enough to process video stream frame rate of images live, making it a go-to solution for a plethora of CV applications.

More recently, the Transformer architecture, which achieved state-of-the-art results in natural language processing, has been adapted for computer vision tasks. The transformer architecture and its attention mechanism is further studied in NLP and GenAI context in Section 4.4. The Vision Transformer (ViT) [270] model, introduced in 2020, demonstrated that a pure transformer architecture could perform on par with or better than CNNs on image classification.

Unlike CNNs, which have a strong inductive bias for locality, ViTs treat an image as a sequence of patches. To process 2D images, the ViT reshapes an input image $x \in \mathbb{R}^{H \times W \times C}$ into a sequence of flattened 2D patches $x_p \in \mathbb{R}^{N \times (P^2 \cdot C)}$, where (H, W) is the resolution of the original image, C is the number of channels, (P, P) is the resolution of each image patch, and $N = \frac{HW}{P^2}$ is the resulting number of patches. These patches are then mapped to a constant latent vector size D using a trainable linear projection, effectively treating the image patches as tokens in a sequence. Conceptually, the idea of splitting the image into patches is, in a way, similar to the YOLO functionality of splitting the image into the $S \times S$ grid, as both methods move away from having a computationally taxing sliding window solution or a pixel-by-pixel approach.

The network then uses a self-attention mechanism [9] to weigh the importance of different patches when creating a representation of the image. The core of this mechanism is the scaled dot-product attention, which computes a weighted sum of values (V), where the weight assigned to each value is determined by the dot-product of a query (Q) with all keys (K). This is formally defined as:

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V, \quad (20)$$

where Q , K , and V are matrices derived from the input token embeddings multiplied by learned weight matrices, and d_k is the dimension of the keys. The scaling factor $\frac{1}{\sqrt{d_k}}$ is applied to prevent the dot products from growing too large in magnitude, which would otherwise push the softmax function into regions with vanishing gradients [9].

This global attention mechanism allows ViTs to capture long-range dependencies within an image more effectively than standard CNNs, as every patch can theoretically attend to every other patch in the image sequence from the very first layer. Hybrids and variants like the Swin Transformer [271] have further improved efficiency and performance on tasks like object detection and segmentation.

While YOLO excels at real-time inference, the reliance on large annotated datasets is the primary bottleneck. To overcome this, there have been advances in self-supervised Vision Transformers, such as the DINO (Self-Distillation with No Labels) family of models [272; 273]. Unlike YOLO, DINO does not require bounding box annotations as it leverages self-supervised learning to extract rich, universal feature representations from unlabeled imagery. These self-supervised models can

subsequently be fine-tuned for specific operational tasks using only a fraction of the labeled data that traditional supervised networks require.

The progress driven by DL has enabled a vast array of real-world applications. Autonomous systems, such as self-driving cars and drones, rely heavily on CV to perceive their environment, detect obstacles, recognize traffic signs, and navigate safely [269]. In medical imaging CV algorithms assist radiologists in analyzing medical scans (X-rays, CTs, MRIs) for the early detection and diagnosis of diseases like cancer, segmenting tumors, and quantifying anatomical structures [274]. Facial recognition technology is used for identity verification and access control [275; 276]. Video surveillance systems employ computer vision to detect anomalous activities, count people, and monitor crowds. In agriculture, drones and cameras equipped with CV systems monitor crop health and diseases [277; 278]. Augmented Reality (AR) and Virtual Reality (VR) technologies use computer vision for spatial mapping, object tracking, and surface detection to seamlessly blend digital content with the real world [279; 280; 281]. In manufacturing, automated quality control on production lines uses CV to inspect products for defects at high speed, surpassing human capabilities [282; 283; 284; 285].

Despite significant progress, CV faces several ongoing challenges and is evolving in new directions. Training state-of-the-art DL models often requires massive, meticulously labeled datasets. A key area of research is developing methods that can learn from less labeled data, such as few-shot learning [286; 287; 288], self-supervised learning [289; 290; 291; 272], and the use of synthetic data [292; 293].

Models can be brittle and fail when presented with inputs that differ slightly from their training data due to, e.g., changes in lighting, viewpoint, or context, but can also be affected by intentional injection of suitable noise unnoticeable for human eye, known as adversarial examples [294]. Improving model robustness to these variations is a critical challenge. Likewise, moving beyond 2D images to understand the 3D world is a major frontier. This includes tasks like 3D object reconstruction from images, depth estimation, and scene understanding from point clouds. In addition, the integration of vision with other modalities, such as language and audio, is a growing trend. This leads to more comprehensive AI systems that can, for example, answer questions about an image or generate textual descriptions of a video.

From a practical standpoint, deploying complex computer vision models on resource-constrained devices like mobile phones and embedded systems requires model optimization, efficient hardware and preferably co-optimizing both in a hardware aware manner [295]. Edge AI is the research area that focuses on running these models locally for real-time processing and improved privacy. The rise of generative models, including Generative Adversarial Networks (GANs) [6] and diffusion models [173], has enabled the creation of highly realistic synthetic images and videos. This has applications in data augmentation, media creation, and also raises societal challenges related to deepfakes and misinformation.

In conclusion, computer vision has evolved from a nascent field of academic inquiry into a mature and impactful technology. While the DL paradigm currently dominates, the field continues to advance rapidly, driven by new architectures, larger datasets, and the pursuit of more robust, efficient, and comprehensive visual understanding.

4.1.1 CV in military domain

In the military context, CV is one of the most mature and widely applied areas of ML, used for tasks ranging from automated target recognition (ATR) in satellite and drone imagery to facial recognition and vehicle identification. The core challenge often lies in acquiring sufficient high-quality, labeled data for training and ensuring the model's robustness in diverse and adverse environmental conditions.

This thesis first explores the practical application of CV to the challenging naval domain of littoral warfare. In Publication I, "DeepMix: AI in Littoral Sonar Operations," a novel approach is presented for detecting objects from sonar images. The study addresses the significant challenges inherent in the underwater domain, such as high levels of noise, varying environmental conditions, and the scarcity of available, high-quality sonar data.

The research detailed in Publication I applies DL techniques to enhance object detection capabilities in shallow-water sonar operations. This serves as a foundational applicability assessment for the dissertation, grounding the research in a tangible and difficult military problem which not only highlights the results but the underlying issues in creating AI capabilities for military units. The research investigates how advanced AI techniques can be leveraged to process complex sensor data, thereby aiming to improve situational awareness and reduce the cognitive load on human operators in critical maritime environments. The findings from this study provide a low-level, practical perspective on the opportunities and limitations of applying CV in a data-scarce, high-stakes military setting.

From a CRISP-DM point of view, the purpose of the study was founded in the doctrinal need to speed up data processing to enhance Mine Countermeasure (MCM) operations. The current process of MCM was assessed and the relevant data identified, observing that the MCM units produce a lot of post-processed data that has been annotated by the subject matter experts, i.e., the operators. Therefore, this was a suitable research area to apply state-of-the-art ML methods such as ViT [270] and Mixture of Experts (MoE) [296; 297] in modeling to create novel value from old data.

The data understanding was created in collaboration with military personnel, while data preparation was alike any CV problem: the data was cropped and harmonized for modeling. The modeling itself had several hypotheses, utilizing individual AI models from conventional ML models to novel DL models.

The modeling itself tested several hypotheses, utilizing individual AI models ranging from conventional ML models, specifically Random Forest (RF) [298] and SVM [182; 299], to novel DL architectures.

The conventional models were selected to serve as robust, non-deep-learning baselines, establishing a performance floor for standard classification before introducing more complex neural architectures. Formally, an SVM is a maximum-margin classifier that seeks the optimal hyperplane to separate classes by minimizing $\frac{1}{2}\|w\|^2$ subject to the constraint $y_i(w^T x_i + b) \geq 1$ for separable data. Meanwhile, an RF operates as an ensemble learning method that constructs a multitude of decision trees during training, utilizing bootstrap aggregating (bagging) and random feature selection to output the mode of the classes, thereby reducing variance and overfitting.

The selection of the individual DL models was methodologically deliberate to contrast distinct computer vision paradigms. Specifically, the comparison included Visual Geometry Group (VGG) [300] and ViT [270] to highlight the transition from strictly localized feature extraction to global contextual representation.

VGG serves as a robust, well-understood and widely used baseline that represents deep CNN architectures which rely heavily on spatial inductive bias to capture local pixel hierarchies. In contrast, the ViT architecture acts as the state-of-the-art alternative that lacks this inherent spatial bias as explained above, instead using self-attention to capture long-range, global dependencies across the entire sonar image. Comparing these two fundamentally different approaches provides critical insight into whether the heavily obscured and noisy nature of underwater sonar imagery benefits more from strict local feature extraction or global contextual awareness.

The most promising approach combined the performance of several models using a MoE architecture [297]. Formally, an MoE system consists of N expert networks, E_1, \dots, E_N , and a gating network, G . For a given input x , each expert produces an output $E_i(x)$, and the gating network outputs an N -dimensional vector $G(x)$ representing the probability distribution over the experts. The final aggregated prediction y is computed as the linearly weighted sum of the experts' outputs:

$$y = \sum_{i=1}^N G(x)_i E_i(x) \quad (21)$$

In Publication I, this approach was used to weight the classification predictions of the best-performing individual models. By allowing the gating network G to learn which underlying models (experts) were most reliable for specific types of features in the sonar data, the MoE ensemble effectively mitigated the difficulty of the problem, surpassing individual model performances on all metrics.

The evaluation was done with k -fold metrics, using a tenth of the data as a test set for each round and taking the average performance over ten rounds of modeling for each solution. Therefore, the evaluation was operationally relevant and used the

small data set to maximum extent.

The key finding was that a mixture of several capable DL models was able to mitigate the difficulty of the problem to an extent, but the resulting models did not exhibit adequate reliability to be deployed for operational use. This is due to two main reasons. First of all, the data set was very small, only hundreds of images. Secondly, the data was of poor quality, as it was extracted from a post-processing system. This decreased the resolution of the data, potentially resulting in losing usable features. In addition, the data did not include the accurate location of the objects but a noisy location, so even evaluation for accurate object detection was difficult.

As such, the research invokes a hypothesis that military units create vast amounts of data without a data governance architecture that would allow and mandate using it for the development and fine-tuning of AI models afterwards or even hypothetically in a continuous setting. In order to increase the organizational maturity to leverage AI solutions, the data pipelines need to be considered holistically, so that the work of the operators who initially process the data is not lost in the process. Data pipeline, in this instance, is not just the software-based data integration and processing but rather covers the operational process from the user or operator to the AI engineers and developers. This poses requirements for the operated systems, which must allow saving the operator insight and data annotations as meta data to enable effective post-processing and model development with AI methods.

On a conceptual level, this hypothesis calls for the development of a *data capability*, where the systemic exploitation of data is treated as an active military function rather than a temporary or ephemeral input-output feed.

4.2 Reinforcement Learning

RL is a paradigm of machine learning where an agent or several agents learn to make sequential decisions by interacting with an environment.[27] Unlike supervised learning, the agent is not told which actions to take but instead discovers which actions yield the most reward or least penalty through a process of trial and error. The core components include an agent, the environment, states, actions, and rewards, with the agent's goal being to learn an optimal policy that maximizes cumulative rewards over time. This framework is exceptionally well-suited for modeling the dynamic and uncertain nature of military operations, which are fundamentally sequential decision-making problems in a dynamic environment.

RL origins are closely connected to the history of optimization. A unified framework that encompasses both stochastic optimization and RL has been advocated by Powell [301]. Fundamentally, despite differences, stochastic optimization and RL share a lot in common, although stochastic optimization emphasizes model-based approaches and RL is largely focused on model-free learning and value-based meth-

ods.

The foundation of modern RL is the framework developed by Sutton and Barto [27]. In their formulation, RL is presented as a problem of learning to map situations to actions to maximize a numerical reward signal. The agent is not provided with explicit instructions on which actions to take. Key aspects of this problem are that actions may affect not only the immediate reward but also subsequent situations and, consequently, all future rewards. These two characteristics, trial-and-error search and delayed reward, are the distinguishing features of reinforcement learning.

The core components of an RL problem are

- The policy (π), which defines the learning agent's behavior at a given time. It maps perceived states of the environment to the actions to be taken
- The reward (R_t), which signals immediate desirability, whereas a value function estimates long-term desirability. The value of a particular state represents the total amount of reward an agent can expect to accumulate over the future of sequential actions beyond that current state
- The model of the environment (optional)¹, which captures the behavior of the environment, allowing inferences to be made about how it will respond to actions.

The RL problem is formalized through a Markov Decision Process (MDP) [27]. An MDP is defined as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$, where \mathcal{S} is the set of states, \mathcal{A} is the set of actions, \mathcal{P} is the transition probability function for each state and \mathcal{R} is the reward function while γ is a discount factor used to estimate accumulating rewards with certain weight beyond the immediate return, i.e., $\gamma = 0$ would neglect future rewards past the next state, which is denoted as s' . An MDP possesses the Markov property, meaning that future is independent of the past given the present.

To contextualize this formalization within the military complexity conceptually described in Subsection 3.2.6, the true scale of the state space \mathcal{S} and action space \mathcal{A} in military operations is subject to the curse of dimensionality. When quantifying the operational environment, the state space expands exponentially based on the number of units, capabilities, and environmental variables. The complexity can be formally defined as:

$$|S_{\text{unit}}| = |V_{\text{sensors}}| \times |V_{\text{vars}}|, \quad (22)$$

where $|V_{\text{sensors}}|$ and $|V_{\text{vars}}|$ represent the number of possible combined states for the unit's sensors and its internal and environmental variables, respectively. The internal variables include human aspects, technical requirements such as power and maintenance and so on, while the environmental variables include weather, impacts of opposing force etc.

¹There is always an environment, but a model of the environment is optional.

Similarly, the action space \mathcal{A} available to a commander scales combinatorially with the forces deployed. For a headquarters managing n similar units, the total state space, S_{HQ} , grows exponentially, as the headquarter has to theoretically deal with all possible combinations of sensors and variables for each unit that are not dealt with at the lower level. In reality, the unit has been delegated the autonomy to act on certain inputs on its own, but even these autonomous actions create inputs to higher echelons. It can be deemed that the set of a_i independent actions in the set of all actions \mathcal{A} does not alter the rate of changes in the $|S_{\text{unit}}|$, all of which may or may not concern the headquarters. As the headquarter decides for itself whether or not all or a subset of changes in state are a concern, this relationship is the primary driver of complexity.

Furthermore, the transition probability function \mathcal{P} in a military MDP is highly volatile due to the friction of war and unexpected events. Rather than treating these as static probabilities, the introduction of unexpected events over time can be modeled as a Poisson process. Let λ be the average rate of "shock" events (such as equipment failure or enemy contact) per unit of time for a single unit. The probability of observing exactly $K = k$ shocks over a time horizon of length T is:

$$P(K = k, T) = \frac{(\lambda T)^k e^{-\lambda T}}{k!} \quad (23)$$

Hence, for example, if $T = 1$ and $\lambda = 0.5$, $P(K = 1, T = 1) \approx 0.3$ for one event ($k = 1$). For a system of n units, the aggregate event rate is $\lambda_{\text{HQ}} = n\lambda$. The expected number of shocks, $E[K]$, i.e., the mean of Poisson distribution $P(K = k, T)$ in Equation 23, increases linearly with both the number of units and the time horizon:

$$E[K] = \lambda_{\text{HQ}} T = n\lambda T \quad (24)$$

The total complexity of the planning problem, \mathcal{C} , is therefore a function of the exponential growth in the state-action space and the temporal expansion of the decision tree due to uncertainty. If each unit has a actions available, the total number of joint actions is a^n . The complexity can thus be conceptualized as:

$$\mathcal{C} \propto f(|S_{\text{unit}}|^n, a^n, n\lambda T). \quad (25)$$

As formulated in Equations 22 through 25, the sheer mathematical scale of military operations necessitates abstraction and sophisticated approximation methods, as exact tabular solutions are infeasible to compute. Instead of calculating every possible future permutation within this exponentially growing decision tree, an autonomous agent must learn to estimate the expected long-term utility of its current state and chosen actions under uncertainty. The mathematical foundation for breaking down these infinitely complex sequential decisions into recursive, evaluable steps

is the Bellman equation for the optimal state-value function, v_* [302], of which the original notation is shown in Equation 26:

$$f(p) = \underset{q}{Sup}[g(p, q) + h(p, q)f(T(p, q))] \quad (26)$$

This early version provided by Bellman gives a mathematical formulation for his principle of optimality which states that "An optimal policy has the property that whatever the initial state and initial decisions are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision" [302]. The $f(p)$ denotes value function at state p , q is an action variable, $g(p, q)$ is the immediate reward function for state p and action q , $T(p, q)$ is the transition function for the new state given respective state p and action q for which $f(T(p, q))$ is the value of next state, $h(p, q)$ is the discount function and Sup denotes supremum, i.e., maximum over actions. The modern, widely recognized versions of the same equation are

$$v_*(s) = \mathbb{E}[R_{t+1} + \gamma v_*(S_{t+1}) | S_t = s, A_t = a] \quad (27)$$

$$= \max_a \sum_{s', r'} p(s', r' | s, a) [r + \gamma v_*(s')], \quad (28)$$

where in the latter, original $f(p)$ is denoted as $v_*(s)$, $h(p, q)$ as γ , $g(p, q)$ ultimately as $r + \gamma v_*(s')$ and $f(T(p, q))$ as $p(s', r' | s, a)$.

Sutton and Barto [27] have expanded the original Bellman equation to incorporate full stochastic MDP model as well as policy and expectations over actions. An optimal action-value function is defined as

$$\begin{aligned} q_*(s, a) &= \mathbb{E} \left[R_{t+1} + \gamma \max_{a'} q_*(S_{t+1}, a') \mid S_t = s, A_t = a \right] \\ &= \sum_{s', r} p(s', r | s, a) \left[r + \gamma \max_{a'} q_*(s', a') \right] \end{aligned} \quad (29)$$

where $q_*(s, a)$ is action-value function, $p(s', r | s, a)$ is the probability of ending in state s' and receiving reward r given current state s and action a , with the discounted optimal future value $\gamma \max_{a'} q_*(s', a')$. The introduction of probability distribution p induces stochasticity. For a state-value function $v_\pi(s)$ under policy π , the formulation is

$$v_\pi(s) = \sum_a \pi(a|s) \sum_{s', r} p(s', r | s, a) [r + \gamma v_\pi(s')], \quad (30)$$

which states that the value of a state, $v_\pi(s)$, under policy π is the expected discounted return from that starting state onward following the policy.

Combining Bellman equation, and its variations, to (stochastic) MDP is the foundation for RL, which aims to optimize a target value in a sequential problem. RL introduces another fundamental concept, which is the exploratory, trial-and-error learning. The learning aims to an optimal policy, in other ML denoted as a model, and the learning of the policy is guided or updated usually according to a value function.

In Sutton and Barto [27] definitions, RL algorithms can be classified into tabular and approximation methods. Tabular methods relate to problems that are discrete in nature or can be discretized to form a table that encompasses the values related to each state or state-action pair. If the state or action spaces are continuous, an approximation method such as an NN is needed. Additionally, algorithms can be classified into direct and indirect methods, where direct refers to learning a policy from direct interactions with the environment, and indirect approach involves learning a model of the environment that is used to learn the policy from. The common understanding, which deviates slightly from this interpretation, is that model-based learning is regarded as indirect and model-free learning is direct. The third distinction is between on-policy and off-policy learning. On-policy learning means that the algorithm addresses updates directly to the policy that is used to explore the state-action space, while off-policy learning utilizes a separate policy for exploration and the learned parameters are iteratively transferred to the policy that is the end result of the training.

With its lengthy history, rooted in optimization and dynamic programming, RL has made a tremendous impact in recent years. Sophisticated RL solutions have shown performance beating human capabilities in games like Go [303] and StarCraft II [304]. These advances have been advocated by adopting sophisticated ML approaches to RL, such as Monte Carlo Tree Search (MCTS) [305; 306; 307] and neural networks from DL. Most recently, RL has been bridged to Reinforcement Learning from Human Feedback (RLHF) [308]. In essence, human comparisons between trajectory segments are used to train a reward model, which is used as a proxy reward function in RL. This has been extended to GenAI, where pretrained language models are fine-tuned on a smaller datasets of human-generated response. Another approach is Inverse Reinforcement Learning (IRL) [309], which aims to infer the reward function of an agent given its policy or observed behavior. In the military setting, the approach could be used to analyze the goals and commander's intent of an adversary from observed behavior of troops, given a suitable environment to enable such an analysis.

4.2.1 RL in military domain

This research focus is on the adaptability and adaptation of RL into military domain and problems. The Publication II explores the integration of Multi-Agent Reinforcement Learning (MARL) into the domain of littoral naval warfare, focusing on its

capacity to generate tactical Course of Actions (COAs) under littoral, adversarial maritime conditions. The study utilizes a simplistic, bespoke simulator to model naval engagements as a Partially Observable Stochastic Game (POSG) and tests two MARL algorithms, Multi-Agent Double Deep Q Network (MADDQN) and Multi-Agent Proximal Policy Optimization (MAPPO).

The selection of these two specific algorithms is methodologically deliberate to contrast two different RL paradigms. MADDQN serves as a well-established, value-based, off-policy baseline that operates over a discrete action space, representing a standard Deep Q-learning solution in multi-agent form, while MAPPO represents a state-of-the-art, on-policy actor-critic approach that natively utilizes a continuous action space. By comparing a value-based off-policy model against an advanced policy-gradient method, the research evaluates how different underlying learning mechanics adapt to the uncertainties of the naval POSG environment.

The POSG formulation expands the stochastic MDP to include partial observability characteristic to real-life military scenarios, where the transparency of the environment is clouded by at least three features: the lack of insight to adversarial planning, force composition, and capabilities, the alternating surrounding environment which encompasses for example the weather, non-combatants and third parties, as well as uncertain outcomes of different interactions, from warfighting functions to equipment malfunctions.

With MADDQN, the policy update is based on the Bellman equation shown in Equation 26, where the update is selected between the current policy and target network as

$$Y_t = R_{t+1} + \gamma Q(S_{t+1}, \operatorname{argmax} Q(S_{t+1}, a; \theta_t), \theta_{t-1}), \quad (31)$$

where R_{t+1} is the reward, γ is the discount factor, S is the set of states, and θ_t is the policy network and θ_{t-1} is the target network which is used to stabilize the original DQN training. Double Deep Q-Networks (DDQN) [310] separates the maximum operation in the target network into action selection and action evaluation, as depicted in Equation 31.

MAPPO, the multi-agent version of a PPO algorithm [311], is an actor-critic algorithm that has two separate NNs, the actor which executes the policy and the critic that serves as a value function. In this case, all the agents on the same side shared their actor and critic networks, and the multi-agent setting comes from having two opposing sides with their respective settings. As in research by Yu et al. [312], the actor network is trained to maximize the objective function

$$L(\theta) = \left[\frac{1}{Bn} \sum_{i=1}^B \sum_{k=1}^n \min(r_{\theta,i}^{(k)} A_i^{(k)}, \operatorname{clip}(r_{\theta,i}^{(k)}, 1-\epsilon, 1+\epsilon) A_i^{(k)}) \right] + \sigma \frac{1}{Bn} \sum_{i=1}^B \sum_{k=1}^n \mathcal{H}[\pi_{\theta}(o_i^{(k)})], \quad (32)$$

where π_θ is the actor network, o denotes observations and a actions, $r_{\theta,i}^{(k)}$ is the ratio $\frac{\pi_\theta(a_i^{(k)}|o_i^{(k)})}{\pi_{\text{old}}(a_i^{(k)}|o_i^{(k)})}$ between old and current network output values, θ denotes actor network parameters, B is the batch size, n is the number of agents, advantage $A_i^{(k)}$ is computed using a Generalized Advantage Estimation common for PPO algorithms [313], \mathcal{H} is the policy entropy and $\sigma \in [0, 1]$ is the entropy coefficient parameter, $\epsilon \in [0, 1]$ is the clipping ratio to stabilize the training and prevent excessive updates. The critic network, V_ϕ , that maps states to rewards, is trained to minimize the loss function

$$L(\phi) = \frac{1}{Bn} \sum_{i=1}^B \sum_{k=1}^n (\max [(V_\phi(s_i^{(k)}) - \hat{R}_i)^2, \text{clip}(V_\phi(s_i^{(k)}) - \epsilon, V_{\phi_{\text{old}}}(s_i^{(k)}) + \epsilon) - \hat{R}_i)^2]), \quad (33)$$

where \hat{R}_i is the discounted reward-to-go, meaning the reward that is accumulated from the current timestep onward.

The simulation environment is a 100x100 grid based on the northern Baltic Sea, with terrain and navigability derived from an aerial image. The simulation accounts for observation sharing, anti-surface warfare in the form of missile salvos and naval artillery, electronic warfare and littoral tactics like utilizing the cover of archipelago. The agents share data between the units of the particular side to create observations of the environment, using a CNN to interpret terrain and environment data, followed by an aggregated input of additional data before fully connected MLP that is used for action prediction or probability distribution approximation, depending on the chosen algorithm.

The study demonstrates that transitioning from standard RL environments to military POSGs requires careful algorithmic selection, as different paradigms handle partial observability and sparse military rewards differently. The comparative evaluation of how these models perform in generating naval COAs is detailed in chapter 5.

Consequently, the research shows that with proper reward engineering RL methods can discover valid tactics and strategies without doctrinal encoding, and may enable tactical decision-making edge if novel solutions emerge in the process. Otherwise, the research highlights and validates the upcoming findings of the Publication III: lack of suitable simulators and operational data inhibits advancing research and advocates proofs of concept instead of creating an operationally relevant path from hypothetical to applications in the real-world systems and proper evaluation.

To address the technology assessment objective of this thesis, a systematic review of RL for decision support in defence and security was conducted. This systematic literature review, presented in Publication III *“Reinforcement Learning for decision support in defence and security”*, was authored by a multinational research task group SAS-181, convened under NATO’s Science and Technology Organiza-

tion (STO). The aim is to examine how RL has been applied to support decision-making within defence and security contexts across NATO member states. The methodology of the survey involved the search strategy, where articles were selected with structured queries focused on three topics: decision support, RL, and defence and security, filtered to include only relevant peer-reviewed works with clear descriptions of RL methodologies.

The chosen 20 articles were classified by evaluating across 19 dimensions including time horizon, decision type and scale, uncertainty modeling, use of simulators, RL algorithm, type and presence of explainability features. An analysis using UMAP [314] was used for dimensionality reduction to visualize the research landscape and identify clusters of trends.

The survey highlights the applicability of RL to military domain where decision-making is characteristically partially observable, uncertain and time-sensitive, but applicable to be formulated as sequential decision-making problems. As such, RL is viewed as a tool that can support the decision-making by learning policies from simulations or operational data.

As with this dissertation, the Publication III framework is strongly rooted in MDMP and OODA-loop, which align RL cycles with real-world command decision-cycles. Another theoretical framework used is Powell's universal modeling framework for sequential decision [301; 315]. Against these frameworks, the survey covers 20 articles from application domains including defence planning, force generation, force projection and force employment, over multiple warfare domains albeit lacking in Multi Domain Operations.

As RL usually translates, in Powell's framework, to value function approximation (VFA) algorithms, VFA algorithms were the dominant policy type. Most frequent algorithms were Q-learning and Proximal Policy Optimization (PPO), mirroring the solutions in Publication II, which notably is also one of the 20 articles that were analyzed. Nearly all implementations relied on bespoke simulators, with few interfaces to standard RL environments. This underlines the lack of common simulation architecture for the defence use cases.

According to the findings, most reviewed systems had direct policy usage, meaning that the RL model output directly supports a decision or suggests it. Indirect use, where the RL model informs a broader planning process, were in minority. Likewise, single agent setups were in majority but several adversarial and cooperative multi-agent approaches were also explored.

From evaluation perspective, all articles used simulation for policy training but real-world validations were non-existent and results were evaluated either hypothetically or qualitatively instead of empirical evaluations in a relevant environment.

The research contributed to ML and military AI by mapping the state of RL research in military decision-support systems and identifying priority areas for research and technical innovation. The critical needs exist in explainability, simulator

integration, continual learning and multi-agent coordination. Regarding simulators, standardization is called for to enable interoperable simulators and shared data formats to enable multinational collaborations.

While the survey confirms RL potential to improve military decision-making, despite the challenges of uncertainty, partial observability and adversarial dynamics there are several persisting challenges. These include operational integration, simulation-to-real world transfer and ethical and legal considerations. Future of RL will likely be characterized by continued experimentation through simulations and military exercises and development of modular, interoperable and explainable RL systems designed in a bottom-up manner. The operational integration calls for trusted, transparent and explainable systems and closer collaboration between researchers and stakeholders to ensure these aspects.

In summary, the Publication III provides the big picture of RL in military decision-making in a cross-sectional manner, providing a technology assessment. The Publication II is included in the reviewed articles of Publication III, where it is highlighted as including most uncertainty factors from all the reviewed 20 articles. It brings the conceptual insights into a single decision-support task within one specified application and warfare domain from a technical and practical perspectives. Consequently, these publications are supplementary researches into the RL paradigm in the military context, encompassing the wide spectrum of problem areas throughout different levels of granularity.

From the CRISP-DM perspective, the major findings between these two publications relate to the business understanding and data. In both cases, it is highlighted that the data is scarce, in silos, and often classified, which hampers the research efforts. From the business understanding perspective, the heterogeneity of research due to the limitations does not promote an impact that would otherwise be achievable with more focused efforts and better infrastructure to execute meaningful, operationally relevant research in close collaboration with military stakeholders. Hence, while RL exhibits promise to solve complex decision-making problems, advancing from theory to impact requires further effort, funding and a common framework to enable research collaboration between all stakeholders and requirements of end-users to function towards advancing the state of research and application beyond current state.

As pointed out in the previous Section 4.1, it also applies to RL that data should be viewed as a capability. Collecting data from exercises and enabling RLHF [308] in military systems would be crucial for major advancements. Therefore, the collection, processing and storing of data should be mandated in the procurement and upgrade processes of military systems, as well as the interface requirements to enable learning from human feedback. In this sense it is important to understand that data is not just the data gathered from sensors, as shown in Figure 8, but also the human insight should be gathered.

4.3 Federated Learning

FL, established by McMahan et al. [316], is a machine learning paradigm that enables collaborative model training across multiple decentralized devices or servers holding local data samples, without exchanging the data. This approach is particularly critical in contexts where data privacy, security, and governance are paramount. The FL framework addresses three core challenges: it reduces the required communication overhead and cost as only model parameters are exchanged instead of entire datasets, it enables mitigating issues with heterogeneity and non-Independent and Identically Distributed (IID) data, and promotes privacy protection.

As a technology assessment, a comprehensive survey of this rapidly evolving learning paradigm was conducted. Publication IV, "Emerging trends in federated learning: from model fusion to federated X learning," provides a comprehensive review of the field. While not military specific, the study investigates the progression of FL from Federated Averaging (FedAVG) [316] to more advanced concepts, establishing the state-of-the-art and future trends. In its basic form, FedAVG objective is formulated as

$$f(w) = \sum_{k=1}^K \frac{n_k}{n} F_k(w), \text{ where } F_k(w) = \frac{1}{n_k} \sum_{i \in \mathcal{P}_k} f_i(w). \quad (34)$$

In Equation 34, the global model is updated with the average parameter weights over K clients and a data partitioning \mathcal{P}_k on client $k \in K$, with $n_k = |\mathcal{P}_k|$. The data partitioning is done in accordance with the amount of data per client, instead of uniform random sampling, to mitigate the non-IID setting the FedAVG is aimed to tackle.

Succeeding model fusion techniques, that are elaborated in Publication IV, include advanced methods for aggregating local client models to create an improved global model and vice versa, moving beyond the standard FedAVG algorithm. The goal is to create a more robust and optimal combined model despite the statistical heterogeneity of client data. Key approaches in model fusion include:

- **Adaptive and Attentive Aggregation:** These methods assign client contribution weights based on performance metrics, such as model parameter distance or accuracy, rather than simply the quantity of client data.
- **Regularization Methods:** To mitigate the "client drift" caused by non-IID data, regularization terms are added to the local or global objectives to constrain the local training process and enhance convergence.
- **Clustered Methods:** This approach groups clients with similar data distributions into clusters and performs model aggregation within each cluster, better capturing the heterogeneity across the entire network.

- **Bayesian Methods:** Probabilistic approaches are used for model fusion, such as aggregating neurons based on, e.g., maximum a posteriori estimation of global neurons and minimization of Kullback-Leibler (KL) divergence between global and local distributions to handle the architectural diversity of neural networks.
- **Fairness:** To prevent the global model from being skewed towards over-represented clients, fairness-aware algorithms adjust the optimization objective to ensure a more uniform distribution of performance gains across all participants.

Publication IV explores the integration of FL with other machine learning paradigms to create more flexible and powerful systems. This combination allows FL to be adapted to a wider range of real-world challenges. Federated Transfer Learning (FTL) and Knowledge Distillation (KD) are used to handle statistical and model heterogeneity. FTL transfers knowledge between clients with different datasets and feature spaces, while KD allows smaller client "student" models to learn from a larger, more powerful server model, accommodating diverse hardware capabilities. Federated Multi-Task and Meta learning approaches treat each client as a distinct but related task, which enables greater personalization. In multi-task learning, Separate models can be trained for each client with some shared structure between models exploiting related tasks, while meta learning aims to adapt a model to a new task by, for example, learning an initial shared model and a meta updating scheme. This allows for the development of models that can be rapidly adapted to new tasks with minimal data. Federated unsupervised, semi-supervised learning and RL integrations adapt the FL framework to scenarios where data labels are scarce or non-existent, as well as for distributed agent-based learning where clients learn from rewards of their actions.

While the baseline FL architecture inherently improves privacy by keeping raw data localized, the exchange of model parameters, i.e., gradients or weights, is not immune to privacy leakage. Adversaries can reverse-engineer sensitive information from these updates through model inversion [317] or membership inference attacks [318]. To robustly counter this, FL is frequently augmented with advanced privacy-enhancing technologies. Differential Privacy (DP) [319] is a prominent mathematical framework utilized in FL to provide formal privacy guarantees. By systematically injecting calibrated statistical noise into the local model updates before transmission, or into the global aggregation process, Differential Privacy (DP) obscures the contribution of any single data point. This ensures that the aggregated model does not memorize or leak individual records, albeit introducing a trade-off between strict privacy and model accuracy. Furthermore, cryptographic techniques such as secure aggregation [320] or homomorphic encryption [321] are commonly employed. These methods allow the central server to mathematically compute the aggregated global model without ever decrypting or exposing the individual, client-specific updates.

Future directions include on-device personalization, unified benchmarks for evaluation, and focus towards unsupervised methods, as these methods require less manual data processing, i.e., labeling, to be effective, as label scarcity is a fundamental challenge. Other challenges include collaboration of various techniques within the FL framework, as well as the mitigated but persistent issues with heterogeneity in both the data, the models and the devices, security and privacy issues, and communication efficiency. The research points out recommendation systems, healthcare, and, e.g., open banking as beneficiary real-world applications, but sharing similarities to military domain due to the impact on stated challenges, FL positions itself as a good fit for military purposes as well.

4.3.1 FL in military domain

The key characteristics of FL align well with many fields, such as healthcare and personal devices, but one can argue that the emphasis is possibly the greatest for military use cases. In a military setting, FL paradigm is highly relevant for training models on data that is distributed across different units, platforms, or even allied nations, and where security constraints and policies prevent data centralization and sharing. Likewise, reduced communication overhead is critical in contested environments, as military forces face far worse conditions with regard to connectivity than civilian or enterprise use cases. The military units are usually spread out over a large theater of operations, where the communications utilize a cascade of methods but are constantly challenged by the conditions and the very likely by the adversary.

In addition, the military data is at least as heterogeneous as any other real-world domain, but so are the use cases between different warfare branches, operational theaters and singular users. To elaborate on previously addressed ML paradigms, FL can be used, for example, for

- training shared CV target recognition models across allied nations, warfare branches or units, without sharing classified sensor data
- training RL decision support models in collaboration between several headquarters, possibly having also broader supply for RLHF that is explained in Section 4.4

Hence, the premise of FL in enabling better security and data privacy, handling of heterogeneous data combined with the possibility for personalized local models and their presumably elevated applicability to local user needs, as well as the reduction of communication overhead all serve well-recognized military requirements that promote secure, robust, and highly relevant ML adaptation through collaboration within military forces.

As Publication V will propose in the next section, FL is a powerful solution for developing a state-of-the-art military base model, i.e., a family of shared global

models. The analyzed solutions, such as FTL, KD and meta learning answer the individual, national or force-specific needs between allied countries while simultaneously preserving data privacy and reducing communication overhead. Therefore, FL offers a framework to provide forces with different features but a common objective to improve on their individual needs, and FL can be deemed as a potent research area to enhance the development and adaptation of military AI systems. Likewise, enabling distributed training can help solve the data scarcity issue to an extent. The distributed, heterogeneous computational resources and hardware requirements also support the claim, while simultaneously providing a more robust, distributed architecture to withstand a possible conflict scenario, in which the adversary is certainly likely to target concentrated AI capabilities such as large data centers.

From the CRISP-DM perspective, FL can be seen as the collaboration and co-operation of multiple parallel processes, which all handle their respective data in their respective systems, but share at least some business understanding or doctrinal task that can be enhanced by distributing the modeling efforts to all participants. FL counters several issues from the data understanding to modeling, as it aims to handle heterogeneous, non-IID data, in a distributed manner, but requires governance, standardization, and infrastructure to be realized effectively.

4.4 Generative Artificial Intelligence

Generative AI is a subfield of ML that aims to produce ML models which generate new, synthetic data based on their training data. Bishop and Bishop [173] do not treat GenAI as a single technique but as a family of modeling strategies whose common goal is to learn a data-generating distribution capable of synthesizing novel samples, optionally under user control through conditioning variables. This perspective makes “generation” a probabilistic task: the object of learning is $p(x)$ or conditionally $p(x|c)$, while the practical question is to define, train, and sample from models that approximate such distributions well enough to produce convincing data or to support downstream inference.

This section explores first NLP, as it has a strong link to generative text models and current trend of utilizing LLMs, after which the focus is directed to generative models.

4.4.1 Natural Language Processing

NLP refers to processing human language with a computer. Many NLP applications are based on probability distribution over sequences of words, characters, or bytes, in natural language. For example, fixed-length sequence based n -grams [6] are combinations of words, or tokens, with n tokens for each and they can be used in parallel, for example 2-grams and 3-grams. In essence, the n can be viewed as the context

window from which the context is derived from. These combinations of n length are used to define a conditional probability of the n -th token, a discrete entity representing, e.g., a word or a character, given the preceding $n - 1$ tokens as displayed in the training set. The fundamental limitation for this simple approach is the dimensionality, where even a large training set and modest n , most n -grams will not occur in the training set, resulting in zero probabilities and non-sensible outputs. This is due to the sparsity of natural language, where words can be combined in combinatorially vast and creative ways. If a combination is not included in the training data, it receives a zero probability, failing to generate a sensible continuation. Succeeding Neural Language Models (NLMs) were designed to overcome the curse of dimensionality problem with a distributed representation of words, called embeddings, not unlike hidden layers of CNN, representing words as dense vectors in a continuous, multi-dimensional space.

For neural machine translation, the early approach was to use an MLP for the common input-output training, so that the model would translate the received input into the target language. The problem is that for an MLP the sequences have to be of fixed length, which is gravely suboptimal for natural language. However, Recurrent Neural Network (RNN) have the ability to accommodate variable input and output lengths, as the recurrent network outputs the next token based on the certain sliding window that it uses to process the input [6]. For natural language, an RNN or CNN can be used to encode an input sequence to capture the context, which is then decoded into the target language with an RNN. This approach also requires a fixed-size representation and it is difficult to grasp all semantic details of a long input. The alternative is to produce the output in the sequential manner described, but focusing on different parts of the input for each output to consider the semantic details effectively. This mechanism, where different parts of the input affect the output, is called attention, essentially a weighted average over feature vectors and weights associated with each input position [6; 322; 9]. An early transformer-based solution that utilizes self-attention was Bidirectional Encoder Representations from Transformers (BERT) [323], which improved for example translation accuracy considerably, and acted as an early implementation of the larger generative language models that are addressed below.

4.4.2 Generative Models

The conceptual base for GenAI is latent-variable modeling. With discrete latent variables Bishop and Bishop [173] show how marginalizing over hidden assignments yields mixture distributions. The hidden assignment is the process of assigning each individual data point to one of the discrete categories it has learned. When using continuous latent variables, however, the conceptual framework shifts to a manifold viewpoint. This approach posits that complex data, such as realistic images, does

not fill its high-dimensional space randomly but is concentrated on a much simpler, smoother surface, i.e. a manifold, with the continuous latent variables serving as a coordinate system for this underlying structure.

GAN is a nonlinear latent-variable approach, where the generator transforms latent noise $z \sim p(z)$ into data space while a discriminator provides the training signal by learning to distinguish real from synthesized examples. The generator learns adversarially by aiming to minimize the difference detected by the discriminator [173]. While GANs do not typically offer a tractable likelihood, their implicit distribution and adversarial objective can yield high-fidelity samples.

The second approach, normalizing flows, constrains the generator to be invertible, allowing exact log-likelihoods via the change-of-variables formula $p_x(x|w) = p_x(g(x|w)) |\det J(x)|$, where $J(x)$ is the Jacobian matrix of partial derivatives $\frac{\partial g(x,w)}{\partial x}$ [173]. It enables straightforward maximum-likelihood training without likelihood approximation. The latent and data spaces must then have the same dimensionality, placing a structural cost on the approach. The third alternative approach includes autoencoders and Variational Autoencoder (VAE). Autoencoders are a useful, non-probabilistic precursor to VAE, which restore a generative interpretation by learning both an approximate posterior $p_\phi(z|x)$ and a generative model $p_\theta(x|z)$. The final approach, diffusion models, also known as denoising diffusion probabilistic models, inverts a multi-step corruption process. In the process, data is gradually perturbed towards a simple noise distribution and a NN is trained to reverse this process step by step. It can be viewed as a hierarchical VAE with a fixed encoder defined by the noise process and a learned decoder, and its training is stable and scales effectively on parallel hardware. Diffusion models are currently the state of the art in image domains.

While diffusion models are the current state of the art for images, transformer based LLMs are the text equivalent. Text models, specifically autoregressive transformers, are generative in the strict probabilistic sense. These models factorize a sequence distribution into token-level conditionals and sample one token at a time, where token is the mathematical representation of a word or a syllable, in this case [173]. LLMs are generative models over discrete token sequences. Given a prefix, denoted $x_{1:t}$, the LLM defines a conditional distribution $p(x_{t+1}|x_{1:t})$, which means that the probability for the next token x_{t+1} is calculated based on the preceding tokens from 1 to t . The model is trained to maximize the likelihood of the next token across the vast corpora it has been trained on. The sampling from the learned distribution outputs open-ended text, which can be conditioned on prompts as displayed in practice with the literature review analysis in chapter 2.

A decisive architectural shift was brought by the introduction of a self-attention mechanism [9], which replaces recurrence with parallelized attention over token positions and adds positional encodings to retain word order. Decoder-only transformers implement the autoregressive factorization used for text generation while

encoder-decoder variants handle conditional generation such as translation and summarization within the same attention framework. Still, the core learning signal is next-token prediction.

The tokenization bridges raw text and model inputs by providing said representations of words or other inputs. The design choice of tokenization, for example between byte-pair encoding [324; 325] or SentencePiece [326] balance open-vocabulary coverage with manageable vocabulary sizes, and shape the discrete space over which the distribution is learned.

Empirically, increasing model size, data, and compute improves LLM perplexity and downstream generalization, as explained by scaling laws [327]. Scaling, along with broad, self-supervised pretraining, has resulted in foundation models. These models are systems trained on broad data and at scale, which has enabled adaptation to many tasks through prompting or lightweight fine-tuning. Regarding foundation models, the transformers with attention mechanism are the currently dominant paradigm [173].

The trained, raw LLMs optimize the said likelihood of the next token, not usability. To enable a system that can, for example, follow guidelines and execute tasks, the pretrained model needs post-training stages such as instruction tuning or multitask prompted training on instruction-formatted datasets [173].

Likewise, RLHF [308] can be used to align the model with human preferences over outputs for which the base model is optimized. The core of RLHF relies on training a reward model $r_\phi(x, y)$ parameterized by ϕ , which takes a prompt x and a generated response y to output a scalar reward. In their seminal work, Christiano et al. [308] followed Bradley-Terry (BT) model, formulating the cross-entropy loss between the predictions and human labels as

$$\mathcal{L}_{BT}(\phi) = - \sum_{(\sigma^1, \sigma^2, \mu) \in \mathcal{D}} \mu(1) \log \hat{P}[\sigma^1 \succ \sigma^2] + \mu(2) \log \hat{P}[\sigma^2 \succ \sigma^1], \quad (35)$$

where

$$\hat{P}[\sigma^1 \succ \sigma^2] = \frac{\exp \sum \hat{r}(o_t^1, a_t^1)}{\exp \sum \hat{r}(o_t^1, a_t^1) + \exp \sum \hat{r}(o_t^2, a_t^2)},$$

and \hat{r} is a preference-predictor, \mathcal{D} is a database of triples $(\sigma^1, \sigma^2, \mu)$, where σ^1 and σ^2 are the two segments or sequences of observations o_t and actions a_t over a segment of time as Christiano et al. [308] applied this on a continuous RL environment, and μ is the distribution over $\{1, 2\}$ that indicates which one the user preferred.

To improve training efficiency and prevent overfitting in LLMs, modern approaches like InstructGPT [328] generate K different responses for a single prompt. Given a dataset where a human annotator ranks these K responses, the reward model is trained using a pairwise ranking loss, denoted here as \mathcal{L}_{rank} . The loss function is formulated to evaluate all $\binom{K}{2}$ pairs for a given prompt, maximizing the difference

in expected reward between the preferred response y_w and the rejected response y_l in each pair:

$$\mathcal{L}_{rank}(\phi) = -\frac{1}{\binom{K}{2}} \mathbb{E}_{(x, y_w, y_l) \sim D} [\log(\sigma(r_\phi(x, y_w) - r_\phi(x, y_l)))], \quad (36)$$

where D is similarly the dataset of human comparisons, $r_\theta(x, y)$ is the scalar output of the reward for prompt x and completion y under parameters θ , while y_w is the preferred completion between y_w and y_l . Once the reward model is trained, a reinforcement learning algorithm uses this reward signal to fine-tune the LLM policy to generate responses that maximize human preference.

To avoid repetitive and dull responses from the LLM, stochastic truncation methods such as top- k and top- p sampling are used to improve diversity by sampling among the probable region of the distribution. This can be controlled by alternating temperature parameter, which controls the randomness of the output, repetition penalties and other constraints.

Furthermore, the responses can be grounded using Retrieval Augmented Generation (RAG) [329], which is also explored in the context of military applications in Publication V. RAG is a paradigm that mitigates hallucinations [22], i.e., erroneous generation, and data staleness. Formally, RAG consists of two core components: a retriever $p_\eta(z|x)$ with parameters η , and a generator $p_\theta(y|z, x)$ with parameters θ . Given an input query sequence x , the retriever searches a large document index to return a set of top- K relevant text chunks, denoted as z . The generator then conditions its output sequence y on both the original query x and the retrieved context z . The probability of generating a target sequence y is approximated by marginalizing over the highest-scoring retrieved documents:

$$p(y|x) \approx \sum_{z \in \text{top-}K} p_\eta(z|x) \prod_{i=1}^N p_\theta(y_i|x, z, y_{1:i-1}). \quad (37)$$

By injecting external, verifiable knowledge z into the generative process, RAG allows the model to utilize tools such as external Application Programming Interfaces (APIs) for lookup, calculation, and code execution. This anchors the probabilistic text generation closer to the source data without requiring other measures such as retraining the parameter weights θ .

4.4.3 GenAI in military domain

GenAI represents a recent and significant development in the field of AI, characterized by models capable of creating novel content and following instructions rather than simply analyzing or classifying existing data. These models, most notably foundation LLMs built on the transformer architecture, can produce and process human-

like text, images, code, and other forms of data. This has led to a proliferation of applications that are rapidly altering both civilian and military domains. In the military context, the potential of GenAI extends from enhancing intelligence analysis and report generation to creating synthetic training data and augmenting command and control systems.

This dissertation explores this disruptive technology by identifying its current trajectory and potential impact on defence. Publication V, "*GenAI in Military: Trends and Opportunities*", provides a dedicated analysis of this topic. The structured review of recent literature (2022–2025) on GenAI in military applications reveals a significant disconnect between state-of-the-art academic and private innovations and practical defence adoption. While GenAI shows clear potential to enhance, for example, decision-making, simulation, and cyber security, current research is largely experimental and highlights two major obstacles: reliance on partially unsuitable proprietary models and lack of military-specific resources and infrastructure, such as datasets and military-secure computation capacity. To bridge this gap, a collaborative approach to developing a family of military-specific base models using federated learning is proposed as a key path forward.

Analysis of recent application studies shows a dual trend in military GenAI research. The first involves leveraging large, proprietary, off-the-shelf models (e.g., GPT-4) for high-level decision-support tasks. These models have been used to accelerate COA generation and to create multi-agent simulations for analyzing historical battles and strategic conflicts. The second trend involves fine-tuning smaller, open-weight models for narrow, domain-specific tasks, such as military equipment entity extraction and automated data tagging for cybersecurity frameworks. Conceptual research complements these applications by proposing strategic frameworks for GenAI's role in military competition, cross-domain ethical principles, and collaborative architectures.

A primary finding is the disconnect between the cutting edge of GenAI research that is driven by industry and academia breakthroughs like the transformer with attention mechanism itself, MoE architectures, unsupervised reinforcement learning, and models capable of test-time learning, and their limited adoption in military studies. This gap is largely due to significant operational, ethical, and security barriers.

The analysis identifies two major obstacles to successful GenAI adaptation in the military:

- **Reliance on proprietary models:** Large proprietary models, while powerful, are unsuitable for critical military use as-is. They lack explainability due to their vast size and undisclosed training data, and they exhibit undesirable behaviors, such as a tendency toward escalation in simulations, because they are not necessarily aligned with military doctrine.
- **Lack of resources and infrastructure:** The immense data and computation re-

quirements for training foundational models from scratch are significant bottlenecks. Even when fine-tuning smaller models, their performance falls short of larger systems, and their applicability is limited. This is compounded by the scarcity of military-specific datasets and benchmarks for proper evaluation.

These limitations do not concern large, technologically advanced nations in the same manner it concerns smaller ones or those lacking sovereign computation and data resources. For example, Anthropic has deployed Claude Gov models for national security customers in the United States [330], stating that these were built on direct feedback from government customers. Similar statements have been given by other large AI companies including OpenAI, Google and Meta [331; 332; 333]. While this highlights the inherent issues of proprietary, off-the-shelf models, the displayed service and supply of governmental models is only convenient for a large customer with national hyperscaler companies to provide the services and deployments without risking sovereignty. Both of the aforementioned obstacles hold for nations with more limited resources and technological sovereignty.

From the CRISP-DM perspective, the obstacles can be perceived from multiple points of view. It can be argued that the current process of creating capable GenAI models with readily available internet data inherently lacks the operational understanding, the data understanding, the modeling insight and perhaps foremost the evaluation. While these models have been proven useful in military tasks as well, the exhibited military understanding is shallow, starting from the corpus and tactical insight that exists in the available training data. To improve such models, military-specific data sets for base or foundation model training and instruction-tuning cannot be fully produced on behalf of military forces, as the required data has to serve the desired tasks, the proper corpus and the thought-through military intent.

To overcome these obstacles, the development of a family of military-specific base, or foundation, models, such as a "NATO base model", is identified as a top priority for allied nations. This would require collaboration with industry and academia, as well as investment in data, computation, and personnel. A key enabler for this is FL, which is proposed as a secure, system-level architecture for training LLMs collaboratively. Under an FL paradigm, allied nations could train a shared global model using their own local, sensitive military data without ever exposing the data itself. This approach would address critical issues of data privacy and ownership while reducing communication overhead and issues with heterogeneous data, enabling the creation of robust, client-specific military-grade AI capabilities. In the best case scenario, military-specific foundation model could be utilized to serve narrow and user-specific problems with greater efficiency than aiming to leverage larger, general models in a variety of tasks.

Additionally, the instruction-tuning data should be generated and stored *in operation*. This means that military forces would require information systems that store

the data that leads to a conclusion to create proper instruction-output datasets. For a simplistic example, within a particular MDMP, there are certain inputs that lead to certain types of COAs. In order to enhance the COA preparation, the data from inputs to suggested alternatives for actions should be stored. This, of course, neglects to an extent the chance to produce greater variance and makes the model converge towards a certain way of tactical or strategic "thinking". To enable further deviation, the systemic approach can be split into two parallel approaches: a legacy-based approach that aims for feasibility according to known doctrines, and exploratory approach that aims to leverage at least partially different knowledge-base to generate more ungrounded solutions.

4.5 Ethical considerations regarding AI systems

In the wake of wide-spread adaptation of AI-based or -enabled systems in the military the ethics of automation in warfare have been considered. As shown in chapter 2, the regulatory field is still in its infancy, and legal guidelines are scarce. However, Western nations widely share the idea presented in U.S. Department of Defense [124]: a human must be responsible for the use of lethal force. This has been echoed in other nations as well, for example, in the Finnish government defence report, which states that "While the decision to use lethal force must always be made by a human, the appropriate level of human involvement will not in all situations require on-line communications connection if responsible behaviour is ascertained in other ways." [334]. The very same report declares that ethical and juridical challenges of artificial intelligence must be resolved. Despite the consensus on the matter, the outspoken intents have not driven other ethical guidelines apart from the dictated human oversight.

The form of "meaningful human control" has been debated and researched [335], but propositions in general still share the idea that human control and oversight are fundamental concepts for the safe deployment of autonomous weapon systems. For example, the meaningful human control is perceived to include three components: that humans make informed, conscious decisions, that humans have sufficient information to ensure compliance with requirements of law, weapons, and context, that the weapon itself is designed and tested in realistic operational environment, and that humans are properly trained to ensure the judicious manner [336].

Publication VI claims that the use of AI in military systems creates a fundamental "Catch-22"-like paradox regarding reliability and human oversight. This dilemma places human operators in an untenable position, making them responsible for systems that are designed to outperform them. A proposed solution involves an ideological shift in the human-machine relationship, which reframes the human's role from one of scrutiny to one of support.

The central paradox emerges from two conflicting requirements:

- The need for perceived reliability: For an autonomous system to be deployed

and trusted, it must be perceived as exceptionally reliable, to the point that it surpasses human judgment.

- The requirement for human oversight: Legal and ethical mandates, such as the U.S. Department of Defense Directive 3000.09, require that a human operator exercise "appropriate levels of human judgment over the use of force".

This creates a paradox: the perception of high reliability needed for deployment undermines the rigorous questioning required for meaningful human oversight. An operator is unlikely to question a system that seems to operate flawlessly in most cases. This leaves the human operator in an impossible role: either they apply too much distrust and nullify the AI's performance advantage, or they apply too little and become a "surrogate scapegoat" for the system's errors. This results in a critical failure in the deployment and evaluation stages of CRISP-DM, but not from a technical perspective. The current model for deployment, requiring human oversight, is flawed as it misunderstands the operational goal of leveraging superhuman performance instead of replicating human performance. Likewise, the evaluation stage is flawed if it focuses on the AI system instead of the actual human-machine team that creates the operational capability in cohesion.

The root of this dilemma is identified as the "perceived sanctity of human intelligence". This analysis contrasts the ideal of "humane" action with the reality that human cognition is flawed by biases, emotion, and fatigue. A properly programmed AI, free from these distortions, could potentially adhere more objectively to ethical guidelines. Furthermore, machines are often held to a much higher standard of flawlessness than people, with every autonomous accident receiving harsh scrutiny while widespread human errors are more accepted.

To resolve this dilemma, an ideological shift in the human-AI relationship is proposed, outlined in a three-step model:

- Acknowledging the limitations of human cognition
- Aligning idealistic expectations with realistic AI capabilities, basing acceptance on measurable improvement over the human baseline
- Shifting focus from the "means", i.e., technology, to the "results", i.e., the improved outcomes

This shift reframes the interaction from the human oversight of a machine to human support for an intelligent system. In this model, the human's primary role is not to second-guess the AI's every decision but to cover its "blind spots" by providing additional, context-rich information that the system might lack. This process creates a feedback mechanism not very distinct from RLHF, but one performed during inference with a direct impact on the current output. The critical question for the human becomes: "Does he or she know something the system does not?", i.e., what are the limitations of the data pipeline and how can that impact the system. This approach

aims to leverage the distinct strengths of both human and machine intelligence for more effective and genuinely "humane" outcomes.

It can be argued that the very notion of putting humans to safeguard systems that are supposed, and expected, to outperform them in the same task is more idealistic than realistic. The approach, by design, hinders the ability to use the system to the capacity where it exceeds human performance: it can be said to strap the benefits of AI and ML to the maximum performance enabled by humans. The more beneficial, and thus arguably favorable from military perspective, approach is to recognize the limits of both, the human and machine, and enable both of these to contribute to the task in maximum respective capacity. As such, the suggested solution also hints guides to accept the imperfectness of all systems and models, which allows for more transparent and realistic conversation regarding AI systems such as autonomous drones and AI agents. This approach could also enable advancing the regulatory and juridical aspects without the risk of falling behind in the so called AI arms race.

4.6 Testing, evaluation, validation, and verification

Test, Evaluation, Validation and Verification (TEVV) practices have been a foundation for many fields, from robotics to software engineering, and it is just as crucial for ML and AI. Within the CRISP-DM framework, it is confined to the evaluation step, but the complete process includes a variety of angles in addition to mere evaluation and testing of a particular ML model [337].

The testing, as in ML training, is usually done by executing the model on a set of inputs to enable analyzing the resulting outputs. In evaluation, the empirical test results are reviewed against a testing criteria to assess the performance of the model from quantitative and qualitative perspectives. After evaluating metrics, verification aims to demonstrate that the system conforms to specifications, i.e, that is was "built right", and validation that it fits to the intended operational context. Overall, TEVV is not a single event but a continuous process that aligns with, e.g., CRISP-DM. As a simple example, an inference model can be used to classify a test set of images to produce empirical evidence, which is then analyzed with testing criteria, and then verified and validated for the intended context.

As an ML specific TEVV twist, deviating from traditional software, ML learns patterns from data, and as a result the data is an integral part of the specification. Quality, representativeness, provenance, and lineage must be tested and documented for proper TEVV. The ML performance is distribution dependent and systems employing ML models may fail under distribution shift, as with different sensors, weather, and deception. Likewise, the calibration of out-of-distribution detection for anomaly detection and mitigation is important.

AI and ML solutions also bring about new threat vectors to systems, for example adversarial examples [294], data poisoning [220; 338], and inverse attacks [317].

Likewise, cumulating new data and updating models invalidates prior assurance or, at least, provokes re-evaluation and testing.

Drawing from this widely applied TEVV practice with a long history in other technical fields there are currently multiple, complementary standards and policies that anchor TEVV for AI:

- NIST AI Risk Management Framework (AI RMF 1.0) [339] provides outcomes and practices to drive trustworthy AI, which ought to be valid, reliable, safe, secure, resilient, explainable, privacy-enhanced, fair and harmful bias managed, and explicitly recommends testing and monitoring across the lifecycle with a risk-based approach.
- ISO/IEC 23894:2023 [340] gives process guidance for AI-specific risk management (hazard identification, risk analysis/evaluation, treatment, and monitoring), integrating TEVV evidence into risk decisions.
- ISO/IEC 42001:2023 (AIMS) [341] establishes an AI management system standard that institutionalizes governance, role ownership, and continuous improvement, providing organizational scaffolding for repeatable TEVV.
- EU Artificial Intelligence Act (2024) [24] requires risk management, high-quality data governance, technical documentation, transparency, human oversight, and post-market monitoring for high-risk systems; conformity assessment hinges on TEVV evidence.
- Defence acquisition policy DoDI 5000.89 [342] defines Test and Evaluation (T&E) policy across acquisition pathways, while DoDI 5000.98 [343] governs Operational Test and Evaluation (OT&E) and live-fire T&E, emphasizing science-based Test and Evaluation Master Plan (TEMP) or T&E strategies, sequential testing with Bayesian or similar inference methods, and integration of developmental, cyber, and operational evidence.

These frameworks converge on a consistent view that TEVV must produce auditable claims, arguments, and evidence that a system is sufficiently safe, effective, fair, secure, and reliable for its intended use, but defence AI is mainly excluded. As pointed out in chapter 2, the laws and regulations for defence AI are still in infancy and close to non-existent, apart from certain guidelines for LAWS. For example, the EU AI Act [24] rules out military AI. TEVV for defence AI has to draw applicable procedures and solutions from the existing framework and adopt it to the military and defence context.

In order to execute TEVV for defence AI, the operational context has to be defined so that, e.g., the intended environment, platform, sensors, conditions, and human roles are considered, as well as the current doctrine, CONOPS, and possible SOPs. The test design and scenario coverage has to rely on this thorough context, simultaneously enabling all the TEVV steps.

4.7 Field observations from Ukraine

Author spent ten days in Ukraine, Kyiv region, between 15th and 25th of July 2025, meeting and collaborating with both local and international entrepreneurs, military personnel, academia and defence technology startups. The main purpose of the visit was a startup venue participation, which aimed to connect novel ideas to battlefield needs as well as suppliers to procurement managers.

While most of the details cannot be disclosed, there were several effecting observations regarding the use of AI and ML models.

First and foremost, the perceived use of GenAI capabilities was nonexistent. The main focus of effort was in CV, GNSS independent navigation, and tracking. As the current frontline combat is dominated by drones, electronic warfare capabilities have become extremely important. This is due to the fact that the drones are piloted by soldiers, and the most cost-effective countermeasure is to sever the connection with, for example, jamming. Same applies to GNSS based navigation for long range targets: jamming GPS signals is relatively easy even on large scale.

Due to these combat realities, the main focus for drone development was on GNSS independent navigation, which is usually based on CV methods of identifying the map location from the visual information available. This usually requires uploading satellite mapping or drone imagery from the area of operations before hand, so that the drone can perform cross-checking navigation on the edge, without connectivity back to a pilot. Essentially, the idea is to compare the drones camera view to the uploaded map, and infer the most probable location. The accuracy can be enhanced with sporadic GNSS signals (fixes) or ground beacons that enable triangulating the position.

For both situational picture compilation and engagement, the automatic identification of targets is an essential task. All the companies met that were developing automated tracking and engagement properties relied on YOLO models, or possible derivatives and alternatives, that were tuned for the particular purpose, such as engaging enemy drones. After classifying and identifying a target, the drone could approach it by aiming to keep it in the center of the field of view while proceeding directly towards the camera direction. In an interception mode, trying to catch a fast target mid-flight, a simple heuristic can be used. For example, at sea, a collision is deemed visually imminent if another ship is in the starboard or port quarter and the vector to that ship does not change over time. Similar heuristic can be applied to the drones, to keep the target in a fixed angle to ensure a simple optimization to collide with the target.

Only a few companies stated that they focus on collecting data from the field and utilizing that in their AI or ML efforts. This seemed like a highly sensitive topic, so the analysis is largely speculative. However, simultaneously, the cooperation between the industry and frontline units seemed seamless but heavily siloed.

The coordination of efforts, to avoid overlap and having a concentrated effect, was not imminent. Simultaneously, despite its tragic nature, the conflict itself is the most prominent field to gather operationally relevant data for AI development. The configurations of drone pilot stations did support collecting the video data, for example, but it was also indicated that the data is not exploited afterwards, at least in a coordinated manner.

While these observations are bound to cover only a fraction of the activities on-site, both spatially and temporally, they still confirm certain findings of this dissertation. First, ML development relies heavily on transfer learning by leveraging pretrained foundation models, especially in CV, and fine-tuning those to perform in a military-specific task such as navigation or engagement. Second, the integration from state-of-the-art breakthroughs to battlefield deployment is not a plug-and-play fit. Third, all of the observed development and innovation was focused on identified bottlenecks in current operations echoing a classic weapon-countermeasure development cycle. This is most likely due to the reality of the ongoing conflict, where innovations stem bottom-up from concrete problems on the battlefield and tactics. While understandable and effective on the battlefield for immediate survival, realizing a broader revolution in AI enhanced warfare requires systemic, far-reaching efforts that aim to integrate and transform the higher levels of the military information system of systems.

4.8 Summary of Findings

Across all the analyses, the decisive factor is not a specific algorithm but the system around the model: data itself, data pipelines (including feedback), interoperable simulators and tooling, privacy-preserving collaboration, non-IID data, heterogeneous environments and use cases as well as evaluation in relevant environments. When these are left to little consideration the results are meager and even strong AI models underperform. When these factors are considered and amplified, the researched paradigms (CV, RL, GenAI, FL) can be combined to deliver comprehensive military capability with a decisive impact. As the fundamental aspects and root causes are similar for each paradigm, this finding can be generalized to consider all existing and emerging AI paradigms and research areas. The unifying imperative is to treat data as a first-class capability and engineer the end-to-end learning loop displayed in Figure 9.

The dominant constraint through all research areas is data quantity, quality and availability. CV in sonar results, RL in littoral warfare as well as overall decision support systems, and GenAI for defence all converge on the same bottleneck: scarce, fragmented, or unsuitable data. Therefore, *data capability* is the central thesis. As noted in Subsection 3.2.2, a capability is defined not as a platform or system itself, but as the ability to execute a military task and generate fighting power. While data

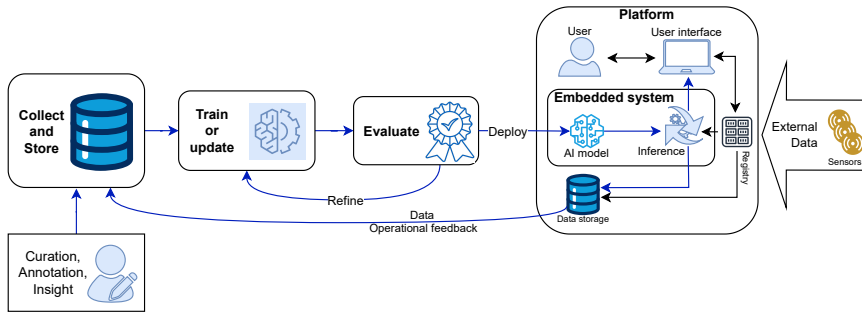


Figure 9. The Learning Loop

is simply an inert asset and cannot be defined as a capability in itself, the concept of a *data capability* transforms it into an active military function. A data capability denotes the systemic ability to collect, govern, and exploit data. Data capability should enable the execution of algorithmic tasks and the continuous improvement of systems, or systems of systems, with AI. In operational terms, this data-driven, continuous improvement translates into strategic, operational, and tactical superiority, a goal framed in contemporary discourse as *closing the kill chain* [163], compressing the OODA loop, and securing a decision advantage.

It has also been suggested that the data quality is more crucial than data quantity, as well curated, carefully crafted small data sets can yield impressive results [344; 345; 346] despite training with considerably fewer samples. As an adversarial result, threats such as data poisoning are also feasible with very small, targeted data sets, that at least for LLMs are able to bypass safety measures and enables the model to comply with harmful requests [338].

All paradigms and research areas require feedback signals, whether labels, rewards, preferences or federated gradients. This creates concrete requirements for systems to establish data capability, specifically the infrastructure to collect operational data and train AI models. Novel warfare and military information systems depicted in Subsection 3.2.5, from combat to managerial functions and from technical to strategic levels, should be designed to support the accumulation and exploitation of high quality data, preferably with annotations and other metadata, at scale. Expanding the observations regarding the data, evaluation requires benchmarks. For example, foundation models in GenAI are measured over various benchmarks, but military requires its own benchmarks to test both proprietary models as well as possible own developments in the field, in a conceptually similar manner to traditional military field tests, though the validation of intelligent, probabilistic systems is considerably more complex than testing conventional military systems such as kinetic assets.

On the other hand, while benchmarks or other test sets provide a shortcut to eval-

uate model performance, it may also create a spiral of convergence of solutions, as the benchmark results create requirements that need to be met and preferably surpassed, as is intended. Recently Gu et al. [347] pointed out that medical benchmarks seem to be measuring wrong things, and that LLM/LRM performance is very brittle. Their paper worryingly notes that LLMs provided correct answers even when removing critical parts of the input such as images. This indicates that the models have learned the answers "by heart", i.e., via rote memorization, instead of understanding, a test taking strategy that is shunned in human higher education.

The TEVV and benchmark approach needs to be carefully considered for each use case, especially with regard to high level decision support, as the other upside of applying AI consists of creating novel solutions. In the field of defence, novel solutions in, for example, decision making may create considerable tactical, operational or even strategic edge over the adversary. If the AI models are constraint to a certain, narrow distribution of known solutions, the evaluation ensures that novel solutions cannot be encouraged by the deployed AI systems. The tradeoff between measurable certainty of outcomes against the unforeseeable benefits of variance needs to be carefully considered, preferably implementing the solution on two parallel rails: the approved go-to solution and the possible, creative solution that stems from the margins of the probability distribution.

Regarding other weaknesses, distribution shift is a shared issue over all research areas. CV brittleness, RL generalization failures, GenAI erroneous generation, and FL client drift all stem from mismatch between training, deployment and use. Augmentation, self-supervision, and uncertainty estimation are crucial elements in bridging AI capabilities to operational reality. Additionally, the infrastructure is lacking. RL requires simulators, CV and GenAI need curated data in vast quantities or impeccable quality, as well as metadata schemes, while FL requires secure orchestration as well as network and device management. Lack of standards blocks multi-nation collaboration and repeatable evaluation, as seen with the absence of unified military benchmarks.

Rapid adaptation via transfer learning and fine-tuning generally outpaces the building of narrow base models in all fields, provided that the expertise, domain data and infrastructure exist. Likewise, explainability and trust are universal requirements for all research areas, as users require trust, confidence calibration and rationale. Otherwise, systems stall at proof-of-concept level. It is proposed that human-machine teaming should be designed for support rather than brittle "oversight", which necessitates trust calibration but also insight to the AI systems functioning. This necessitates technical understanding and expertise for the personnel, so that the advanced AI systems are not perceived as black box solution automatons but rather as inference machines that operate under certain limitations, requirements and constraints.

On a final note, the influential essay from Richard Sutton, "The Bitter Lesson" [348], posits that the two methods that "seem to scale arbitrarily" are *search* and

learning. He argues, with the academic weight of being one of the most celebrated pioneers of RL, that sophisticated solutions that mimic human knowledge or cognition usually fail, in time, against search and learning backed up by large enough computation capacity. The essay highlights "the great power of general purpose methods" as the key takeaway from the 70 years of history in AI and ML. An obvious counterargument would highlight the transformer architecture and its current triumphs, but it can also be described as a novel way to *search* vast knowledge bases after *learning* a comprehensive representation of them. Sutton's idea underpins the example laid out in Figure 9 and the conclusions of the following chapters: *there has to be data to learn from, and there has to be an environment to search within*. Whatever the algorithm or the model architecture, these principles hold, and according to Sutton, in due time given computation, robust data capability, and suitable simulation environments, it is inevitable that general methods to scale into the next generation of AI achievements.

5 Contribution of this thesis

The research presented in this dissertation addresses the challenge of developing and integrating machine learning capabilities within the military domain. The contributions are not confined to a single technical area but span conceptual, technical, applicatory, and ethical dimensions and considerations, backed up by field observations in Ukraine described in Section 4.7.

5.1 Publication I: *Deep Mix: AI in Littoral Sonar Operations*

5.1.1 Summary

Publication I investigates the application of deep learning techniques to enhance naval MCM operations in the challenging littoral environment of the Baltic Sea. The primary goal is to automate the detection and classification of underwater objects from Side Scan Sonar (SSS) imagery, aiming to reduce the workload on human operators and speeding up the overall mine hunting process. The authors explore the effectiveness of several AI models, including CNNs and ViT, for this task. A key focus of the research is the development of a MoE system, which combines the strengths of multiple models to improve classification accuracy. The paper also emphasizes the importance of explainable AI (XAI), using methods like Grad-CAM to provide visual feedback to operators, ensuring the system is transparent and operationally useful.

5.1.2 Methods and Data

The research utilized a unique, operationally relevant dataset provided by the Finnish Navy, consisting of 1,299 processed side-scan sonar images collected with a Klein 5500 system. The images are categorized into four classes:

Mines, Mine-like Contacts (MILCOs), Rocks, and Wrecks. Due to the limited size of the dataset, data augmentation techniques were employed to expand the training set.

The study's methodology involved several machine learning models:

- Baselines: A SVM and a RF were used as classical baseline models for per-

formance comparison.

- Deep Learning Models: The core of the research involved transfer learning with pre-trained models, specifically the VGG16 and VGG19 CNN architectures and a ViT (ViT-B/16).
- MoE: The primary contribution was testing an MoE framework that combines the outputs of the best-performing models (VGG16, VGG19, ViT, and SVM) using a gating network to weigh their predictions to improve overall accuracy.

The models were evaluated on two classification tasks: a four-class problem and a three-class problem where the 'Mine' and 'MILCO' classes were combined to reflect operational procedures.

5.1.3 Results and contribution

The experimental results demonstrated the difficulty of the classification task, with individual models achieving modest performance due to the challenging nature and limited size of the dataset. The Vision Transformer (ViT) generally outperformed the VGG models and the baseline SVM.

The main contribution of this work is the successful application of the MoE framework, which significantly improved classification performance over any single model. The best-performing configuration, a three-expert MoE (MoEv3) on the three-class task, achieved an overall accuracy of 73.29%. This result, while lower than accuracies reported in studies using cleaner or synthetic data, is significant given it was achieved using real-world, operational data from the cluttered Baltic Sea environment.

Furthermore, the paper successfully demonstrated the feasibility of adding an explainability tool for operators. By using Grad-CAM and attention heatmaps, the system can visually highlight the regions in a sonar image that led to a particular classification, providing valuable and transparent decision support for MCM operations. This research is reportedly the first to apply a transformer architecture and an MoE framework to mine detection using real-world sonar data from this specific region.

5.1.4 Author's contribution

Author was in charge of the research project and the methodology, contributing to model implementation and evaluation as well as writing, editing and submission of the final report. Author also applied for, in cooperation with professor Heikkonen, the research funding, which enabled the research in the first place.

5.2 Publication II: *Strategizing the Shallows: Leveraging Multi-Agent Reinforcement Learning for Enhanced Tactical Decision-Making in Littoral Naval Warfare*

5.2.1 Summary

Publication II develops and evaluates a MARL approach for generating COAs in a littoral warfare context representative of the Baltic Sea. The combat environment is formalized as a POSG with reactive agents that base decisions on local and joint observations under explicitly modeled uncertainties. Two families of RL methods are instantiated and compared in the same environment: (i) a discrete-action pipeline centered on DDQN, where the state-space consists of surrounding terrain and map objects as well as other values such as radar state, number of missiles, location of friendly units and enemy bearings, and is processed with a CNN and linear NN to produce actions, and (ii) a continuous-control or probability distribution pipeline using MAPPO with reward normalization and parameter-space noise to address sparse rewards. Experiments on a 100×100 grid with $\approx 2.7\text{nm}$ cells, 15-minute time steps) produce plausible, tactically interpretable COAs. While MADDQN tends to converge to simple policies under uncertainty, MAPPO yields more stable learning and qualitatively stronger COAs aligned with established littoral principles.

5.2.2 Methods and Data

The RL environment instantiates a POSG, an MDP extension, with multiple agents per side, partial observability, and stochastic effects reflecting naval operations in cluttered, shallow waters. The state is rendered as a 2D “game grid” over the Baltic Sea (100×100 grid, $\approx 2.7\text{nm}$ cells), advanced in 15-minute ticks. Obstacles and electronic-warfare considerations are encoded; inter-cell movement feasibility is checked with A* pathfinding. Exogenous uncertainties (e.g., sensing and engagement outcomes) are injected as probability distributions. Agents receive observations, communicate within side, and select actions that produce rewards tied to mission outcomes (including victory signals and penalties for losses).

Two algorithms from different RL methodologies are implemented and compared. For discrete action spaces, a DDQN is trained with a CNN that processes a slice of the observation state as a 2D image prior to convolution, from which the output is concatenated with other feature inputs feeding fully connected layers. For continuous control, MAPPO is employed in similar manner with only a deeper fully connected linear layer NN and with PopArt normalization of return (value) signals and parameter-space noise for exploration under sparse rewards; clipped, zero-mean Gaussian perturbations are adapted during training. Training tactics include

- Adjusting learning rate when repeated victories emerge to curb overfitting and

guide the exploration towards the found solutions.

- Increasing per-rollout epochs when success is detected to exploit promising policies.

For scenarios, core experiments use 3-vs-3 Blue/Red surface combatants alternating the turns of trained agents, first with predetermined Red policies for initial training of Blue agents, then optimizing Red tactics without updating Blue agents and finally tuning Blue tactics against learned Red tactics. Multiple rollouts per seed are executed to cope with stochasticity. Resulting COAs are visualized as trajectories from start points to nearest engagement clusters over the Baltic grid.

5.2.3 Results and contribution

In the discrete-action setting, MADDQN converges fast but utilizes simplistic policies and struggles with generalization, exhibiting a tendency to fixate on overly simple policies in the face of non-stationary, uncertain dynamics. In the reported runs, Blue achieves victory in 33.9% of trials while Red wins 11.7%, the remainder being draws. The aggressive, straightforward maneuver that disregards complex environmental interaction in favor of immediate engagement conceptually mirrors Admiral Horatio Nelson's famous tactical doctrine at the Battle of Trafalgar, where he discarded traditional parallel formations to sail perpendicularly, straight at the adversary's line to force a decisive close quarter combat [349; 350]. In contrast, MAPPO demonstrates improved training stability and performance, adapting better to sparse rewards and environment complexity. The exploration was aided by parameter-space noise instead of action noise, which amplified the results. The learned COAs for Blue are tactically sensible, showcasing emergent behavior that correlates with established tactical naval doctrine e.g., holding units back to exploit the cover of archipelago while spreading some units towards the open sea to enable tracking. The agents were utilizing terrain, coordinating unit movements and distributing force posture to balance concealment and target acquisition. This resembles an established littoral doctrine and offered interpretable options for a commander. Direct quantitative comparison between the results of the algorithms is not suitable, as the measured metrics highlight engagements and victories, but the policies differ greatly in aggressiveness, rendering more cautious MAPPO policies less effective despite being considerably more robust in survival rates and qualitatively evaluated tactical coherence.

The publication provides an operationally grounded MARL testbed for littoral warfare with uncertainty modeling and POSG formalization. As a result, it produces a practical comparison between MADDQN and MAPPO in the same environment, including network architecture, reward handling, and exploration design choices that address sparse signals.

The research gives evidence that MARL can synthesize plausible, commander-

useful COAs rather than only maximizing abstract scores, enabling visual products that support human reasoning. Reproducibility details such as scenario, grid scale, timing, hardware, enable follow-on experimentation. Collectively, the work shows how MARL can augment COA development for higher-level decision-making in shallow-water naval contexts.

5.2.4 Author's contribution

Author created the research setting, developed the Python simulation environment as well as the RL algorithm implementation and model training and testing, and wrote the initial research report. PhD Saastamoinen summarized, edited, submitted and presented the research at 20th AIAI conference in June 2024.

5.3 Publication III: *Reinforcement Learning for decision support in defence and security: A systematic review*

5.3.1 Summary

This article is a systematic review of how RL is being used to provide decision support in defence and security. It surveys public, post-2000 work with authorship from NATO SAS-181 nations, and complements the review with primers on military decision making (MDMP/OODA), DSS, RL fundamentals, simulation, and explainability. The core contributions are:

- A curated set of 20 defence decision-support applications using RL.
- A 19-dimension classification mapped to Powell's unified framework for sequential decisions.
- A Uniform Manifold Approximation and Projection (UMAP) landscape of the field.
- A synthesis of gaps and recommendations.

The paper anchors readers with an OODA-to-RL mapping and an extended unified framework that adds TEVV and explainability to Powell's original pipeline. The literature landscape is then summarized via UMAP, with tables cataloging studies and cross-tabs.

5.3.2 Methods and Data

The review uses a multi-national search strategy over national databases, Web of Science and Scholar, combining three facets: decision support, RL/ADP, and defence/security. Inclusion criteria included public availability, year ≥ 2000 , defence

or security focus, explicit use of RL for support, and at least one SAS-181 author. Screening removes behavior-automation work (no human-in-the-loop support), non-RL optimizers, and entertainment-only game studies. Each retained article is labeled on 19 characteristics including application domain, policy class, uncertainty handled, simulator type, maturity, explanation capability), which the authors map to stages of the extended framework that includes modeling, uncertainty quantification, policy design, algorithm strategy, TEVV, and explainability. The field is then projected with UMAP: nominal features are one-hot encoded, ordinal ones are ordinal-encoded, and two projections are unioned to reveal clusters. Methodological scaffolding includes an OODA–RL process diagram, the extended framework, and a complete study matrix, with targeted cross-tabs for maturity by domain, uncertainty coverage, simulation horizon versus human time available, and direct and indirect policy usage versus MARL setting.

5.3.3 Results and contribution

Results show that most decision-support applications concentrate on force employment at the tactical level. Value-function policies dominate the solutions, and nearly all studies use bespoke simulators. Exogenous uncertainty is commonly modeled but other forms are underrepresented and algorithmic explainability is rare. Maturity skews towards theoretical and proofs of concept with no fielded systems. The UMAP view separates studies primarily by presence/absence of evaluation, with a secondary split by direct versus indirect policy usage. Additional relationships include a largely aligned simulation horizon and human decision window, and that indirect uses more often coincide with adversarial MARL settings.

The paper raises four practitioner challenges:

1. Complex, multi-actor, non-stationary scenarios.
2. Scarce, siloed, and sensitive data.
3. Weak RL–simulator interoperability.
4. Human trust and explainability.

These findings are translated into concrete recommendations, which include to model to the decision, which can be described as fidelity by purpose or Occam’s razor principle, to perform early, iterative TEVV, standardize data sharing, e.g., through NATO Alliance Data Sharing Ecosystem, pursue a wargaming cloud and RL-ready interfaces, and operationalize explainability. Collectively, the review provides a reproducible taxonomy, a field map, and a deployment-oriented agenda for RL-based decision support in defence and security.

5.3.4 Author's contribution

Author was in charge of sections 2.1 and 5.2 while contributing as part of the research group to the whole review article, it's writing, related meetings, working groups and submission. Weekly meetings revolved around reporting and discussing the progress while the writing progressed. The work also included week long workshops that focuses solely on reviewing and progressing the paper. All sections were reviewed and commented by all the participants.

5.4 Publication IV: *Emerging trends in federated learning: from model fusion to federated X learning*

5.4.1 Summary

Publication IV is a focused survey on FL viewed through the lens of model fusion and its intersections with other paradigms, hence labeled as “federated X learning.” It organizes advances beyond FedAvg into five fusion families that include adaptive and attentive aggregation, regularization, clustered, and Bayesian methods, with the sixth cross-sectional focus on fairness. The research then maps how FL couples with transfer learning, knowledge distillation, multi-task and meta-learning, adversarial, semi- and unsupervised learning, and reinforcement learning. The survey contrasts this vantage point with general FL surveys, emphasizes statistical heterogeneity, communication and privacy as core drivers, and closes with challenges and future directions such as label scarcity, on-device personalization, unsupervised/self-supervised FL, combining paradigms, benchmarks, and production readiness.

5.4.2 Methods and Data

This is a narrative, scope-delimited literature review instead of a PRISMA-style systematic review. The authors formalize the standard FL objective and FedAvg workflow to anchor comparisons, propose a taxonomy of model-fusion strategies, and survey “federated X” couplings with concrete formulations, e.g., transfer objectives, KD losses, multi-task formulations over client-task matrices, meta-learning updates, adversarial learning for bias mitigation, semi- and unsupervised losses, and FedRL coordination. Evidence is drawn from peer-reviewed venues and well-cited preprints up to early 2024, summarized in tables and topical sections, with brief mathematical formulations to clarify algorithmic families. Application highlights cover recommendation systems, healthcare, IoT, and edge scenarios. As a note, the author of this thesis promoted military as a prime application field to be mentioned, but the focus was kept on aforementioned subjects due to their universally accepted and understood nature.

5.4.3 Results and contribution

Across fusion methods, the survey finds that

- Adaptive and attentive aggregation can temper non-IID drift by learning client weights from parameter distance, recency, accuracy, or attention.
- Regularization (proximal, momentum/mime-style, dynamic penalties, contrastive, prototype) aligns local and global objectives and mitigates client drift and privacy noise.
- Clustered FL (two-stage, multi-center, IFCA-style, ensembles) yields multiple globals to capture client subpopulations.
- Bayesian approaches (neural matching, variational personalization, ensembles) address permutation invariance and uncertainty.
- Fairness objectives (MiniMax/Q-fairness, collaborative/group fairness) rebalance gains for under-represented clients.

For “federated X” and learning paradigms, the review catalogs workable patterns and proposed solutions for FTL, KD, Federated Multi-Task Learning (FMTL), meta-learning, adversarial learning, semi- and unsupervised FL as well as Federated Reinforcement Learning (FedRL).

The paper contributes a taxonomy centered on model fusion and FL’s couplings that complements broad FL surveys, a curated map of algorithm research with concise objective forms that clarify where resilience to heterogeneity, security and privacy, and communication efficiency is introduced, as well as future directions and challenges regarding label scarcity, on-device personalization, proliferating unsupervised learning, and collaboration of multiple federated paradigms. The final conclusions call for a unified benchmark to better enable research and improvement, as well as an agenda for production FL to showcase practical applications of label-efficient training and on-device personalization that utilize the aforementioned unsupervised or self-supervised pretraining, combined paradigms, unified benchmarks and tooling, and deployment patterns robust against real-world problems such as drift, diurnal effects, and cold-starts. These elements provide a deployment-oriented framework that can be used to position other contributions within model-fusion choices and “federated X” integrations.

5.4.4 Author’s contribution

Author organized the framework and architecture of the paper as well as collected, categorized and analyzed the background data from esteemed conference papers, as well as reviewed the manuscript. Originally, a hundred high-level conference papers, with methodological significance, were retrieved from, for example, AAAI and NeurIPS proceedings, which were then analyzed by the author and clustered into

different groups depending on the methodology, proposed solutions, performance metrics and applicability.

5.5 Publication V: *GenAI in Military: Trends and Opportunities*

5.5.1 Summary

Publication V surveys how Generative AI (GenAI) is entering military use and where it realistically adds value now. It opens with a concise state-of-the-art introduction that covers Transformers with self-attention mechanism, MoE, RAG, Chain-of-Thought (CoT), knowledge distillation, unsupervised RL, Titans architectures, and agentic AI. The research then reviews 29 military relevant publications from 2022 to early 2025, spanning decision-support and COA generation, wargaming and simulation, information extraction and fine-tuning, as well as cybersecurity. The review highlights the application field and utilized methodology, assessing opportunities and promises such as faster planning, better intel synthesis, training, cyber defence, as well as risks including hallucinations, escalation tendencies, doctrine and corpus misalignment, security and ethics gaps. The paper argues that development and deployment will hinge on domain-specific data and interoperable, secure architectures, and makes a case for allied, federated approaches to build military-grade base models to reduce reliance on proprietary systems.

5.5.2 Methods and Data

The study conducts a structured literature scan across RAND and RUSI outputs and academic sources from Scholar, with the query (*Military OR Defense*) AND (*GenAI OR LLM OR Generative Artificial Intelligence*) filtered to 2022–2025. From an initial pool of top 200 results inspected according to relevance 48 were preliminarily retained, and 29 publications met inclusion criteria of GenAI focus and military relevance. Each paper was double-coded into seven categories: Survey, Review, Policy Analysis, Application, Proposition, Overview, Other, where application and proposition papers received deeper assessment for readiness, feasibility, and implementation issues. As stated above, the article supplements the review with a state-of-the-art primer to contextualize trends and a cross-cutting analysis of gaps and recommendations.

5.5.3 Results and contribution

As a result, the research shows that current applicatory research focuses on decision-support, domain information extraction and small-model fine-tuning, as well as se-

lected cyber tasks. Decision making and decision support includes COA generation as well as agentic settings simulating diplomatic discourses, strategic, and tactical scenarios. Smaller-scale approaches focus on operator level systems for unmanned assets and, e.g., intent detection.

Risks and limits denote that LLM agents can over-escalate in simulated geopolitics. COA generators accelerate planning but can raise friendly casualties. Studies rely on unclassified data, are simulator-bound, and not integrated end-to-end with actual C2 and ISR workflows. Likewise, security, ethics, and explainability remain underrepresented. The trend suggests that practice gravitates to either proprietary, general-purpose LLMs used via prompting for exploratory surrogates, and in a parallel lane towards smaller open-weight models fine-tuned for missions— or task-specific purposes, while cutting-edge advances such as unsupervised RL, Titans, and long-context, are mostly commercial and not yet militarized at scale.

The paper curates and classifies recent GenAI-for-military literature into an actionable map, integrates a state-of-the-art primer to connect frontier techniques to military constraints, and translates gaps into concrete guidance: build secure, interoperable pipelines, operationalize governance and human-AI teaming, develop mission-aligned benchmarks, and pursue coalition training of military base models via federated learning to overcome data-sharing limits. These elements form a deployment-oriented agenda for making GenAI a trustworthy decision-support capability rather than an off-the-shelf curiosity.

5.5.4 Author's contribution

Author collected and analyzed the data to gather the used knowledge base of the research field, wrote the manuscript, edited and submitted the research and applied the reviewer comments from revision. PhD Koski supervised the article and contributed in the methodology, co-analysis of the data and the synthesis of findings.

5.6 Publication VI: *The dilemma of AI reliability*

5.6.1 Summary

Publication VI analyzes a core paradox in deploying AI for military use: to be fielded, an AI system must be perceived as highly reliable, yet the very perception of reliability erodes the ethical-legal demand for “meaningful human oversight.” The author frames this as a Catch-22: either the human slows the system enough to negate its benefits, or the human becomes a nominal gatekeeper who defers to the seemingly reliable machine and only absorbs responsibility when things go wrong. The argument is developed across autonomy configurations generally stated as in/on/off-the-loop, decision-support as well as weapon systems, the limits of explainability for

modern ML, and the practical constraints of testing and generalization under scarce, heterogeneous military data. Historical episodes such as Petrov’s false-alarm intervention and Patriot fratricides illustrate asymmetries in risk tolerance for humans versus machines. The paper proposes an ideological shift to treat the human as a support to an intelligent system, supplying context and missing information, rather than as an all-knowing overseer, and adopt a three-step model to accept human cognitive limits, align ideals with realistic improvements, and focus on outcomes over means.

5.6.2 Methods and Data

This is a concept-driven, argumentative essay grounded in policy, e.g., DoD 3000.09, ethics literature, and illustrative cases rather than empirical experiments. The method is analytic and synthetic. It defines the deployment context, including autonomy levels, DSS and weapons, assesses explainability and validation limits for complex ML models such as Grad-CAM and CoT aids, over- and underfitting issues, out-of-distribution risk under non-IID data, and reasons from historical incidents and comparative risk tolerance in civilian autonomy versus human drivers, missile warning and air-defence cases to articulate the oversight paradox. The proposal is then structured into a normative, three-step integration model that reassigns the human role from scrutinizer to context amplifier and reorients acceptance criteria toward measurable improvements in performance and error-to-risk rates relative to the human baseline.

5.6.3 Results and contribution

Perceived reliability is both necessary and corrosive to oversight. The more “reliable” a system appears, the more humans rationally defer, especially under time pressure, creating a de facto responsibility gap. Explainability-first remedies scale poorly for deep models and can undercut the very scalability advantages that motivate autonomy. TEVV limits in military contexts that include data scarcity, heterogeneity, and non-IID drift make flawless assurance infeasible. Acceptance must hinge on comparative outcomes, not impossibly absolute guarantees. Risk tolerance asymmetry means machines are judged against near-zero-error expectations, while humans routinely err without equivalent scrutiny.

The publication makes an ethical and philosophical contribution by identifying and proposing a solution to the “reliability-oversight paradox” in autonomous military systems. The results claim that a clarified *Catch-22* of AI reliability for military AI that unifies oversight, trust, and responsibility into a single deployment dilemma requires a role reallocation, framing human as information augments for the AI, resolving the oversight-to-performance tension without abandoning ethical intent.

Three-step adoption model acknowledges human cognitive limits, calls to project ideals into realistic, measurable baselines, and aims to focus evaluation on outcomes. This is claimed to yield practicable acceptance criteria and policy-relevant guidance that shifts governance from scrutinizing every decision toward ensuring interfaces and processes that surface blind spots and feed the system timely, context-rich data or showcase the absence of it. Together, these elements offer a path to integrate AI that is more humane in result, even when less “human” in mechanism. In this new model, the human operator is tasked not with second-guessing the AI system, but with covering its inherent blind spots by providing and assessing wider context and information the system lacks.

5.7 Conceptual framework

Derived from the original publications, the primary conceptual contribution of this dissertation is the formulation and validation of the thesis that for military AI, data capability is the critical enabler. This has been the case from the information security point of view, as the classification and security of information have been a key military aspect for, arguably, as long as military forces have existed. However, the point of view has been limited to the particular information in a particular sample, and its value has been assessed in isolation. For example, information regarding technical capability of a weapon system is strictly classified, as it gives away the performance metrics of the system. The wider context of the information, as machine-readable data, has not been a priority nor a capability concern. At most, data has been viewed as an information flow inwards, and AI as a technology has been seen as a way to handle the increasing amounts of data. Simultaneously, operational data has been ephemeral and its value has been measured with regard to the specific use case for which it was obtained.

This work argues for a paradigm shift away from a narrow, algorithm-centric view of AI development toward a holistic, system-centric perspective focused on the entire learning loop that transforms the ephemeral use of data into a continuous cycle of data acquisition, preparation, continuous improvement, and deployment of operational AI models.

This framework, displayed in Figure 10 and substantiated across the analyses of CV, RL, FL, and GenAI, indicates that the dominant constraint and simultaneously the greatest opportunity in military AI is the data ecosystem. The contribution lies in defining this ecosystem not merely as a repository of information, but as an active, operational capability that encompasses data governance and pipelines, feedback mechanisms, and interoperable infrastructure.

The data governance and pipelines are the processes and infrastructure required to collect, process, and exploit high-quality data at scale. Data pipelines require integrations to connect the collection of data to its preprocessing, preparation and

utilization. Data governance is also linked to accessibility, as a significant portion of the military data is classified and accessible on a need-to-know basis. Hence, a rigorous framework is required to enable the best tradeoff between data availability, its use, and data security.

System design must incorporate built-in feedback mechanisms to capture human operator inputs and feedback, such as labels, rewards, and preferences, to accumulate operationally relevant, enriched data that enables continual model improvement without placing an unnecessary strain on already limited resources such as personnel.

Interoperable infrastructure refers to both digital and physical infrastructures. The digital infrastructure includes common simulators, benchmarks, and tooling to support training, testing, evaluation, and collaboration. The physical infrastructure encompasses data centers, computation capacity, and secure operational networks that enable proper connectivity. Crucially, the physical infrastructure also requires robustness through distribution and federation to ensure operational capabilities during conflict.

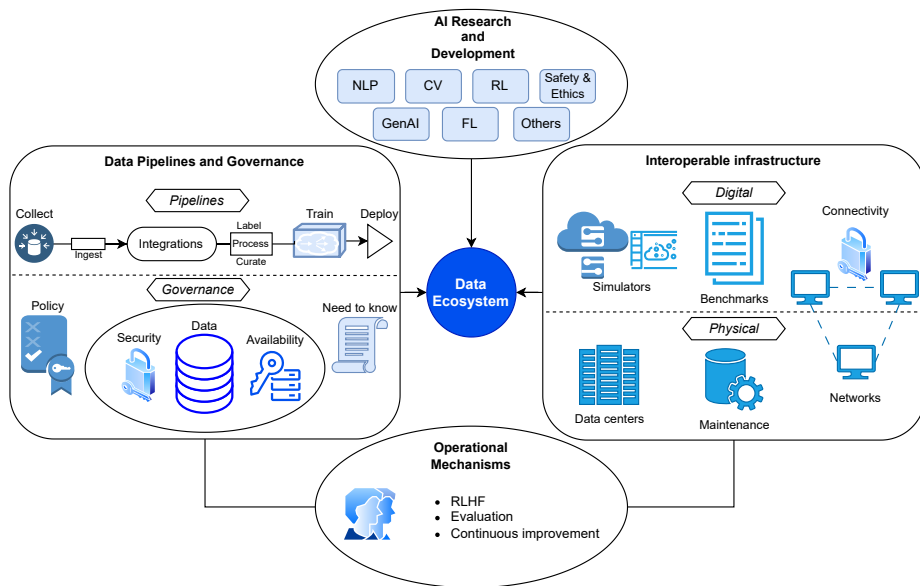


Figure 10. The Data Ecosystem framework

The complete conceptual framework is, non-exhaustively, visualized in Figure 10, aiming to capture the key points. The development, deployment and improvement of military AI capabilities through research and development relies on the ecosystem, which is structurally governed by data pipelines and policy frameworks. These include the non-technical processes and policies as well as technical procedures and solutions that first dictate what each process aims to achieve as well as the constraints

and requirements that are placed on it. These include, for example, the accessibility and security policies. The interoperable infrastructure comprises digital and physical entities, where physical structure serves as the capacity for digital systems. The digital capacities include the stated simulators, benchmarks, and connectivity protocols, but also their own policies and guidelines on usage and development. Finally, the operational mechanisms include the stated system requirements that aim to enhance data collection, preprocessing, and storing to create valuable data capabilities in order to advance the AI research and development efforts.

Despite being conceptual, the physical and digital infrastructure of this ecosystem must correlate with the military information system described in Subsection 3.2.5. To avoid treating AI as an external add-on, the model deployment, data pipelines and feedback mechanisms must be embedded directly into the DSS, C2 and ISR systems that constitute the military's existing information architecture.

Finally, while the emergence of AGI capabilities in AI is speculative at best, not being supported by current broader understanding across the scientific field, a great shift could occur if an RL based paradigm could be instantiated to continuously learn from (near) real-world like interactions. This was briefly iterated in Section 4.8, quoting Richard Sutton's claim that only search and learning, as general purpose methods, scale arbitrarily to arbitrarily complex problems. This possible future course does not render the Figure 10 obsolete, but merely changes the way it would be viewed from human perspective and operated by the ML algorithm or model that could evolve either *in silico* or even when operational.

5.8 Methodological Framework for Synthesis

This thesis presents a methodological contribution by adapting and applying the CRISP-DM as a tool for synthesizing disparate military AI research. While traditionally used for enterprise data mining projects, this work shows its utility as an academic framework for connecting low-level technical findings to high-level strategic and operational objectives.

By consistently applying the military adapted version of CRISP-DM stages across the different publications, this dissertation provides a structured and repeatable methodology for assessing the operational relevance of technical AI research and development, identifying systemic bottlenecks such as data availability and infrastructure gaps, and most importantly providing a framework to translate theoretical and technical algorithm or model performance into tangible, operational military capability.

6 Conclusion

This dissertation set out to examine how machine learning can be developed and integrated within the military domain and to determine what truly constrains and enables military AI capabilities. Across four research areas, namely computer vision, reinforcement learning, federated learning, and generative AI, the analyses converge on a single finding that the data capability, realized as the data ecosystem, is both the dominant constraint and the greatest opportunity. For advancing military AI, specific models and algorithms matter less than a robust data ecosystem where data itself is treated, not as a by-product or an ephemeral feedstock, but as a first-class operational resource.

Empirically and methodologically, the work contributes in three ways. First, it synthesizes evidence from applied studies and literature surveys to show that all practical ML paradigms require mechanisms to collect, govern, and exploit high-quality data at scale, under strict security and classification constraints. Second, the dissertation adapts CRISP-DM as a synthesis framework, connecting problem definition to data understanding, preparation, modeling, evaluation, and deployment in military settings. This common lens, adapted from the industry, makes visible where bottlenecks accumulate and what are the issues regarding successful deployment of ML in military domain. The highlighted bottlenecks exist especially at three points: the business understanding, which can be understood as doctrines and policies in the military sense, as well as the mere existence and availability of quality data, and finally the data understanding. The business understanding that guides the goal setting is consequently goal-focused: even ancient doctrines highlight speed and quality of data processing, although in different wordings. However, the understanding lacks the expertise to identify the systemic limitations that hinder the full process of advancing from data to deployment.

Likewise, the existence and availability of data as well as the related data understanding is limited, which means that data is seen as an ephemeral tool for decision-making unless it consists of detailed capability information, such as intelligence reports and technical details of weapon systems. This prevents storing high-quality data for preprocessing into valuable data products that could be used for modeling purposes. The data understanding appears two-fold: at first, if the value of certain data is not recognized, it will not be stored. If the data is not stored, the emergent understanding is never reached, as usually the data can be explored with, for example,

unsupervised methods to reveal patterns and value that is not trivially evident. For the user, the value of data is indeed ephemeral, as the task is not the data itself but what the data immediately enables, with regard to a certain desired end state.

These three are simultaneously the choke points where investments have the highest leverage. From a supplementary perspective, the model architectures and algorithms are not as much of hindrance compared to the mentioned three aspects.

As a third contribution and a solution for the aforementioned issues, this dissertation proposes and motivates an integrated Data Ecosystem framework, shown in Figure 10 that combines governance and pipelines, operator-in-the-loop feedback mechanisms, and interoperable infrastructure spanning digital (simulators, benchmarks, tooling) and physical (compute, networks, secure facilities) layers.

6.1 Summary of Key Findings

As the hypothesis is that AI is, or inevitably becomes, a major military capability that largely determines the future of warfare, the central counterargument of this thesis is that progress in military AI is fundamentally constrained by systemic issues related to data availability, quality, governance, and infrastructure. This was consistently shown across all research areas in the system-level investigation into the applicability, impact and challenges of current AI methods and research areas.

Publication I, a CV study on sonar imagery, highlighted how even state-of-the-art ML models fail or fall short when confronted with small, low-quality, and fragmented real-world sensor data that represents the operational reality, underscoring the critical need for integrated data pipelines and a curated data repository. The analyses of RL in Publication II and Publication III revealed that the field's potential is hampered by the scarcity of high-fidelity, interoperable simulators and the lack of operational data, which are essential for training, validation, and real-world transfer. The Publication I and Publication II provided empirical evidence of the challenges in solving even narrow problems with a limited resources regarding data and supporting systems, fulfilling the empirical research objective.

The investigation into GenAI in Publication V provided insight into the future direction and challenges of this novel ML research area, identifying a critical dependency on proprietary models that are misaligned with military requirements, arguing that progress hinges on developing military-specific foundation models, a goal that is currently hindered by the lack of curated data and a collaborative infrastructure. In part, this meets the future directions and challenges insight objective.

In response to these challenges, this dissertation puts forward two key solutions. First, FL that is introduced in Publication IV was identified as a critical enabling paradigm for secure, collaborative model development among allied nations, substantially mitigating the tension between data sharing and security as well as displaying a possible solution to non-IID data and user specific requirement chal-

lenges. Second, the ethical analysis in Publication VI proposed a resolution to the “reliability-oversight paradox” by reframing the human role from simple oversight to one of active support, thereby creating a more effective and realistic model for human-machine teaming within the operational learning loop and integrating ethical considerations as a system-level factor into military AI capability development.

From the system-level, top-down perspective, the adaptation of CRISP-DM methodology to bridge the AI research into tangible operational capability was, within the limitations of this thesis, shown feasible and effective. Other frameworks can be adapted in the similar manner, but CRISP-DM was selected primarily to bridge the theoretical research into practice in an industry equivalent manner.

Overall, the objective of a system-level investigation into the applicability, impact and challenges of developing military AI capabilities has been met alongside other objectives.

6.2 Implications

The primary, tangible contribution is the conceptual framework of the data ecosystem that synthesizes the findings into a holistic, actionable model for military AI capability development. It aims to shift the focus from isolated ML models and AI paradigms to the surrounding infrastructure, processes, and policies required to sustain a continuous learning loop. As pointed out in Section 4.8, historical perspective posits that in the long run, advancing computation capacity enables success with general methods, given that there is sufficient high-quality data and, in some cases like RL, an environment to search within. In this context, the robust data ecosystem that continuously improves data capability is the core component in advancing military AI. This does not render algorithmic innovations and novel model architectures obsolete in any way, but enables and strengthens the ability to adapt these in an efficient manner.

The practical implication is that military AI programs should prioritize building data capabilities, which include policies, processes, and infrastructure, so that continual data acquisition, curation, secure sharing, and reuse are enabled across missions and organizations. This includes systematic feedback capture including labels, rewards and preferences to support ongoing model improvement without placing unnecessary strain on already limited personnel. Interoperability and robustness must extend across both digital standards, APIs, and physical networks, with distribution for resilience under contested conditions.

Ultimately, these implications extend across the military domain on different tasks and functions shown in Table 2, and even surface a critical ethical and organizational point regarding the design of human-machine teams.

Table 2. Implications of a Data-Centric Approach for the Military

Domain	Implication
Doctrine and Policy	Military doctrine must recognize data as a strategic asset and a core capability component, on par with traditional platforms and weapon systems. Because commercial markets cannot supply the highly specific, classified data required for tactical operations, allied forces must organically generate this data capability through their own operations.
Acquisition and Procurement	New systems must be procured with "data readiness" as a core requirement to actively realize the proposed Data Ecosystem. Moving beyond passive data logging, procurement mandates must specify interoperable data pipelines, open architectures, and built-in interfaces that capture operator use and feedback (e.g., for RLHF). Enforcing these requirements and standards ensures that ephemeral operational information is systematically transformed into an active data capability for continuous AI model development and improvement.
Research and Development	R&D efforts should prioritize the creation of common, interoperable infrastructure, including shared benchmarks and simulation environments, to enable repeatable evaluation. Furthermore, multinational collaboration should be fostered through privacy-preserving paradigms like FL, advancing military-specific foundation models, complementing and reducing reliance on proprietary, commercial alternatives.
Human-Machine Teaming	The design of AI-enabled systems should move beyond the brittle concept of "human-on-the-loop" oversight and towards a model of human support, where operators are trained and equipped to augment and assess the AI system by covering its known limitations.

6.3 Limitations and Future Research

This dissertation, being a cumulative work, synthesizes findings from specific ML subfields. While this provides a broad overview, it does not encompass every potential military application of AI. The research was primarily focused on Western military contexts, inevitably reflecting an authorial bias rooted in European, and especially Nordic, strategic and operational perspectives and realities. Further work is needed to explore these concepts in other strategic and operational settings. It is also noteworthy that the analyses are necessarily bounded by the availability of public sources, small-scale experiments, and domain-specific constraints.

Building on the Data Ecosystem framework, future research should pursue several key directions. Firstly, the development of military-specific, standardized benchmarks for core military tasks, such as target recognition and COA generation, is critical for objectively measuring progress and validating ML models and AI systems. Secondly, future research should focus on architectures for federated operations. Research is needed to develop and test robust FL architectures in operationally relevant, multi-national exercise environments to address challenges of network latency, security, personalization, knowledge distillation, as well as model and hardware heterogeneity. Further research is required to validate the proposed "human-as-supporter" model, measuring the combined performance of human-AI teams against conventional human-in-the-loop and human-on-the-loop approaches.

Finally, bridging the simulation-to-reality gap is a considerable research area outside the military, but a concerted effort is required to create high-fidelity, validated simulation environments that can serve as reliable training and testing grounds. Not just for RL agents, but for TEVV purposes and, for example, to validate COA generation in a quantified manner as well as to research autonomous system deployment cost-effectively.

In conclusion, the successful integration of state-of-the-art AI into military forces will not be achieved through a technological silver bullet, especially one that is not built from within. It requires a deliberate, systemic, and sustained effort to build an ecosystem that can cultivate, process, and leverage data at the speed and scale of modern conflict. The path to effective military AI runs through data capability: governed, pipelined, fed by human feedback, and supported by interoperable infrastructure. Acquisition programs that build these foundations will be best positioned to field trustworthy, resilient, and continually improving AI systems, closing the gap between the promise of AI and its operational reality.

Declarations

Declaration of Funding and Conflicts of Interest

There are no conflicts of interests to be declared, and this research acquired no funding apart from Publication II, for which the statement is declared in the paper itself.

Research ethical statement

The summary of this dissertation was written solely by the author, adhering to common research ethics. AI was used in the making of this research in its obvious role as the primary research interest as displayed in original publications and as a tool as displayed in the Chapter 2. Large, proprietary LLMs, Gemini and OpenAI GPT, were used in formatting, spell-checking and cross-checking the main text. These models were also used to format and check bibliography and to search for research papers such as foundation sources for original ideas. Despite these tools, the research effort and the written text is solely work of the author.

Bibliography

- [1] Friedrich Naumann Foundation for Freedom. Lethal autonomous weapons systems: Challenges for regulation and the role of the european union. Technical report, Friedrich Naumann Foundation for Freedom, 2023. URL <https://shop.freiheit.org/download/P2@953/335163/Policy%20Paper%20LAWS-ENG-Final.pdf>.
- [2] Forrest E. Morgan, Benjamin Boudreaux, Andrew J. Lohn, Mark Ashby, Christian Curriden, Kelly Klima, and Derek Grossman. Military applications of artificial intelligence: Ethical concerns in an uncertain world. Technical Report RR-3139-1-AF, RAND Corporation, Santa Monica, CA, 2020. URL https://www.rand.org/pubs/research_reports/RR3139-1.html.
- [3] The Editors of Encyclopaedia Britannica Encyclopedia Britannica. Bombe. Encyclopedia Britannica, 2025. URL <https://www.britannica.com/topic/Bombe>. Accessed 24 August 2025.
- [4] Francis Harry Hinsley and Alan Stripp. *Codebreakers: the inside story of Bletchley Park*. Oxford University Press, 2001.
- [5] Alexander Jung. Machine learning: The basics, 2022. URL <https://arxiv.org/abs/1805.05052>.
- [6] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [7] Nestor Maslej, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Njenga Kariuki, Emily Capstick, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Juan Carlos Niebles, Yoav Shoham, Russell Wald, Tobi Walsh, Armin Hamrah, Lapo Santarlasci, Julia Betts Lotufo, Alexandra Rome, Andrew Shi, and Sukrut Oak. The AI Index 2025 Annual Report. Technical report, Institute for Human-Centered AI, Stanford University, Stanford, CA, 2025.
- [8] Bradley Martin, Danielle C. Tarraf, Thomas C. Whitmore, Jacob Deweese, Cedric Kenney, Jon Schmid, and Paul Deluca. Advancing Autonomous Systems: An Analysis of Current and Future Technology for Unmanned Maritime Vehicles. Technical Report RR-2751-NAVY, RAND Corporation, Santa Monica, CA, jan 2019. URL https://www.rand.org/pubs/research_reports/RR2751.html.

- [9] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, page 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964. URL https://papers.nips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.
- [10] OpenAI. Introducing chatgpt (research preview). OpenAI Blog, November 2022. URL <https://openai.com/fi-FI/index/chatgpt/>. Accessed last on 2026-03-07.
- [11] OpenAI. Chatgpt. <https://chat.openai.com>, 2023. Accessed: 2025-06-20.
- [12] NATO. Emerging and disruptive technologies, 2024. URL https://www.nato.int/cps/en/natohq/topics_184303.htm. Accessed 7th March 2026.
- [13] Francis Bacon. *Novum Organum Scientiarum*. John Bill, London, 1620.
- [14] René Descartes. *Discours de la Méthode pour bien counduire sa raison, et cherched la vérité dans les sciences*. Jan Maire, Leiden, 1637.
- [15] René Descartes. *Key Philosophical Writings*. Wordsworth Editions Limited, 1997. ISBN 9781853264702. Translated by Elizabeth S. Haldane and G.R.T. Ross.
- [16] Defence Acquisition University (DAU). AI Glossary for the DoD, 2025. URL <https://www.dau.edu/sites/default/files/Migrated/CopDocuments/DAU%20AI%20Glossary.pdf>. Cited 20.7.2025.
- [17] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Pearson, 3 edition, 2016. ISBN 9781292153964.
- [18] Elaine Rich. *Artificial Intelligence*. McGraw-Hill series in artificial intelligence, 1983. ISBN 0-07-052261-8.
- [19] Arleen Salles, Kathinka Evers, and Michele Farisco. Anthropomorphism in AI. *AJOB Neuroscience*, 11(2):88–95, 2020. doi: 10.1080/21507740.2020.1740350. URL <https://doi.org/10.1080/21507740.2020.1740350>. PMID: 32228388.
- [20] Drew V. McDermott. Artificial Intelligence Meets Natural Stupidity. *SIGART Newsletter*, 57:4–9, apr 1976. doi: 10.1145/1045339.1045340. URL <https://dl.acm.org/doi/epdf/10.1145/1045339.1045340>.
- [21] Melanie Mitchell. *Artificial Intelligence: A Guide for Thinking Humans*. Farrar, Straus and Giroux, New York, 2019. ISBN 978-0374257835.
- [22] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Comput. Surv.*, 55(12), March 2023.

- ISSN 0360-0300. doi: 10.1145/3571730. URL <https://doi.org/10.1145/3571730>.
- [23] Muru Zhang, Ofir Press, William Merrill, Alisa Liu, and Noah A. Smith. How language model hallucinations can snowball. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 59670–59684. PMLR, 21–27 Jul 2024. URL <https://proceedings.mlr.press/v235/zhang24ay.html>.
- [24] European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689>, June 2024. OJ L, 13 June 2024.
- [25] Philip M. Morse and George E. Kimball. *Methods of Operations Research*. MIT Press, Cambridge, MA, 1951.
- [26] Frederick S. Hillier and Gerald J. Lieberman. *Introduction to Operations Research*. McGraw-Hill, 11 edition, 2021.
- [27] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, 2 edition, 2018. ISBN 9780262039246.
- [28] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. MIT Press, Cambridge, MA, 2nd edition, 2018. ISBN 9780262039406. URL <https://cs.nyu.edu/~mohri/mlbook/>.
- [29] Arnold Wolfers. “national security” as an ambiguous symbol. *Political Science Quarterly*, 67(4):481–502, dec 1952. doi: 10.2307/2145138. URL <https://doi.org/10.2307/2145138>.
- [30] Carl von Clausewitz. *On War*. Princeton University Press, 1984. Originally published posthumously in 1832; this is the authoritative English edition.
- [31] Max Weber. Politics as a vocation. In H. H. Gerth and C. Wright Mills, editors, *From Max Weber: Essays in Sociology*. Oxford University Press, New York, 1946. Reprinted from Max Weber’s *Essays in Sociology*. Translated, edited, and with an introduction by H. H. Gerth and C. Wright Mills.
- [32] U.S. Joint Chiefs of Staff. Doctrine for the Armed Forces of the United States. Technical Report JP 1, Joint Chiefs of Staff, Washington, DC, March 2013. Incorporating Change 1, 12 July 2017.
- [33] U.S. Joint Chiefs of Staff. Joint planning. Technical Report JP 5-0, Joint Chiefs of Staff, Washington, DC, December 2020. URL https://www.esd.whs.mil/Portals/54/Documents/FOID/Reading%20Room/Joint_Staff/18-F-1152_JP_5-0_Joint_Planning_2020.pdf.

- [34] Headquarters, Department of the Army. The operations process. Technical Report ADP 5-0, Department of the Army, Washington, DC, July 2019. URL https://rdl.train.army.mil/catalog-ws/view/100.ATSC/E4166A5D-0FE7-4780-916A-A7E9B227147C-1337689957702/adp5_0.pdf.
- [35] Anthony Bellione. The heart of decision superiority: Evolve or lose – why your next war may be won or lost in seconds. Joint Air Power Competence Centre (JAPCC), Journal Edition 36, October 2023. URL <https://www.japcc.org/articles/the-heart-of-decision-superiority/>. Author listed as Col (ret.) Anthony Bellione, USAF. Accessed 2026-02-21.
- [36] Sun Tzu. *The Art of War*. Arcturus Publishing Limited, London, England, 2014. ISBN 9781784042028.
- [37] Aleksandr Vasilévič Suvorov. *Suvorov's Art of Victory*. H. Charles-Lavauzelle, Paris, 1899. URL <http://catalogue.bnf.fr/ark:/12148/cb30353002b>. Digital preservation: <ark:/12148/bpt6k86570m>.
- [38] Joseph Clark. Pentagon Official Lays Out DOD Vision for AI, February 2024. <https://www.defense.gov/News/News-Stories/Article/article/3682355/pentagon-official-lays-out-dod-vision-for-ai/>. DOD News. Accessed 17.10.2025.
- [39] North Atlantic Treaty Organization. Summary of the NATO Artificial Intelligence Strategy, October 2021. URL <https://www.nato.int/en/about-us/official-texts-and-resources/official-texts/2021/10/22/summary-of-the-nato-artificial-intelligence-strategy>. Accessed 2026-03-07.
- [40] U.S. Department of Defense. Department of Defense Data, Analytics, and Artificial Intelligence Adoption Strategy. Technical report, U.S. Department of Defense, November 2023. URL https://media.defense.gov/2023/Nov/02/2003333300/-1/-1/1/DOD_DATA_ANALYTICS_AI_ADOPTION_STRATEGY.PDF.
- [41] Koichiro Takagi. Is the PLA overestimating the potential of artificial intelligence? *Joint Force Quarterly*, 116(4):71–78, 2025. URL <https://digitalcommons.ndu.edu/joint-force-quarterly/vol116/iss4/10>. 1st Quarter 2025.
- [42] Lauri Vasankari. Tekoäly ja automaatio tulevaisuuden laivastojoukoissa. Pro gradu -tutkielma (master's thesis), Maanpuolustuskorkeakoulu (National Defence University), Helsinki, Finland, 2022. URL <https://www.doria.fi/handle/10024/185612>. Department of Military Technology (So-

- tatekniikan laitos). Persistent identifier: URN:NBN:fi-fe2022080953410. Accessed 2026-02-21.
- [43] Peter B. Checkland. *Systems Thinking, Systems Practice*. John Wiley & Sons, Chichester, UK, 1981. Revised edition, 1999.
- [44] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, and Guillaume Lample. LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. URL <https://arxiv.org/abs/2302.13971>.
- [45] Mistral AI. Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023. URL <https://arxiv.org/abs/2310.06825>.
- [46] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, and *et al.* Gemma 3 technical report. Technical Report arXiv:2503.19786, Google DeepMind, March 2025.
- [47] Shu hsien Liao. Case-based decision support system: Architecture for simulating military command and control. *European Journal of Operational Research*, 123(3):558–567, 2000. ISSN 0377-2217. doi: [https://doi.org/10.1016/S0377-2217\(99\)00109-5](https://doi.org/10.1016/S0377-2217(99)00109-5). URL <https://www.sciencedirect.com/science/article/pii/S0377221799001095>.
- [48] J. Scrimgeour. Open surveys: An information system for the improvement of international stability. *Control Engineering Practice*, 2(5):791–802, 1994. ISSN 0967-0661. doi: [https://doi.org/10.1016/0967-0661\(94\)90344-1](https://doi.org/10.1016/0967-0661(94)90344-1). URL <https://www.sciencedirect.com/science/article/pii/0967066194903441>.
- [49] Shashi D. Buluswar and Bruce A. Draper. Color machine vision for autonomous vehicles. *Engineering Applications of Artificial Intelligence*, 11(2):245–256, 1998. ISSN 0952-1976. doi: [https://doi.org/10.1016/S0952-1976\(97\)00079-1](https://doi.org/10.1016/S0952-1976(97)00079-1). URL <https://www.sciencedirect.com/science/article/pii/S0952197697000791>.
- [50] Luiz Bortolan Neto, Michael Saleh, Vanessa Pickerd, George Yianakopoulos, Zenka Mathys, and Warren Reid. Rapid mechanical evaluation of quadrangular steel plates subjected to localised blast loadings. *International Journal of Impact Engineering*, 137:103461, 2020. ISSN 0734-743X. doi: <https://doi.org/10.1016/j.ijimpeng.2019.103461>. URL <https://www.sciencedirect.com/science/article/pii/S0734743X19301708>.
- [51] Mehdi Hosseinzadeh, Jawad Tanveer, Amir Masoud Rahmani, Khursheed Aurangzeb, Efat Yousefpoor, Mohammad Sadegh Yousefpoor, Aso Darwesh, Sang-Woong Lee, and Mahmood Fazlali. A Q-learning-based smart clustering routing method in flying Ad Hoc networks. *Journal of King Saud University - Computer and Information Sciences*, 36(1):

- 101894, 2024. ISSN 1319-1578. doi: <https://doi.org/10.1016/j.jksuci.2023.101894>. URL <https://www.sciencedirect.com/science/article/pii/S1319157823004482>.
- [52] Rehan Akbani, Turgay Korkmaz, and G.V. Raju. EMLTrust: An enhanced Machine Learning based Reputation System for MANETs. *Ad Hoc Networks*, 10(3):435–457, 2012. ISSN 1570-8705. doi: <https://doi.org/10.1016/j.adhoc.2011.08.003>. URL <https://www.sciencedirect.com/science/article/pii/S1570870511001867>.
- [53] Manjit Kaur, Deepak Prashar, Leo Mrcic, and Arfat Ahmad Khan. Machine learning-based routing protocol in flying ad hoc networks: A review. *Computers, Materials and Continua*, 82(2):1615–1643, 2025. ISSN 1546-2218. doi: <https://doi.org/10.32604/cmc.2025.059043>. URL <https://www.sciencedirect.com/science/article/pii/S1546221825001298>.
- [54] Miguel Ângelo Lellis Moreira, Guilherme Vinagre Pinto de Souza, Igor Pinheiro de Araújo Costa, Wilson Tarantin Junior, Luiz Paulo Fávero, Marcos dos Santos, and Carlos Francisco Simões Gomes. Defense perception in the geopolitical scope: An exploratory study through unsupervised machine learning. *Procedia Computer Science*, 221:689–696, 2023. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2023.08.039>. URL <https://www.sciencedirect.com/science/article/pii/S1877050923007974>. Tenth International Conference on Information Technology and Quantitative Management (ITQM 2023).
- [55] Igor Pinheiro de Araújo Costa, Gabriel Custódio Rangel, Arthur Pinheiro de Araújo Costa, Gabriel Pereira de Oliveira Capela, Luiz Paulo Fávero, Carlos Francisco Simões Gomes, Marcos dos Santos, and Luiz Frederico Horácio de Souza de Barros Teixeira. Multi-criteria decision-making and machine learning techniques: A multidisciplinary analysis of the world military scenario. *Procedia Computer Science*, 242: 184–191, 2024. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2024.08.263>. URL <https://www.sciencedirect.com/science/article/pii/S1877050924019823>. 11th International Conference on Information Technology and Quantitative Management (ITQM 2024).
- [56] J.R. James and C.J. Herget. Software tools for distributed intelligent control systems. *IFAC Proceedings Volumes*, 24(10):87–90, 1991. ISSN 1474-6670. doi: <https://doi.org/10.1016/B978-0-08-041698-4.50018-4>. URL <https://www.sciencedirect.com/science/article/pii/B9780080416984500184>. 3rd IFAC Workshop on Artificial Intelligence in Real-Time Control 1991, California, USA, 23-25 September 1991.
- [57] Wen Jiang, Yihui Ren, and Yanping Wang. Improving anti-jamming

- decision-making strategies for cognitive radar via multi-agent deep reinforcement learning. *Digital Signal Processing*, 135:103952, 2023. ISSN 1051-2004. doi: <https://doi.org/10.1016/j.dsp.2023.103952>. URL <https://www.sciencedirect.com/science/article/pii/S1051200423000477>.
- [58] Alexandra Zabala-López, Mario Linares-Vásquez, Sonia Haiduc, and Yezid Donoso. A survey of data-centric technologies supporting decision-making before deploying military assets. *Defence Technology*, 42: 226–246, 2024. ISSN 2214-9147. doi: <https://doi.org/10.1016/j.dt.2024.07.012>. URL <https://www.sciencedirect.com/science/article/pii/S221491472400182X>.
- [59] Antonio A. Sánchez-Ruiz and Maximiliano Miranda. A machine learning approach to predict the winner in StarCraft based on influence maps. *Entertainment Computing*, 19:29–41, 2017. ISSN 1875-9521. doi: <https://doi.org/10.1016/j.entcom.2016.11.005>. URL <https://www.sciencedirect.com/science/article/pii/S1875952116300647>.
- [60] Jay Liebowitz and Laura C. Davis. Sharing the solution: The need for generic artificial intelligence decision support development tools in battle management. *Computers & Industrial Engineering*, 16(4):587–593, 1989. ISSN 0360-8352. doi: [https://doi.org/10.1016/0360-8352\(89\)90176-9](https://doi.org/10.1016/0360-8352(89)90176-9). URL <https://www.sciencedirect.com/science/article/pii/0360835289901769>.
- [61] Shuangxi Liu, Zehuai Lin, Wei Huang, and Binbin Yan. Current development and future prospects of multi-target assignment problem: A bibliometric analysis review. *Defence Technology*, 43: 44–59, 2025. ISSN 2214-9147. doi: <https://doi.org/10.1016/j.dt.2024.09.006>. URL <https://www.sciencedirect.com/science/article/pii/S2214914724002228>.
- [62] Abu S.M. Masud, Paul Metcalf, and Don Hommertzheim. A knowledge-based model management system for aircraft survivability analysis. *European Journal of Operational Research*, 84(1):47–59, 1995. ISSN 0377-2217. doi: [https://doi.org/10.1016/0377-2217\(94\)00317-6](https://doi.org/10.1016/0377-2217(94)00317-6). URL <https://www.sciencedirect.com/science/article/pii/0377221794003176>. Decision Technology and Intelligent Decision Support.
- [63] Sven Nõmm and Adrian Venables. Towards generation of synthetic data sets for hybrid conflict modelling. *IFAC-PapersOnLine*, 55(29): 25–30, 2022. ISSN 2405-8963. doi: <https://doi.org/10.1016/j.ifacol.2022.10.226>. URL <https://www.sciencedirect.com/science/article/pii/S2405896322022510>. 15th IFAC Symposium on Analysis, Design and Evaluation of Human Machine Systems HMS 2022.

- [64] Jean Oh, Felipe Meneguzzi, and Katia Sycara. Chapter 11 - Probabilistic Plan Recognition for Proactive Assistant Agents. In Gita Sukthankar, Christopher Geib, Hung Hai Bui, David V. Pynadath, and Robert P. Goldman, editors, *Plan, Activity, and Intent Recognition*, pages 275–288. Morgan Kaufmann, Boston, 2014. ISBN 978-0-12-398532-3. doi: <https://doi.org/10.1016/B978-0-12-398532-3.00011-7>. URL <https://www.sciencedirect.com/science/article/pii/B9780123985323000117>.
- [65] Jacques M. Perry, Raffaele Galliera, and Nirranjan Suri. A Machine Learning Approach to the Determination of Value of Information to Operators and Applications on the Tactical Edge. *Procedia Computer Science*, 205: 137–146, 2022. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2022.09.015>. URL <https://www.sciencedirect.com/science/article/pii/S1877050922008808>. 2022 International Conference on Military Communication and Information Systems (ICMCIS).
- [66] Dave Mechergui and Paramsothy Jayakumar. Efficient generation of accurate mobility maps using machine learning algorithms. *Journal of Terramechanics*, 88:53–63, 2020. ISSN 0022-4898. doi: <https://doi.org/10.1016/j.jterra.2019.12.002>. URL <https://www.sciencedirect.com/science/article/pii/S0022489819301454>.
- [67] Matheus R.F. Mendonça, Heder S. Bernardino, and Raul Fonseca Neto. Reinforcement learning with optimized reward function for stealth applications. *Entertainment Computing*, 25:37–47, 2018. ISSN 1875-9521. doi: <https://doi.org/10.1016/j.entcom.2017.12.003>. URL <https://www.sciencedirect.com/science/article/pii/S1875952117300587>.
- [68] Yan Xia, S.S. Iyengar, and N.E. Brener. An event driven integration reasoning scheme for handling dynamic threats in an unstructured environment. *Artificial Intelligence*, 95(1):169–186, 1997. ISSN 0004-3702. doi: [https://doi.org/10.1016/S0004-3702\(97\)00035-0](https://doi.org/10.1016/S0004-3702(97)00035-0). URL <https://www.sciencedirect.com/science/article/pii/S0004370297000350>.
- [69] Pamul Yadav and Shiho Kim. Chapter Four - OODA loop for learning open-world novelty problems. In Shiho Kim and Ganesh Chandra Deka, editors, *Artificial Intelligence and Machine Learning for Open-world Novelty*, volume 134 of *Advances in Computers*, pages 91–130. Elsevier, 2024. doi: <https://doi.org/10.1016/bs.adcom.2023.06.002>. URL <https://www.sciencedirect.com/science/article/pii/S0065245823000451>.
- [70] David W. Aha. The omnipresence of case-based reasoning in science and application. *Knowledge-Based Systems*, 11(5):261–273, 1998. ISSN 0950-7051. doi: [https://doi.org/10.1016/S0950-7051\(98](https://doi.org/10.1016/S0950-7051(98)

- 00066-5. URL <https://www.sciencedirect.com/science/article/pii/S0950705198000665>.
- [71] Siyuan Zhao, Jiapeng Liu, Miłosz Kadziński, Xiuwu Liao, and Yao Wang. A probabilistic preference learning approach for multiple criteria ranking in dynamic decision context. *European Journal of Operational Research*, 2025. ISSN 0377-2217. doi: <https://doi.org/10.1016/j.ejor.2025.08.008>. URL <https://www.sciencedirect.com/science/article/pii/S0377221725006241>.
- [72] Ayhan Altınors, Ferhat Yol, and Orhan Yaman. A sound based method for fault detection with statistical feature extraction in uav motors. *Applied Acoustics*, 183:108325, 2021. ISSN 0003-682X. doi: <https://doi.org/10.1016/j.apacoust.2021.108325>. URL <https://www.sciencedirect.com/science/article/pii/S0003682X21004199>.
- [73] Gracieth Cavalcanti Batista, Johnny Öberg, Osamu Saotome, Haroldo F. de Campos Velho, Elcio Hideiti Shiguemori, and Ingemar Söderquist. Machine learning algorithm partially reconfigured on FPGA for an image edge detection system. *Journal of Electronic Science and Technology*, 22(2): 100248, 2024. ISSN 1674-862X. doi: <https://doi.org/10.1016/j.jnlest.2024.100248>. URL <https://www.sciencedirect.com/science/article/pii/S1674862X24000168>.
- [74] Thierry D Fualdes and Claude J Barrouil. A common framework for reasoning on uncertainty both at symbolic and numerical levels. *Future Generation Computer Systems*, 9(4):339–347, 1993. ISSN 0167-739X. doi: [https://doi.org/10.1016/0167-739X\(93\)90036-O](https://doi.org/10.1016/0167-739X(93)90036-O). URL <https://www.sciencedirect.com/science/article/pii/0167739X93900360>.
- [75] John F. Gilmore. Military applications of expert systems. *Future Generation Computer Systems*, 1(6):403–410, 1985. ISSN 0167-739X. doi: [https://doi.org/10.1016/0167-739X\(85\)90024-X](https://doi.org/10.1016/0167-739X(85)90024-X). URL <https://www.sciencedirect.com/science/article/pii/0167739X8590024X>.
- [76] R. Sutton and G.N. Roberts. Approaches to fuzzy autopilot design optimization. *IFAC Proceedings Volumes*, 30(22):77–82, 1997. ISSN 1474-6670. doi: [https://doi.org/10.1016/S1474-6670\(17\)46493-7](https://doi.org/10.1016/S1474-6670(17)46493-7). URL <https://www.sciencedirect.com/science/article/pii/S1474667017464937>. 4th IFAC Conference on Manoeuvring and Control of Marine Craft (MCMC '97), Briujuni, Croatia, 10-12 September.
- [77] Amir Masoud Rahmani, Saqib Ali, Efat Yousefpoor, Mohammad Sadegh Yousefpoor, Danial Javaheri, Pooia Labakhsh, Omed Hassan Ahmed, Mehdi Hosseinzadeh, and Sang-Woong Lee. OLSR+: A new routing method based on fuzzy logic in flying ad-hoc networks (FANETs). *Vehicular Communi-*

- cations*, 36:100489, 2022. ISSN 2214-2096. doi: <https://doi.org/10.1016/j.vehcom.2022.100489>. URL <https://www.sciencedirect.com/science/article/pii/S2214209622000365>.
- [78] Wenyu Cai, Ziqiang Liu, Meiyang Zhang, and Chengcai Wang. Cooperative artificial intelligence for underwater robotic swarm. *Robotics and Autonomous Systems*, 164:104410, 2023. ISSN 0921-8890. doi: <https://doi.org/10.1016/j.robot.2023.104410>. URL <https://www.sciencedirect.com/science/article/pii/S0921889023000490>.
- [79] Erhan Akbal, Ayhan Akbal, Sengul Dogan, and Turker Tuncer. An automated accurate sound-based amateur drone detection method based on skinny pattern. *Digital Signal Processing*, 136:104012, 2023. ISSN 1051-2004. doi: <https://doi.org/10.1016/j.dsp.2023.104012>. URL <https://www.sciencedirect.com/science/article/pii/S1051200423001070>.
- [80] Jianan Wei, Ling Zhang, Junchao Yang, Molin Qin, Binyue Fan, Liu Yang, and Shuya Cao. Machine learning-based six-channel dual-peak photonic nose for identifying real organophosphorus nerve agents and their simulants. *Sensors and Actuators B: Chemical*, 444:138275, 2025. ISSN 0925-4005. doi: <https://doi.org/10.1016/j.snb.2025.138275>. URL <https://www.sciencedirect.com/science/article/pii/S0925400525010512>.
- [81] Jinhong K. Guo, David Van Brackle, Nicolas LoFaso, and Martin O. Hofmann. Extracting meaningful entities from human-generated tactical reports. *Procedia Computer Science*, 61:72–79, 2015. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2015.09.153>. URL <https://www.sciencedirect.com/science/article/pii/S187705091502983X>. Complex Adaptive Systems San Jose, CA November 2-4, 2015.
- [82] Xinjie Zhao and So Morikawa. Rapid assessment of large-scale urban destruction in conflict zones using hypergraph-based visual-structural machine learning. *Journal of Engineering Research*, 2024. ISSN 2307-1877. doi: <https://doi.org/10.1016/j.jer.2024.08.006>. URL <https://www.sciencedirect.com/science/article/pii/S2307187724002189>.
- [83] Mahdi Hashemi and Margeret Hall. Detecting and classifying online dark visual propaganda. *Image and Vision Computing*, 89:95–105, 2019. ISSN 0262-8856. doi: <https://doi.org/10.1016/j.imavis.2019.06.001>. URL <https://www.sciencedirect.com/science/article/pii/S0262885619300848>.
- [84] Rabiye Kılıç, Nida Kumbasar, Emin Argun Oral, and Ibrahim Yucel Ozbek. Drone classification using RF signal based spectral fea-

- tures. *Engineering Science and Technology, an International Journal*, 28: 101028, 2022. ISSN 2215-0986. doi: <https://doi.org/10.1016/j.jestch.2021.06.008>. URL <https://www.sciencedirect.com/science/article/pii/S2215098621001403>.
- [85] Hyun Kwon and Sanghyun Lee. Novel Rifle Number Recognition Based on Improved YOLO in Military Environment. *Computers, Materials and Continua*, 78(1):249–263, 2024. ISSN 1546-2218. doi: <https://doi.org/10.32604/cmc.2023.042466>. URL <https://www.sciencedirect.com/science/article/pii/S1546221824001747>.
- [86] Riddhi Mehta and Dr. Ankit Shah. An Insight into Real Time Vehicle Detection and Classification Methods using ML/DL based Approach. *Procedia Computer Science*, 235:598–605, 2024. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2024.04.059>. URL <https://www.sciencedirect.com/science/article/pii/S187705092400735X>. International Conference on Machine Learning and Data Engineering (ICMLDE 2023).
- [87] Nedyalko Petrov, Ivan Jordanov, and Jon Roe. Radar emitter signals recognition and classification with feedforward networks. *Procedia Computer Science*, 22:1192–1200, 2013. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2013.09.206>. URL <https://www.sciencedirect.com/science/article/pii/S187705091300999X>. 17th International Conference in Knowledge Based and Intelligent Information and Engineering Systems - KES2013.
- [88] William Baker, Steven Nixon, Jeffrey Banks, Karl Reichard, and Kaitlynn Castelle. Degradation analysis for diagnostic and predictive capabilities: A demonstration of progress in DoD CBM+ initiatives. *Procedia Computer Science*, 168:257–264, 2020. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2020.02.253>. URL <https://www.sciencedirect.com/science/article/pii/S1877050920303926>. Complex Adaptive Systems, Malvern, Pennsylvania, November 13-15, 2019.
- [89] Ram S. Mohril, Bhupendra S. Solanki, Makarand S. Kulkarni, and Bhupesh K. Lad. Residual life prediction in the presence of human error using machine learning. *IFAC-PapersOnLine*, 53(3):119–124, 2020. ISSN 2405-8963. doi: <https://doi.org/10.1016/j.ifacol.2020.11.019>. URL <https://www.sciencedirect.com/science/article/pii/S2405896320301634>. 4th IFAC Workshop on Advanced Maintenance Engineering, Services and Technologies - AMEST 2020.
- [90] Antonio Candelieri, Raul Sormani, Gaia Arosio, Ilaria Giordani, and Francesco Archetti. A Hyper-solution Framework for SVM Classification: Improving Damage Detection on Helicopter Fuselage Panels. *AASRI Procedia*, 4:31–36, 2013. ISSN 2212-6716. doi: <https://doi.org/10.1016/j.aasri.2013.10.006>. URL <https://www.sciencedirect.com/science/>

- article/pii/S2212671613000073. 2013 AASRI Conference on Intelligent Systems and Control.
- [91] Beibei Li, Bin Feng, and Li Chen. A graph network-based learnable simulator for spatial-temporal prediction of rigid projectile penetration. *International Journal of Impact Engineering*, 195:105123, 2025. ISSN 0734-743X. doi: <https://doi.org/10.1016/j.ijimpeng.2024.105123>. URL <https://www.sciencedirect.com/science/article/pii/S0734743X24002483>.
- [92] Donald B. Malkoff. A framework for real-time fault detection and diagnosis using temporal data. *Artificial Intelligence in Engineering*, 2(2):97–111, 1987. ISSN 0954-1810. doi: [https://doi.org/10.1016/0954-1810\(87\)90144-0](https://doi.org/10.1016/0954-1810(87)90144-0). URL <https://www.sciencedirect.com/science/article/pii/0954181087901440>.
- [93] Nikolaos Vasilikis, Rinze Geertsma, and Andrea Coraddu. A digital twin approach for maritime carbon intensity evaluation accounting for operational and environmental uncertainty. *Ocean Engineering*, 288:115927, November 2023. ISSN 0029-8018. doi: <https://doi.org/10.1016/j.oceaneng.2023.115927>. URL <https://www.sciencedirect.com/science/article/pii/S0029801823023119>.
- [94] Petros Boutselis and Ken McNaught. Using Bayesian Networks to forecast spares demand from equipment failures in a changing service logistics context. *International Journal of Production Economics*, 209:325–333, 2019. ISSN 0925-5273. doi: <https://doi.org/10.1016/j.ijpe.2018.06.017>. URL <https://www.sciencedirect.com/science/article/pii/S0925527318302615>. The Proceedings of the 19th International Symposium on Inventories.
- [95] Jason Whelan, Abdulaziz Almeahmadi, and Khalil El-Khatib. Artificial intelligence for intrusion detection systems in unmanned aerial vehicles. *Computers and Electrical Engineering*, 99:107784, 2022. ISSN 0045-7906. doi: <https://doi.org/10.1016/j.compeleceng.2022.107784>. URL <https://www.sciencedirect.com/science/article/pii/S0045790622000842>.
- [96] Maulik Sojitra, Nilesh Kumar Jadav, Rajesh Gupta, Usha Patel, Janam Patel, Sudeep Tanwar, Giovanni Pau, Fayeze Alqahtani, and Amr Tolba. Interplay of ml and blockchain for secure internet of military vehicles communication underlying 5g. *Ad Hoc Networks*, 178:103968, 2025. ISSN 1570-8705. doi: <https://doi.org/10.1016/j.adhoc.2025.103968>. URL <https://www.sciencedirect.com/science/article/pii/S1570870525002161>.
- [97] Clara Maathuis and Kasper Cools. The Role of AI in Military Cyber Security: Data Insights and Evaluation Methods. *Procedia Computer Sci-*

- ence*, 254:191–200, 2025. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2025.02.078>. URL <https://www.sciencedirect.com/science/article/pii/S1877050925004284>. International Conference on Digital Sovereignty (ICDS).
- [98] Bandar Almaslukh. Deep learning and entity embedding-based intrusion detection model for wireless sensor networks. *Computers, Materials and Continua*, 69(1):1343–1360, 2021. ISSN 1546-2218. doi: <https://doi.org/10.32604/cmc.2021.017914>. URL <https://www.sciencedirect.com/science/article/pii/S1546221821011401>.
- [99] Shahaboddin Shamshirband, Nor Badrul Anuar, Miss Laiha Mat Kiah, and Ahmed Patel. An appraisal and design of a multi-agent system based cooperative wireless intrusion detection computational intelligence technique. *Engineering Applications of Artificial Intelligence*, 26(9):2105–2127, 2013. ISSN 0952-1976. doi: <https://doi.org/10.1016/j.engappai.2013.04.010>. URL <https://www.sciencedirect.com/science/article/pii/S0952197613000766>.
- [100] Joseph C. Hoecherl, Matthew J. Robbins, Brett J. Borghetti, and Raymond R. Hill. Partially autoregressive machine learning: Development and testing of methods to predict United States Air Force retention. *Computers & Industrial Engineering*, 171:108424, 2022. ISSN 0360-8352. doi: <https://doi.org/10.1016/j.cie.2022.108424>. URL <https://www.sciencedirect.com/science/article/pii/S0360835222004612>.
- [101] Devin Wasilefsky, William N. Caballero, Chancellor Johnstone, Nathan Gaw, and Phillip R. Jenkins. Responsible machine learning for United States Air Force pilot candidate selection. *Decision Support Systems*, 180:114198, 2024. ISSN 0167-9236. doi: <https://doi.org/10.1016/j.dss.2024.114198>. URL <https://www.sciencedirect.com/science/article/pii/S0167923624000319>.
- [102] Yuhan Zhang, Yishu Wei, Yanshan Wang, Yunyu Xiao, COL Ret. Ronald K. Poropatich, Gretchen L. Haas, Yiye Zhang, Chunhua Weng, Jinze Liu, Lisa A. Brenner, James M. Bjork, and Yifan Peng. Machine learning applications related to suicide in military and veterans: A scoping literature review. *Journal of Biomedical Informatics*, 167:104848, 2025. ISSN 1532-0464. doi: <https://doi.org/10.1016/j.jbi.2025.104848>. URL <https://www.sciencedirect.com/science/article/pii/S1532046425000772>.
- [103] Richard M. Satava. Virtual reality and telepresence for military medicine. *Computers in Biology and Medicine*, 25(2):229–236, 1995. ISSN 0010-4825. doi: [https://doi.org/10.1016/0010-4825\(94\)00006-C](https://doi.org/10.1016/0010-4825(94)00006-C). URL <https://www.sciencedirect.com/science/article/pii/S001048259400006C>. Virtual Reality for Medicine.

- [104] Nizam U. Ahamed, Kellen T. Krajewski, Camille C. Johnson, Adam J. Sterczala, Julie P. Greeves, Sophie L. Wardle, Thomas J. O’Leary, Qi Mi, Shawn D. Flanagan, Bradley C. Nindl, and Chris Connaboy. Using machine learning and wearable inertial sensor data for the classification of fractal gait patterns in women and men during load carriage. *Procedia Computer Science*, 185:282–291, 2021. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2021.05.030>. URL <https://www.sciencedirect.com/science/article/pii/S1877050921011121>.
- [105] Aashay Gondalia, Dhruv Dixit, Shubham Parashar, Vijayanand Raghava, Animesh Sengupta, and Vergin Raja Sarobin. IoT-based Healthcare Monitoring System for War Soldiers using Machine Learning. *Procedia Computer Science*, 133:1005–1013, 2018. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2018.07.075>. URL <https://www.sciencedirect.com/science/article/pii/S1877050918310202>. International Conference on Robotics and Smart Manufacturing (RoSMa2018).
- [106] Mustafa Canan, Andres Sousa-Poza, and Anthony Dean. Complex adaptive behavior of hybrid teams. *Procedia Computer Science*, 114: 139–148, 2017. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2017.09.013>. URL <https://www.sciencedirect.com/science/article/pii/S1877050917318070>. Complex Adaptive Systems Conference with Theme: Engineering Cyber Physical Systems, CAS October 30 – November 1, 2017, Chicago, Illinois, USA.
- [107] T.B. Sheridan. On trusting C3I, particularly in SDI: When the PIE meets the sky. *IFAC Proceedings Volumes*, 19(8):57–62, 1986. ISSN 1474-6670. doi: <https://doi.org/10.1016/B978-0-08-034915-2.50016-0>. URL <https://www.sciencedirect.com/science/article/pii/B9780080349152500160>. IFAC Workshop on Contributions of Technology to International Conflict Resolution, Cleveland, OH, USA, 3-5 June 1986.
- [108] T. Wittig and R. Onken. Knowledge based cockpit assistant for controlled airspace flight operation. *IFAC Proceedings Volumes*, 25(9):195–200, 1992. ISSN 1474-6670. doi: [https://doi.org/10.1016/S1474-6670\(17\)50192-5](https://doi.org/10.1016/S1474-6670(17)50192-5). URL <https://www.sciencedirect.com/science/article/pii/S1474667017501925>. 5th IFAC Symposium on Analysis, Design and Evaluation of Man-Machine Systems (MMS’92), The Hague, The Netherlands, 9-11 June 1992.
- [109] Maria Grazia De Giorgi and Marco Quarta. Hybrid multigene genetic programming - artificial neural networks approach for dynamic performance prediction of an aeroengine. *Aerospace Science and Technology*, 103:105902, 2020. ISSN 1270-9638. doi: <https://doi.org/10.1016/j.ast.2020.105902>.

- 2020.105902. URL <https://www.sciencedirect.com/science/article/pii/S1270963820305848>.
- [110] Rezoanul Hafiz Chandan, Nusrat Sharmin, Muhaimin Bin Munir, Abdur Razzak, Tanvir Ahamad Naim, Tasneem Mubashshira, and Mokhlesur Rahman. AI-based small arms firing skill evaluation system in the military domain. *Defence Technology*, 29:164–180, 2023. ISSN 2214-9147. doi: <https://doi.org/10.1016/j.dt.2023.02.024>. URL <https://www.sciencedirect.com/science/article/pii/S221491472300051X>.
- [111] Lan Yang, Junqi Guo, Rongfang Bie, Anton Umek, and Anton Kos. Machine learning based accuracy prediction model for augmented biofeedback in precision shooting. *Procedia Computer Science*, 174:358–363, 2020. ISSN 1877-0509. doi: <https://doi.org/10.1016/j.procs.2020.06.099>. URL <https://www.sciencedirect.com/science/article/pii/S1877050920316197>. 2019 International Conference on Identification, Information and Knowledge in the Internet of Things.
- [112] Philip Klahr. Artificial intelligence approaches to simulation. In D.J. Murray-Smith, editor, *UKSC 84*, pages 87–92. Butterworth-Heinemann, 1984. ISBN 978-0-408-01504-2. doi: <https://doi.org/10.1016/B978-0-408-01504-2.50014-4>. URL <https://www.sciencedirect.com/science/article/pii/B9780408015042500144>.
- [113] B.M. Knapp, A.R. Dudley, and J.S. Ryder. Modelling techniques for simulation of submarine engagements. *Mathematical and Computer Modelling*, 12(8):1048–1049, 1989. ISSN 0895-7177. doi: [https://doi.org/10.1016/0895-7177\(89\)90219-7](https://doi.org/10.1016/0895-7177(89)90219-7). URL <https://www.sciencedirect.com/science/article/pii/0895717789902197>.
- [114] Geethanjali Govindarajan, Gulam Nabi Alsath Mohammed, Abhishek Premanand, and Kirubaveni Savarimuthu. Optimum design of a novel Ku-band rasorber for RADAR warfare systems using ML neural network. *AEU - International Journal of Electronics and Communications*, 185:155453, 2024. ISSN 1434-8411. doi: <https://doi.org/10.1016/j.aeue.2024.155453>. URL <https://www.sciencedirect.com/science/article/pii/S143484112400339X>.
- [115] Jinrui Li, Guohua Wu, and Ling Wang. A comprehensive survey of weapon target assignment problem: Model, algorithm, and application. *Engineering Applications of Artificial Intelligence*, 137:109212, 2024. ISSN 0952-1976. doi: <https://doi.org/10.1016/j.engappai.2024.109212>. URL <https://www.sciencedirect.com/science/article/pii/S0952197624013708>.
- [116] Moayad Aloqaily, Ouns Bouachir, and Ismaeel Al Ridhawi. UAV-supported communication: Current and prospective solutions. *Vehicular Communications*, 54:100923, 2025. ISSN 2214-2096. doi: <https://doi.org/10.1016/>

- j.vehcom.2025.100923. URL <https://www.sciencedirect.com/science/article/pii/S2214209625000506>.
- [117] Xiaoyan Wang, Jingjing Yang, Zixiao Peng, Shunfang Wang, and Ming Huang. Hilbert signal envelope-based multi-features methods for GNSS spoofing detection. *Computers & Security*, 144:103959, 2024. ISSN 0167-4048. doi: <https://doi.org/10.1016/j.cose.2024.103959>. URL <https://www.sciencedirect.com/science/article/pii/S0167404824002645>.
- [118] Harriet H. Kagiwada. Military modelling and computing: Where do we go from here? *Mathematical and Computer Modelling*, 11:693–698, 1988. ISSN 0895-7177. doi: [https://doi.org/10.1016/0895-7177\(88\)90582-1](https://doi.org/10.1016/0895-7177(88)90582-1). URL <https://www.sciencedirect.com/science/article/pii/0895717788905821>.
- [119] North Atlantic Treaty Organization. Summary of NATO’s revised Artificial Intelligence (AI) strategy, July 2024. URL <https://www.nato.int/en/about-us/official-texts-and-resources/official-texts/2024/07/10/summary-of-natos-revised-artificial-intelligence-ai-strategy>. Accessed 2026-03-07.
- [120] U.S. Department of Defense. Ethical Principles for Artificial Intelligence. <https://www.defense.gov/News/Releases/Release/Article/2091996/>, February 2020. Accessed 2026-03-07.
- [121] National Security Commission on Artificial Intelligence. Final Report. Technical report, NSCAI, 2021. URL <https://reports.nscai.gov/final-report/>.
- [122] North Atlantic Treaty Organization. NATO DIANA Announces Companies Chosen for the Next Phase of the Accelerator Programme. https://www.nato.int/cps/en/natohq/news_228518.htm, 2024. NATO News Release, 9 Oct 2024.
- [123] U.S. Department of Defense. AUKUS Pillar II Milestones Hint at Future Integrated Autonomous, Artificial Intelligence Operations. <https://www.defense.gov/News/Releases/Release/Article/3867890/>, August 2024.
- [124] U.S. Department of Defense. DoD Directive 3000.09: Autonomy in Weapon Systems. Technical report, Department of Defense, January 2023. URL <https://media.defense.gov/2023/Jan/25/2003149928/-1/-1/0/DOD-DIRECTIVE-3000.09-AUTONOMY-IN-WEAPON-SYSTEMS.PDF>.
- [125] International Committee of the Red Cross. Submission on Autonomous Weapon Systems to the United Nations Secretary-General. <https://www.icrc.org/sites/default/files/wysiwyg/war-and-law/>

- icrc_submission_on_autonomous_weapons_to_unsg.pdf, 2024.
- [126] UK Ministry of Defence. Defence Artificial Intelligence Strategy. <https://www.gov.uk/government/publications/defence-artificial-intelligence-strategy>, 2022.
- [127] UK Ministry of Defence. JSP 936 V1.1: Dependable Artificial Intelligence (AI) in Defence—Part 1: Directive. Technical report, UK Ministry of Defence, November 2024. URL https://assets.publishing.service.gov.uk/media/6735fc89f6920bfb5abc7b62/JSP936_Part1.pdf.
- [128] Edmund J. Burke, Kristen Gunness, Cortez A. Cooper III, and Mark Cozad. People’s Liberation Army Operational Concepts. Technical Report RR-A394-1, RAND Corporation, 2020. URL https://www.rand.org/pubs/research_reports/RRA394-1.html.
- [129] Heiko Borchert, Torben Schütz, and Joseph Verbovsky, editors. *The Very Long Game: 25 Case Studies on the Global State of Defense AI*. Contributions to Security and Defence Studies. Springer, 2024. doi: 10.1007/978-3-031-58649-1. URL <https://doi.org/10.1007/978-3-031-58649-1>.
- [130] Peter Highnam. The Defense Advanced Research Projects Agency’s Artificial Intelligence Vision. *AI Magazine*, 41(2):83–85, 2020. doi: 10.1609/aimag.v41i2.5301. URL <https://doi.org/10.1609/aimag.v41i2.5301>.
- [131] Allen Newell. Some problems of basic organization in problem-solving programs. Technical Report RM-3283-PR, RAND Corporation, Santa Monica, CA, 1962. URL https://www.rand.org/pubs/research_memoranda/RM3283.html.
- [132] A. Klinger. Natural language, linguistic processing, and speech understanding: Recent research and future goals. Technical Report R-1377, RAND Corporation, Santa Monica, CA, 1973. URL <https://www.rand.org/pubs/reports/R1377.html>.
- [133] Robert M. Kaplan. The mind system: A grammar-rule language. Technical Report RM-6265/1-PR, RAND Corporation, Santa Monica, CA, 1970. URL https://www.rand.org/pubs/research_memoranda/RM6265z1.html.
- [134] M. E. Maron. Artificial intelligence and brain mechanisms. Technical Report RM-3522-PR, RAND Corporation, Santa Monica, CA, 1963. URL https://www.rand.org/pubs/research_memoranda/RM3522.html.
- [135] Allan M. Din. *Arms and Artificial Intelligence: Weapons and Arms Control Applications of Advanced Computing*. Oxford University Press for SIPRI, Oxford, 1987. ISBN 0-19-829122-1.

- [136] Carl G. Jacobsen. *The Uncertain Course: New Weapons, Strategies and Mindsets*. Oxford University Press for SIPRI, Oxford, 1987. ISBN 0-19-829115-9.
- [137] Li Ang Zhang, Yusuf Ashpari, and Anthony Jacques. Understanding the Limits of Artificial Intelligence for Warfighters: Volume 3, Predictive Maintenance. Research Report RRA1722-3, RAND Corporation, Santa Monica, CA, 1 2024. URL https://www.rand.org/pubs/research_reports/RRA1722-3.html.
- [138] Joshua Steier, Erik Van Hegewald, Anthony Jacques, Gavin S. Hartnett, and Lance Menthe. Understanding the limits of artificial intelligence for warfighters: Volume 2, distributional shift in cybersecurity datasets. Research Report RRA1722-2, RAND Corporation, Santa Monica, CA, 1 2024. URL https://www.rand.org/pubs/research_reports/RRA1722-2.html.
- [139] Theodora Ogden, Anna Knack, Mélusine Lebret, James Black, and Vasilios Mavroudis. The Role of the Space Domain in the Russia–Ukraine War: The Impact of Converging Space and AI Technologies. External Publication EP-70408, RAND Corporation and RAND Europe, feb 2024. URL https://www.rand.org/pubs/external_publications/EP70408.html.
- [140] Vincent Boulanin. Mapping the innovation ecosystem driving the advance of autonomy in weapon systems. SIPRI working paper, Stockholm International Peace Research Institute, dec 2016. URL <https://www.sipri.org/sites/default/files/Mapping-innovation-ecosystem-driving-autonomy-in-weapon-systems.pdf>.
- [141] Vincent Boulanin and Maaïke Verbruggen. Mapping the development of autonomy in weapon systems. SIPRI white paper, Stockholm International Peace Research Institute, nov 2017. URL https://www.sipri.org/sites/default/files/2017-11/siprireport_mapping_the_development_of_autonomy_in_weapon_systems_1117_1.pdf.
- [142] Vincent Boulanin, Netta Goussac, Laura Bruun, and Luke Richards. Responsible military use of artificial intelligence: Can the European Union lead the way in developing best practice? SIPRI policy report, Stockholm International Peace Research Institute, nov 2020. URL https://www.sipri.org/sites/default/files/2020-11/responsible_military_use_of_artificial_intelligence.pdf.
- [143] Vincent Boulanin, Kolja Brockmann, and Luke Richards. Responsible artificial intelligence research and innovation for international peace and security. Policy report, Stockholm International Peace Research Institute (SIPRI), November 2020. URL https://www.sipri.org/sites/default/files/2020-11/responsible_artificial_intelligence_research_and_innovation_for_international_peace_and_security.pdf.

- [//www.sipri.org/publications/2020/policy-reports/responsible-artificial-intelligence-research-and-innovation-i](https://www.sipri.org/publications/2020/policy-reports/responsible-artificial-intelligence-research-and-innovation-i)
- [144] Fei Su, Vladislav Chernavskikh, and Wilfred Wan. Advancing Governance at the Nexus of Artificial Intelligence and Nuclear Weapons. SIPRI insights on peace and security, Stockholm International Peace Research Institute (SIPRI), Stockholm, March 2025. URL <https://www.sipri.org/publications/2025/sipri-insights-peace-and-security/advancing-governance-nexus-artificial-intelligence-and-nuclear-weapons>
- [145] Jon Schmid, Chad J. R. Ohlandt, and Shawn Cochran. Net technical assessment: A methodology for assessing military technology competition. Technical Report RR-A1350-1, RAND Corporation, May 2024. URL https://www.rand.org/pubs/research_reports/RRA1350-1.html.
- [146] Alexander Blanchard and Laura Bruun. Bias in military artificial intelligence. SIPRI background paper, Stockholm International Peace Research Institute (SIPRI), Stockholm, December 2024. URL <https://www.sipri.org/publications/2024/sipri-background-papers/bias-military-artificial-intelligence>.
- [147] Paul K. Davis and Paul Bracken. Artificial Intelligence for Wargaming and Modeling. *Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, 19(3):1–16, 2022. doi: 10.1177/15485129211073126. URL https://www.rand.org/pubs/external_publications/EP68860.html. RAND external publication EP-68860.
- [148] Edward Geist, Aaron B. Frank, and Lance Menthe. Understanding the Limits of Artificial Intelligence for Warfighters: Volume 4, Wargames. Research Report RRA1722-4, RAND Corporation, Santa Monica, CA, Jan 2024. URL https://www.rand.org/pubs/research_reports/RRA1722-4.html.
- [149] Headquarters, Supreme Allied Commander Transformation (HQ SACT). *NATO Wargaming Handbook*. NATO Allied Command Transformation, Norfolk, VA, USA, 2023. URL <https://paxsims.wordpress.com/wp-content/uploads/2023/09/nato-wargaming-handbook-202309.pdf>. First version publicly disclosed.
- [150] James Black, Rebecca Lucas, John Kennedy, Megan Hughes, and Harper Fine. Command and control in the future: Concept paper: Grappling with complexity. Technical Report RR-A2476-1, RAND Corporation, January 2024. URL https://www.rand.org/pubs/research_reports/RRA2476-1.html.
- [151] David Schulker, Matthew Walsh, Avery Calkins, Monique Graham, Cheryl K.

- Montemayor, Albert A. Robbert, Sean Robson, Claude M. Setodji, Joshua Snoke, Joshua Williams, and Li Ang Zhang. Leveraging machine learning to improve human resource management: Volume 1, key findings and recommendations for policymakers. Research Report RR-A1745-1, RAND Corporation, Santa Monica, CA, 2 2024. URL https://www.rand.org/pubs/research_reports/RRA1745-1.html.
- [152] Irineo Cabrerros, Joshua Snoke, Osonde A. Osoba, Inez Khan, and Marc N. Elliott. Advancing Equitable Decisionmaking for the Department of Defense Through Fairness in Machine Learning. Research Report RR-A1542-1, RAND Corporation, 6 2023. URL https://www.rand.org/pubs/research_reports/RRA1542-1.html.
- [153] Vincent Boulanin, editor. *The Impact of Artificial Intelligence on Strategic Stability and Nuclear Risk, Volume I, Euro-Atlantic perspectives*. Stockholm International Peace Research Institute, may 2019. URL <https://www.sipri.org/sites/default/files/2019-05/sipri1905-ai-strategic-stability-nuclear-risk.pdf>.
- [154] Petr Topychkanov, editor. *The Impact of Artificial Intelligence on Strategic Stability and Nuclear Risk, Volume III, South Asian Perspectives*. Stockholm International Peace Research Institute, apr 2020. URL https://www.sipri.org/sites/default/files/2020-04/impact_of_ai_on_strategic_stability_and_nuclear_risk_vol_iii_topychkanov_1.pdf.
- [155] Lora Saalman, editor. *The Impact of Artificial Intelligence on Strategic Stability and Nuclear Risk, Volume II, East Asian Perspectives*. Stockholm International Peace Research Institute, apr 2020. URL https://www.sipri.org/sites/default/files/2019-10/the_impact_of_artificial_intelligence_on_strategic_stability_and_nuclear_risk_volume_ii.pdf.
- [156] Vincent Boulanin, Lora Saalman, Petr Topychkanov, Fei Su, and Moa Peldán Carlsson. Artificial intelligence, strategic stability and nuclear risk. SIPRI report, Stockholm International Peace Research Institute, jun 2020. URL https://www.sipri.org/sites/default/files/2020-06/artificial_intelligence_strategic_stability_and_nuclear_risk.pdf.
- [157] Vladislav Chernavskikh. Nuclear weapons and artificial intelligence: Technological promises and practical realities. SIPRI background paper, Stockholm International Peace Research Institute, sep 2024. URL https://www.sipri.org/sites/default/files/2024-09/bp_2409_ai-nuclear.pdf.
- [158] Nivedita Raju and Wilfred Wan. Escalation risks at the space–nuclear nexus. SIPRI insights on peace and security, Stockholm In-

- ternational Peace Research Institute, feb 2021. URL https://www.sipri.org/sites/default/files/2024-02/2402_rpp_space-nuclear_nexus.pdf.
- [159] Edward Geist and Andrew J. Lohn. How Might Artificial Intelligence Affect the Risk of Nuclear War? Technical Report PE-296-RC, RAND Corporation, April 2018. URL <https://www.rand.org/pubs/perspectives/PE296.html>.
- [160] Center for a New American Security. Paul scharre wins colby award for book “army of none”. Press release, April 2019. URL <https://www.cnas.org/press/press-release/paul-scharre-wins-colby-award-for-book-army-of-none>. Accessed 2026-03-07.
- [161] Paul Scharre. *Army of None: Autonomous Weapons and the Future of War*. W. W. Norton & Company, 2018. ISBN 9780393608984.
- [162] Sam J. Tangredi and George V. Galdorisi. *AI at War: How Big Data, Artificial Intelligence, and Machine Learning Are Changing Naval Warfare*. Naval Institute Press, 2021. ISBN 9781682476345.
- [163] Christian Brose. *The Kill Chain: Defending America in the Future of High-Tech Warfare*. Hachette Books, New York, NY, 2020.
- [164] Jonathan Wong. Book review: “The Kill Chain: Defending America in the Future of High-Tech Warfare”. RAND Commentary, July 2020. URL <https://www.rand.org/pubs/commentary/2020/07/book-review-the-kill-chain-defending-america-in-the.html>. Accessed 2026-03-07.
- [165] NDU Press. The Kill Chain: Defending America in the Future of High-Tech Warfare. National Defense University Press news / review, 2020. URL <https://ndupress.ndu.edu/Media/News/News-Article-View/Article/2541993/the-kill-chain-defending-america-in-the-future-of-high-tech-> Accessed 2026-03-07.
- [166] Paul Scharre. *Four Battlegrounds: Power in the Age of Artificial Intelligence*. W. W. Norton & Company, February 2023. ISBN 9780393866865. Hardcover edition.
- [167] Robert J. Bunker. Book review: Four battlegrounds: Power in the age of artificial intelligence. Parameters Bookshelf (U.S. Army War College Press), October 2023. URL https://press.armywarcollege.edu/parameters_bookshelf/27/. Publication date: 2023-10-17. Accessed 2026-03-07.
- [168] U.S. Department of Defense. Data, Analytics, and Artificial Intelligence Adoption Strategy: Accelerating Decision Advantage. Technical report, Department of Defense, June 2023. URL

- https://media.defense.gov/2023/Nov/02/2003333300/-1/-1/1/DOD_DATA_ANALYTICS_AI_ADOPTION_STRATEGY.PDF.
- [169] Augusta Ada Lovelace. Sketch of the Analytical Engine Invented by Charles Babbage, by L. F. Menabrea, with Notes by the Translator. *Scientific Memoirs, Selected from the Transactions of Foreign Academies of Science and Learned Societies*, 3:666–731, 1843. URL <https://www.gutenberg.org/ebooks/75107>. Translation and notes by Ada Lovelace.
- [170] George Boole. *An Investigation of the Laws of Thought on Which are Founded the Mathematical Theories of Logic and Probabilities*. Walton and Maberly, London, 1854. URL <https://www.gutenberg.org/files/15114/15114-pdf.pdf>. Modern digital version published by Project Gutenberg.
- [171] Thomas Bayes. An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53:370–418, 1763. Edited and published posthumously by Richard Price.
- [172] New World Encyclopedia contributors. Euclid, 2023. URL <https://www.newworldencyclopedia.org/entry/Euclid>. Accessed: 2026-03-07.
- [173] Christopher M. Bishop and Hugh Bishop. *Deep Learning: Foundations and Concepts*. Springer, 2024.
- [174] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge university press, 2 edition, 2013. doi: <https://doi.org/10.1017/CBO9780511803161>.
- [175] Alan M. Turing. Computing machinery and intelligence. *Mind*, 59(236):433–460, 1950. doi: 10.1093/mind/LIX.236.433.
- [176] A. Newell and H. Simon. The logic theory machine—a complex information processing system. *IRE Transactions on Information Theory*, 2(3):61–79, 1956. doi: 10.1109/TIT.1956.1056797. URL <https://ieeexplore.ieee.org/document/1056797>.
- [177] Thomas Haigh. How the ai boom went bust. *Communications of the ACM*, January 2024. doi: 10.1145/3634901. URL <https://cacm.acm.org/opinion/how-the-ai-boom-went-bust/>.
- [178] Edward A. Feigenbaum, Bruce G. Buchanan, and Joshua Lederberg. On generality and problem solving: A case study using the dendral program. In B. Meltzer and D. Michie, editors, *Machine Intelligence 6*, pages 165–190. Edinburgh University Press, Edinburgh, 1971.
- [179] Edward H. Shortliffe. *Computer-Based Medical Consultations: MYCIN*. Elsevier, 1976. ISBN 978-0444569691. URL <https://www.shortliffe.net/Shortliffe-1976/MYCIN%20thesis%20Book.htm>.
- [180] Frank Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958. doi: 10.1037/h0042519.

- [181] Marvin Minsky and Seymour Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA, 1969. ISBN 9780262343930. URL <https://direct.mit.edu/books/monograph/3132/PerceptronsAn-Introduction-to-Computational>.
- [182] Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory, COLT '92*, page 144–152, New York, NY, USA, 1992. Association for Computing Machinery. ISBN 089791497X. doi: 10.1145/130385.130401. URL <https://doi.org/10.1145/130385.130401>.
- [183] John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude E. Shannon. A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence. Technical report, Dartmouth College / Rockefeller Foundation, August 1955. Proposal formalized term “Artificial Intelligence”; workshop held summer 1956.
- [184] DeepMind. Gemini: Google deepmind’s multimodal ai model. <https://deepmind.google/technologies/gemini>, 2023. Accessed: 2026-03-08.
- [185] Anthropic. Model card and evaluations for claude models. Technical report, Anthropic, July 2023. URL <https://www-cdn.anthropic.com/files/4zrzovbb/website/bd2a28d2535bfb0494cc8e2a3bf135d2e7523226.pdf>. Claude 2 released July 2023. Accessed 2026-02-21.
- [186] xAI. Grok by xai. <https://x.ai>, 2023. Accessed: 2026-03-08.
- [187] Financial Times. Meta’s AI chief says large language models will not reach human intelligence. *Financial Times*, May 2024. URL <https://www.ft.com/content/23fab126-f1d3-4add-a457-207a25730ad9>. Accessed 22 July 2025.
- [188] Heise Online. Meta’s head of AI: Yann LeCun does not believe in the future of generative AI, July 2025. <https://www.heise.de/en/news/Meta-s-head-of-AI-Yann-LeCun-does-not-believe-in-the-future.html>. Accessed 2026-03-07.
- [189] Association for the Advancement of Artificial Intelligence (AAAI). AAAI 2025 Presidential Panel on the Future of AI Research. Presidential panel report, Association for the Advancement of Artificial Intelligence, March 2025. URL <https://aaai.org/wp-content/uploads/2025/03/AAAI-2025-PresPanel-Report-Digital-3.7.25.pdf>.
- [190] Sam Altman. Reflections, January 2025. URL <https://blog.samaltman.com/reflections>. Accessed 2026-03-07.
- [191] Anca Dragan, Rohin Shah, Four Flynn, and Shane Legg. Taking a responsi-

- ble path to agi, April 2025. URL <https://deepmind.google/blog/taking-a-responsible-path-to-agi/>. Accessed 2026-03-07.
- [192] Dario Amodei. *Machines of Loving Grace: How AI Could Transform the World for the Better*, October 2024. URL <https://www.darioamodei.com/essay/machines-of-loving-grace>. Accessed 2026-03-07.
- [193] Parshin Shojaee*†, Iman Mirzadeh*, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad Farajtabar. *The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity*, 2025. URL <https://ml-site.cdn-apple.com/papers/the-illusion-of-thinking.pdf>.
- [194] Marina Mancoridis, Bec Weeks, Keyon Vafa, and Sendhil Mullainathan. *Potemkin Understanding in Large Language Models*. In *Proceedings of the 42nd International Conference on Machine Learning (ICML 2025)*, Vancouver, Canada, July 2025. doi: 10.48550/arXiv.2506.21521. URL <https://icml.cc/virtual/2025/poster/44050>. Poster #E-2703.
- [195] ARC Prize Foundation. *Arc-agi leaderboard*. <https://arcprize.org/leaderboard>, 2025. Accessed: 2025-06-20.
- [196] Peter Chapman, Julian Clinton, Randy Kerber, Thomas Khabaza, Thomas Reinartz, Colin Shearer, and Rüdiger Wirth. *CRISP-DM 1.0: Step-by-step data mining guide*. Technical report, CRISP-DM Consortium (NCR, Daimler-Chrysler, SPSS, OHRA), January 2000. Published under ESPRIT project; non-proprietary data mining methodology.
- [197] Flavius Vegetius Renatus. *Epitome of Military Science*. Number 16 in *Translated Texts for Historians*. Liverpool University Press, Liverpool, 1993. Translated with introduction and notes.
- [198] Niccolò Machiavelli. *The Prince*. Dover Publications, 1992. Original work published 1532; translated from Italian.
- [199] Niccolò Machiavelli. *The Art of War*. University of Chicago Press, 2005. Originally published in 1521; translated with introduction and notes by Christopher Lynch.
- [200] Nina Wilén and Lisa Strömbom. *A versatile organisation: Mapping the military’s core roles in a changing security environment*. *European Journal of International Security*, 7(1):18–37, 2022. doi: 10.1017/eis.2021.27.
- [201] Jukka Anteroinen. *Enhancing the Development of Military Capabilities by a Systems Approach*. PhD thesis, National Defence University, Finland, Helsinki, 2013. Doctoral dissertation, Publication Series No. 33, Department of Defence Technology.
- [202] Max Weber. *Economy and Society: An Outline of Interpretive Sociology*. University of California Press, 1978. Edited by Guenther Roth and Claus Wittich; translated by Frank H. Knight, Ephraim Fischhoff et al.

- [203] U.S. Army. FM 3-90 Chapter 1: The Art of Tactics, 2013. URL https://rdl.train.army.mil/catalog-ws/view/100.ATSC/17614720-DF1D-40BE-9123-F80680BF3974-1274406509298/fm3_90.pdf. Approved for public release.
- [204] Joint Chiefs of Staff. *Joint Publication 3-0: Joint Operations*. Joint Chiefs of Staff, 8 2011. URL https://edocs.nps.edu/dodpubs/topic/jointpubs/JP3/JP3_0_110811.pdf. Accessed 2026-02-22.
- [205] USAF College of Aerospace Doctrine, Research and Education (CADRE). Three levels of war. Technical report, Air University Press, Maxwell AFB, AL, 1997. URL <https://faculty.cc.gatech.edu/~tpilsch/INTA4803TP/Articles/Three%20Levels%20of%20War%3DCADRE-excerpt.pdf>. Excerpt from Air and Space Power Mentoring Guide, Vol. 1; accessed 2026-02-22.
- [206] Department of the Army. *Field Manual 100-5: Operations*. Headquarters, Department of the Army, Washington, DC, 8 1982. URL <https://cgsc.contentdm.oclc.org/digital/collection/p4013coll19/id/976>. Accessed 2026-02-22.
- [207] Headquarters, Department of the Army. *Field Manual (FM) 3-0: Operations*. Washington, DC, March 2025.
- [208] U.S. Army Combined Arms Center, Center for Army Lessons Learned. *23-07 (594) Military Decision-Making Process: Organizing and Conducting Planning*, November 2023. URL <https://api.army.mil/e2/c/downloads/2023/11/17/f7177a3c/23-07-594-military-decision-making-process-nov-23-public.pdf>. Available via army.mil (public release).
- [209] NATO Standardization Office. *APP-28: Tactical Planning for Land Forces*, b edition, 2019. Includes continuous planning cycle from receipt of mission to orders; rapid decision-making process outlined.
- [210] John R Boyd et al. *A discourse on winning and losing*, volume 400. Air University Press Maxwell Air Force Base, AL, 2018.
- [211] Grant Hammond. *The mind of war: John Boyd and American security*. Smithsonian Institution, 2001. ISBN 1-58834-178-X. Includes Boyd's Appendix "The OODA Loop" (orig. June 28, 1995 briefing).
- [212] John Ferris. Netcentric warfare, c4isr and information operations: Towards a revolution in military intelligence? *Intelligence & National Security*, 19(2): 199–225, 2004. doi: 10.1080/0268452042000302967.
- [213] Peter Checkland and Sue Holwell. *Information, Systems and Information Systems*. John Wiley & Sons, 1998.
- [214] David S. Alberts, John J. Garstka, and Frederick P. Stein. *Network Centric Warfare: Developing and Leveraging Information Superiority*. National Defense University Press, Washington, DC, 2 edition, 1999. ISBN

9781579060190. URL http://dodccrp.org/files/Alberts_NCW.pdf. CCRP PDF; accessed 2026-02-22.
- [215] Center for Development of Security Excellence. NATO information short: Student guide. Student Guide IFS0007, Defense Counterintelligence and Security Agency, November 2024. URL <https://www.cdse.edu/Portals/124/Documents/student-guides/shorts/IFS0007-guide.pdf>.
- [216] Susmit Sarkar, Aishika Chakraborty, Aveek Saha, and Anushka Bannerjee. Securing air-gapped systems. In *Proceedings of the International Ethical Hacking Conference 2019*, Advances in Intelligent Systems and Computing. Springer, 2020. doi: 10.1007/978-981-15-0361-0_18.
- [217] Robert W. Shirey. Internet security glossary, version 2. IETF RFC 4949, 2007. URL <https://datatracker.ietf.org/doc/html/rfc4949>. Defines “air gap” and related security terminology.
- [218] National Institute of Standards and Technology (NIST). air gap. NIST Computer Security Resource Center (CSRC) Glossary, 2007. URL https://csrc.nist.gov/glossary/term/air_gap.
- [219] James O’Donnell. We saw a demo of the new AI system powering Anduril’s vision for war. *MIT Technology Review*, December 2024. URL <https://www.technologyreview.com/2024/12/10/1108354/we-saw-a-demo-of-the-new-ai-system-powering-andurils-vision-f>. Accessed 2026-03-06.
- [220] Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector machines. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 1467–1474, 2012. URL <https://doi.org/10.48550/arXiv.1206.6389>.
- [221] Marco Barreno, Blaine Nelson, Russell Sears, Anthony D. Joseph, and J. D. Tygar. Can machine learning be secure? In *Proceedings of the 2006 ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, pages 16–25, 2006. doi: 10.1145/1128817.1128824.
- [222] Andrea De Martino. *Introduction to Modern EW Systems*, chapter 1. Artech House, 2 edition, 2018. URL https://api.pageplace.de/preview/DT0400.9781630815158_A35217969/preview-9781630815158_A35217969.pdf.
- [223] Paul Hannen. Introduction to radar and electronic warfare. In *Radar and Electronic Warfare Principles for the Non-Specialist*. The Institution of Engineering and Technology, Edison, 4 edition, 2014. ISBN 9781613530115. URL <https://doi.org/10.1049/SBRA502E>.
- [224] Nigel Walton. ‘Four-Closure’: How Amazon, Apple, Facebook & Google are driving business model innovation. In *2012 International Conference on*

- Innovation Management and Technology Research*. IEEE, May 2012. URL <https://ieeexplore.ieee.org/document/6236368>.
- [225] Ossi Ylijoki. *Big Data – Towards Data-Driven Business*. Doctoral dissertation, Lappeenranta-Lahti University of Technology (LUT), Lappeenranta, Finland, April 2019. URL <https://urn.fi/URN:ISBN:978-952-335-347-3>. Acta Universitatis Lappeenrantaensis 845.
- [226] Thomas H. Davenport and Rajeev Ronanki. Artificial intelligence for the real world. *Harvard Business Review*, 96(1):108–116, 2018.
- [227] Erik Brynjolfsson and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company, New York, 2014.
- [228] Michael Chui, James Manyika, and Mehdi Miremadi. Notes from the AI frontier: Applications and value of deep learning. Technical report, McKinsey Global Institute, 2018. URL <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-lea>
- [229] Joshua S. Gans, Avi Goldfarb, and Ajay K. Agrawal. Theory Is All You Need: AI, Human Cognition, and Causal Reasoning. *Strategy Science*, 9(4):356–365, 2024. doi: 10.1287/stsc.2024.0189.
- [230] Jacob Fraden. *Handbook of Modern Sensors: Physics, Designs, and Applications*. Springer, 5th edition, 2016.
- [231] Jon S. Wilson, editor. *Sensor Technology Handbook*. Newnes, 2005. ISBN 978-0-7506-7729-5. doi: 10.1016/B978-0-7506-7729-5.X5040-X.
- [232] Armada International. The Role of the Electromagnetic Spectrum in Russian Surveillance, Offensive and Defensive Operations in Ukraine, 2024. <https://www.armadainternational.com/2024/11/the-role-of-the-electromagnetic-spectrum-in-russian-surveillance>. Accessed: 2025-06-22.
- [233] GlobalSecurity.org. The role of electromagnetic spectrum control in warfare, 1990. <https://www.globalsecurity.org/military/library/report/1990/RSC.htm>. Accessed: 2025-06-22.
- [234] Yuntao Wang, Zhou Su, Shaolong Guo, Minghui Dai, Tom H. Luan, and Yiliang Liu. A survey on digital twins: Architecture, enabling technologies, security and privacy, and future prospects. *IEEE Internet of Things Journal*, 10(17):14965–14987, 2023. doi: 10.1109/JIOT.2023.3263909. URL <https://ieeexplore.ieee.org/document/10090432>.
- [235] Sushil Jajodia, Paulo Shakarian, V.S. Subrahmanian, Vipin Swarup, and Cliff Wang, editors. *Cyber Warfare: Building the Scientific Foundation*, volume 56 of *Advances in Information Security*. Springer, 2015. doi: 10.1007/978-3-319-14039-1.

- [236] Herbert Lin and Jaclyn Kerr. On cyber-enabled information warfare and information operations. In *The Oxford Handbook of Cyber Security*, chapter 16, pages 251–272. Oxford University Press, 2021. doi: 10.1093/oxfordhb/9780198800682.013.15.
- [237] Vijay Srinivas Agneeswaran. Big-data – theoretical, engineering and analytics perspective. In S. Srinivasan, editor, *Big Data Analytics*, volume 7678 of *Lecture Notes in Computer Science*, pages 8–15. Springer, 2012. doi: 10.1007/978-3-642-35542-4_2.
- [238] Vladimir N Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity: festschrift for alexey chervonenkis*, pages 11–30. Springer, 2015.
- [239] Vladimir N Vapnik. *Statistical learning theory*. Springer, 2 edition, 1999. ISBN 0387987800. URL <https://statisticalsupportandresearch.wordpress.com/wp-content/uploads/2017/05/vladimir-vapnik-the-nature-of-statistical-learning-springer-2015.pdf>.
- [240] Yaser S Abu-Mostafa, Malik Magdon-Ismail, and Hsuan-Tien Lin. *Learning from data*. AMLBook, 2012.
- [241] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963. URL <https://www.jstor.org/stable/2282952>.
- [242] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. URL <https://doi.org/10.48550/arXiv.1412.6980>.
- [243] Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-RMSprop: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural networks for machine learning, 2012.
- [244] Adrien-Marie Legendre. On least squares. University of York, Department of Mathematics, Historical Statistics, 1959. URL <https://www.york.ac.uk/depts/maths/histstat/legendre.pdf>. Reprinted from D.E. Smith, *A Source Book in Mathematics*, Vol. II, pp. 576–579.
- [245] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics. Springer, New York, NY, 2 edition, 2009. ISBN 978-0-387-84857-0. doi: 10.1007/978-0-387-84858-7. URL <https://link.springer.com/book/10.1007/978-0-387-84858-7>.
- [246] Charu C. Aggarwal. *Data Mining: The Textbook*. Springer, 2015. ISBN 978-3-319-14141-1. doi: 10.1007/978-3-319-14142-8.
- [247] J. B. MacQueen. Some methods for classification and analysis of multivariate observations. In Lucien M. Le Cam and Jerzy Neyman, editors, *Proceed-*

- ings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pages 281–297, Berkeley, CA, 1967. University of California Press. URL <https://www.cs.cmu.edu/~bhiksha/courses/mlsp.fall2010/class14/macqueen.pdf>.
- [248] Franz Aurenhammer. Voronoi diagrams — a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3):345–405, 1991. URL <https://dl.acm.org/doi/abs/10.1145/116873.116880>.
- [249] Stuart Lloyd. Least squares quantization in PCM. *IEEE transactions on information theory*, 28(2):129–137, 1982. URL <https://ieeexplore.ieee.org/document/1056489>.
- [250] Karl Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901. doi: <https://doi.org/10.1080/14786440109462720>.
- [251] Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24(6):417–441, 1933. doi: [10.1037/h0071325](https://doi.org/10.1037/h0071325). URL <https://doi.org/10.1037/h0071325>.
- [252] C. J. Van Rijsbergen. *Information Retrieval*. Butterworth-Heinemann, Newton, MA, USA, 2nd edition, 1979. URL https://openlib.org/home/krichel/courses/lis618/readings/rijsbergen79_infor_retriev.pdf.
- [253] David Marvin Green, John A Swets, et al. *Signal detection theory and psychophysics*, volume 1. Wiley New York, 1966.
- [254] Tom Fawcett. An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006.
- [255] James A Hanley and Barbara J McNeil. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1): 29–36, 1982.
- [256] David H. Hubel and Torsten N. Wiesel. Receptive fields of single neurones in the cat’s striate cortex. *The Journal of Physiology*, 148(3):574–591, 1959. doi: [10.1113/jphysiol.1959.sp006308](https://doi.org/10.1113/jphysiol.1959.sp006308).
- [257] David H. Hubel and Torsten N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of Physiology*, 160(1):106–154, 1962. doi: [10.1113/jphysiol.1962.sp006837](https://doi.org/10.1113/jphysiol.1962.sp006837).
- [258] Lawrence G. Roberts. *Machine Perception of Three-Dimensional Solids*. Ph.d. dissertation, Massachusetts Institute of Technology (MIT), Cambridge, MA, 1963. Available from MIT Libraries: <http://hdl.handle.net/1721.1/11589>.
- [259] John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, 1986. doi: [10.1109/TPAMI.1986.4767851](https://doi.org/10.1109/TPAMI.1986.4767851).
- [260] T. Ojala, M. Pietikainen, and T. Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE*

- Transactions on Pattern Analysis and Machine Intelligence*, 24(7):971–987, 2002. doi: 10.1109/TPAMI.2002.1017623. URL <https://ieeexplore.ieee.org/document/1017623>.
- [261] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, second edition, 2004. ISBN 0521540518.
- [262] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. doi: 10.1023/B:VISI.0000029664.99615.94. URL <https://link.springer.com/article/10.1023/B:VISI.0000029664.99615.94>.
- [263] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. SURF: Speeded up robust features. *Computer Vision and Image Understanding*, 110(3):346–359, 2008. doi: 10.1016/j.cviu.2007.09.014.
- [264] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791. URL <https://ieeexplore.ieee.org/document/726791>.
- [265] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. URL <https://ieeexplore.ieee.org/document/7780459>.
- [266] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet classification with deep convolutional neural networks. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf.
- [267] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [268] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation, 2014. URL <https://arxiv.org/abs/1311.2524>.
- [269] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You Only Look Once: Unified, Real-Time Object Detection . In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788, Los Alamitos, CA, USA, June 2016. IEEE Computer Society. doi: 10.1109/CVPR.2016.91. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2016.91>.
- [270] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn,

- Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *arXiv preprint arXiv:2010.11929*, 2021. URL <https://arxiv.org/abs/2010.11929>.
- [271] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows, 2021. URL <https://arxiv.org/abs/2103.14030>.
- [272] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafranec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. DINOv2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. URL <https://doi.org/10.48550/arXiv.2304.07193>.
- [273] Oriane Siméoni et al. DINOv3. *arXiv preprint arXiv:2508.10104*, 2025. URL <https://doi.org/10.48550/arXiv.2508.10104>.
- [274] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud A. A. Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen A. W. M. van der Laak, Bram van Ginneken, and Clara I. Sánchez. A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42:60–88, 2017. doi: 10.1016/j.media.2017.07.005.
- [275] Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. DeepFace: Closing the gap to human-level performance in face verification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1701–1708, 2014. doi: 10.1109/CVPR.2014.220. URL <https://ieeexplore.ieee.org/document/6909616>.
- [276] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, 2015. doi: 10.1109/CVPR.2015.7298682. URL <https://ieeexplore.ieee.org/document/7298682>.
- [277] Sharada P. Mohanty, David P. Hughes, and Marcel Salathé. Using deep learning for image-based plant disease detection. *Frontiers in Plant Science*, 7: 1419, 2016. doi: 10.3389/fpls.2016.01419. URL <https://pmc.ncbi.nlm.nih.gov/articles/PMC5032846/>.
- [278] Andreas Kamilaris and Francesc X. Prenafeta-Boldú. Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147:70–90, 2018. doi: 10.1016/j.compag.2018.02.016.
- [279] Georg Klein and David Murray. Parallel Tracking and Mapping for Small AR Workspaces. In *ISMAR*, pages 225–234, 2007. doi: 10.1109/ISMAR.2007.4538852.

- [280] Steven M. LaValle. *Virtual Reality*. Cambridge University Press, 2023. URL <https://doi.org/10.1017/9781108182874>.
- [281] Christian Zimmermann and Thomas Brox. Learning to Estimate 3D Hand Pose from Single RGB Images. In *ICCV*, pages 4903–4911, 2017. doi: 10.1109/ICCV.2017.523.
- [282] Tamas Czimmermann, Gastone Ciuti, Mario Milazzo, Marcello Chiurazzi, Stefano Roccella, Calogero Maria Oddo, and Paolo Dario. Visual-based defect detection and classification approaches for industrial applications—a survey. *Sensors*, 20(5):1459, 2020. doi: 10.3390/s20051459.
- [283] Ruoxu Ren, Terence Hung, and Kay Chen Tan. A generic deep-learning-based approach for automated surface inspection. *IEEE Transactions on Cybernetics*, 48(3):929–940, 2018. doi: 10.1109/TCYB.2017.2668395. URL <https://ieeexplore.ieee.org/document/7864335>.
- [284] Paul Bergmann, Michael Fauser, David Sattlegger, and Carsten Steger. MVTec AD – A Comprehensive Dataset for Unsupervised Anomaly Detection in Industrial Inspection. In *CVPR*, pages 9592–9600, 2019. doi: 10.1109/CVPR.2019.00982.
- [285] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes. In *Robotics: Science and Systems (RSS)*, 2018. URL <https://www.roboticsproceedings.org/rss14/p19.pdf>.
- [286] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching Networks for One Shot Learning. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper_files/paper/2016/file/90e1357833654983612fb05e3ec9148c-Paper.pdf.
- [287] Jake Snell, Kevin Swersky, and Richard S. Zemel. Prototypical networks for few-shot learning. In *NIPS*, 2017. URL https://papers.nips.cc/paper_files/paper/2017/hash/cb8da6767461f2812ae4290eac7cbc42-Abstract.html.
- [288] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [289] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.
- [290] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. URL <https://ieeexplore.ieee.org/document/9157636>.

- [291] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Pires, Zhaohan Daniel Guo, Mohammad Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning. *NeurIPS*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/f3ada80d5c4ee70142b17b8192b2958e-Abstract.html>.
- [292] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *IROS*, 2017. URL <https://ieeexplore.ieee.org/document/8202133>.
- [293] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A. Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018. URL <https://proceedings.mlr.press/v80/hoffman18a>.
- [294] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and Harnessing Adversarial Examples. *ICLR*, 2015. URL <https://arxiv.org/abs/1412.6572>.
- [295] Hadjer Benmeziane, Kaoutar El Maghraoui, Hamza Ouarnoughi, Smail Niar, Martin Wistuba, and Naigang Wang. Hardware-aware neural architecture search: Survey and taxonomy. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4322–4329. International Joint Conferences on Artificial Intelligence Organization, 8 2021. doi: 10.24963/ijcai.2021/592. URL <https://doi.org/10.24963/ijcai.2021/592>. Survey Track.
- [296] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991. URL <https://doi.org/10.1162/neco.1991.3.1.79>.
- [297] Saeed Masoudnia and Reza Ebrahimpour. Mixture of experts: a literature survey. *Artificial Intelligence Review*, 42(2):275–293, 2014. doi: 10.1007/s10462-012-9338-y. URL <https://link.springer.com/article/10.1007/s10462-012-9338-y>.
- [298] Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001. doi: 10.1023/A:1010933404324.
- [299] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.
- [300] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. URL <https://arxiv.org/abs/1409.1556>.
- [301] Warren B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley Series in Probability and Statistics. John Wiley &

- Sons, 2nd edition, 2011. ISBN 978-0470604458. URL <https://doi.org/10.1002/9781118029176>.
- [302] Richard Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, 1957. ISBN 9780691146683.
- [303] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016. doi: 10.1038/nature16961. URL <https://www.nature.com/articles/nature16961>.
- [304] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019. URL <https://deepmind.com/research/highlighted-research/alphastar>.
- [305] Rémi Coulom. Efficient selectivity and backup operators in Monte-Carlo tree search. In *International conference on computers and games*, pages 72–83. Springer, 2006.
- [306] Guillaume M. J.-B. Chaslot, Jaap-T. Saito, Bruno Bouzy, Jos W. H. M. Uiterwijk, and H. Jaap van den Herik. Monte-Carlo Strategies for Computer Go. In Pierre-Yves Schobbens, Wim Vanhoof, and Glenn Schwanen, editors, *Proceedings of the 18th BeNeLux Conference on Artificial Intelligence (BNAIC'06)*, pages 83–90, Namur, Belgium, 2006. University of Namur.
- [307] Levente Kocsis and Csaba Szepesvári. Bandit Based Monte-Carlo Planning. In Johannes Fürnkranz, Tobias Scheffer, and Myra Spiliopoulou, editors, *Proceedings of the 17th European Conference on Machine Learning (ECML 2006)*, volume 4212 of *Lecture Notes in Computer Science*, pages 282–293, Berlin, Heidelberg, 2006. Springer-Verlag.
- [308] Paul F Christiano, Jan Leike, Tom B Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/d5e2c0adad503c91f91df240d0cd4e49-Paper.pdf.
- [309] Saurabh Arora and Prashant Doshi. A Survey of Inverse Reinforcement Learning: Challenges, Methods and Progress, 2020. URL <https://arxiv.org/abs/1806.06877>.
- [310] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double Q-learning. In *Proceedings of the AAAI Conference on*

- Artificial Intelligence*, AAAI'16, pages 2094–2100, 2016. URL <https://doi.org/10.1609/aaai.v30i1.10295>.
- [311] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization algorithms. *arXiv preprint arXiv:1707.06347*, July 2017. URL <https://arxiv.org/abs/1707.06347>.
- [312] Chao Yu, Akash Velu, Eugene Vinitzky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and YI WU. The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 24611–24624. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/9c1535a02f0ce079433344e14d910597-Paper-Datasets_and_Benchmarks.pdf.
- [313] John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2016. URL <https://arxiv.org/abs/1506.02438v6>.
- [314] Leland McInnes, John Healy, and James Melville. UMAP: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018. URL <https://arxiv.org/abs/1802.03426>.
- [315] Warren B. Powell. *Unified Framework for Optimization under Uncertainty*. INFORMS TutORials in Operations Research, 2022. ISBN 9780984337897. URL <https://doi.org/10.1287/educ.2016.0149>.
- [316] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1273–1282, Fort Lauderdale, FL, USA, April 2017. PMLR. URL <https://proceedings.mlr.press/v54/mcmahan17a.html>.
- [317] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients-how easy is it to break privacy in federated learning? In *Advances in Neural Information Processing Systems*, volume 33, pages 16937–16947, 2020.
- [318] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE*

- Symposium on Security and Privacy (SP)*, pages 3–18. IEEE, 2017. URL <https://ieeexplore.ieee.org/document/7958568>.
- [319] Cynthia Dwork. Differential privacy. In *International Colloquium on Automata, Languages, and Programming*, pages 1–12. Springer, 2006.
- [320] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1175–1191, 2017.
- [321] Le Trieu Phong, Yoshinori Aono, Takuya Hayashi, Lihua Wang, and Shiho Moriai. Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Transactions on Information Forensics and Security*, 13(5):1333–1345, 2018. URL <https://ieeexplore.ieee.org/document/8241854>.
- [322] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*, 2015. URL <https://arxiv.org/abs/1409.0473>.
- [323] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186, 2019.
- [324] Philip Gage. A new algorithm for data compression. *The C Users Journal archive*, 12:23–38, 1994. URL <https://api.semanticscholar.org/CorpusID:59804030>.
- [325] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, 2016. URL <https://aclanthology.org/P16-1162.pdf>.
- [326] Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71. Association for Computational Linguistics, 2018.
- [327] Yasaman Bahri, Ethan Dyer, Jared Kaplan, Jaehoon Lee, and Utkarsh Sharma. Explaining neural scaling laws. *Proceedings of the National Academy of Sciences*, 121(27):e2311878121, 2024. doi: 10.1073/pnas.2311878121. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2311878121>.

- [328] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/blfede53be364a73914f58805a001731-Paper-Conference.pdf.
- [329] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *arXiv preprint arXiv:2005.11401*, 2021. URL <https://arxiv.org/abs/2005.11401>.
- [330] Anthropic. Claude Gov models for U.S. national security customers. <https://www.anthropic.com/news/claude-gov-models-for-u-s-national-security-customers>, June 2025. Accessed: 2026-03-01.
- [331] OpenAI. Introducing ChatGPT Gov. <https://openai.com/global-affairs/introducing-chatgpt-gov/>, January 2025. Accessed: 2026-03-01.
- [332] Google Cloud. Introducing ‘Gemini for Government’: Supporting the U.S. government’s transformation with AI. <https://cloud.google.com/blog/topics/public-sector/introducing-gemini-for-government-supporting-the-us-government>, August 2025. Accessed: 2026-03-01.
- [333] Meta. How meta is supporting US national security with AI. <https://about.fb.com/news/2025/09/meta-supporting-us-national-security-with-ai/>, September 2025. Accessed: 2026-03-01.
- [334] Ministry of Defence, Finland. Government defence report. Technical Report 2024:7, Publications of the Ministry of Defence (Finland), December 2024. URL https://julkaisut.valtioneuvosto.fi/bitstream/handle/10024/166004/PLM_2024_7.pdf?sequence=4&isAllowed=y. Published by the Ministry of Defence; ISBN via URN:URN:ISBN:978-951-663-471-8.
- [335] Filippo Santoni de Sio and Jeroen van den Hoven. Meaningful human control over autonomous systems: A philosophical account. *Frontiers in Robotics*

- and AI*, 5:15, 2018. doi: 10.3389/frobt.2018.00015. URL <https://www.frontiersin.org/articles/10.3389/frobt.2018.00015>.
- [336] Advisory Council on International Affairs (AIV) and Advisory Committee on Issues of Public International Law (CAVV). Autonomous weapon systems: The need for meaningful human control. Technical Report Advisory Report No. 97, AIV/CAVV, The Hague, 2015. URL <https://www.advisorycommitteeinternationallaw.nl/documents/2015/10/12/autonomous-weapon-systems>.
- [337] National Institute of Science and Technology. Outline: Proposed Zero Draft for a Standard on AI Testing, Evaluation, Verification, and Validation. *SuperIntelligence - Robotics - Safety & Alignment*, 2, September 2025. doi: 10.70777/si.v2i5.15513.
- [338] Alexandra Souly, Javier Rando, Ed Chapman, Xander Davies, Burak Hasircioglu, Ezzeldin Shereen, Carlos Mougan, Vasilios Mavroudis, Erik Jones, Chris Hicks, Nicholas Carlini, Yarin Gal, and Robert Kirk. Poisoning Attacks on LLMs Require a Near-constant Number of Poison Samples, 2025. URL <https://arxiv.org/abs/2510.07192>.
- [339] Elham Tabassi. Artificial Intelligence Risk Management Framework (AI RMF 1.0). Technical Report NIST AI 100-1, National Institute of Standards and Technology, Gaithersburg, MD, 2023. URL https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=936225. NIST Trustworthy and Responsible AI. Accessed: 2025-10-09.
- [340] ISO/IEC. Information technology — Artificial intelligence — Guidance on risk management. International Standard, February 2023. URL <https://www.iso.org/standard/77304.html>. Edition 1. Accessed: 2025-10-10.
- [341] ISO/IEC. Information technology — Artificial intelligence — Management system. International Standard, December 2023. URL <https://www.iso.org/standard/42001>. Edition 1. Accessed: 2025-10-10.
- [342] Office of the Under Secretary of Defense for Research and Engineering. DoD Instruction 5000.89: Test and Evaluation. Technical Report DoDI 5000.89, U.S. Department of Defense, Washington, DC, November 2020. URL <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500089p.PDF>. Accessed: 2026-03-07.
- [343] Office of the Under Secretary of Defense for Research and Engineering. DoD Instruction 5000.98: Operational Test and Evaluation and Live Fire Test and Evaluation. Technical Report DoDI 5000.98, U.S. Department of Defense, Washington, DC, December 2024. URL <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500098p.PDF>.

- Supersedes OT&E/LFT&E content previously in DoDI 5000.89. Accessed: 2026-03-07.
- [344] Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, LILI YU, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. LIMA: Less Is More for Alignment. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 55006–55021. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/ac662d74829e4407ce1d126477f4a03a-Paper-Conference.pdf.
- [345] Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A. Efros. Dataset distillation, 2020. URL <https://arxiv.org/abs/1811.10959>.
- [346] Max Marion, Ahmet Üstün, Luiza Pozzobon, Alex Wang, Marzieh Fadaee, and Sara Hooker. When Less is More: Investigating Data Pruning for Pre-training LLMs at Scale, 2023. URL <https://arxiv.org/abs/2309.04564>.
- [347] Yu Gu, Jingjing Fu, Xiaodong Liu, Jeya Maria Jose Valanarasu, Noel CF Codella, Reuben Tan, Qianchu Liu, Ying Jin, Sheng Zhang, Jinyu Wang, Rui Wang, Lei Song, Guanghui Qin, Naoto Usuyama, Cliff Wong, Hao Cheng, Hohin Lee, Praneeth Sanapathi, Sarah Hilado, Jiang Bian, Javier Alvarez-Valle, Mu Wei, Khalil Malik, Jianfeng Gao, Eric Horvitz, Matthew P Lungren, Hoifung Poon, and Paul Vozila. The illusion of readiness: Stress testing large frontier models on multimodal medical benchmarks, 2025. URL <https://arxiv.org/abs/2509.18234>.
- [348] Richard Sutton. The bitter lesson, March 2019. URL <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>.
- [349] Wayne P Hughes. *Fleet Tactics and Coastal Combat*. Naval Institute Press, Annapolis, MD, 2000.
- [350] Alfred Thayer Mahan. *The Life of Nelson: The Embodiment of the Sea Power of Great Britain*. Little, Brown and Company, Boston, 1897.

