

ANNALS OF COMMAND AND CONTROL

VOL 3, NO 1, 2026

REDACTED AND PUBLISHED BY THE
NATO COMMAND AND CONTROL
CENTRE OF EXCELLENCE



Authors:

Editor in Chief: COL Mietta Groeneveld (NLD A)
Deputy-Editor in Chief: LtCol Rabia SAYLAM, PhD (TUR AF)
Managing Editor: LTC Hendrik ten Hoeve (NLD A)
Academic Editors: Dragos Calitoiu, PhD (CAN CIV), Ahmed Ghanmi, PhD (CAN CIV),
Mr. Frank Gubbels (NLD CIV)
Formatting: SM Marvin Schaap (NLD A), Max van Rijn (NLD CIV)

NATO Command and Control Centre of Excellence

Kromhout Barracks, Building K1,
Herculeslaan 1, 3584 AB Utrecht
The Netherlands
www.c2coe.org

This product is co-sponsored by the NATO Command and Control Centre of Excellence (NATO C2COE) and the Canadian Joint Warfare Centre (CJWC). It does not necessarily represent opinions or policies of NATO and/or the Sponsoring Nations of the Centre and it is designed to provide an independent opinion

All articles published by NATO C2COE are licensed under the Creative Commons Attribution 4.0 International License, if otherwise remarked. This permits anyone to copy, redistribute, remix, transmit and adapt the work provided the original work and source is appropriately cited. NATO C2COE does not grant you any other rights in relation to this journal. In other words, all other rights are reserved.

©2026 NATO C2COE. All rights reserved.

This publication may not be reproduced or used in any manner whatsoever without reference to the NATO C2COE.



CONTENT

Editorial

Col Mietta Groeneveld, MSc, Editor in Chief 4

Holonic Architectural Paradigm for Complex Modern C2

Dr. Abder Rezak Benaskeur 7

A Control-Theoretic Unification of Defensive and Offensive C2

Dr. Abder Rezak Benaskeur 45

Clausewitzian Friction in Multi-Domain C2

Antonio Legaz Sanz 63

Russian Command and Control in Information Warfare and Reflexive Influence: Implications or NATO and Counter-Reflexive Control

Dr. Ivana Stradner, Rear Admiral (Ret.) Mark Montgomery 78

Accountability under Autonomy: Human Oversight in AI-Enabled Command and Control

Kjell Schipper MSc 90

A Cognitive Framework of Assessing Leadership: In the Context of Multi-Domain Operations

Mehmet Akif Morkaya 105

Masked speech recognition among multinational talker and listener cohorts

Rebecca E Bieber, Matthew J Makashay, Benjamin M Sheffield, & Douglas S Brungart 123

Set a Goal, Allocate Resources, Don't Get in the Way: 7 Principles for effective acquisition of C2 support systems

Mr Johan Nordström, Dr Susanna Nilsson, Mr Jonas Herkevall and Ms Moamia Gideskog 148



EDITORIAL

Col Mietta Groeneveld, MSc, Editor in Chief

Welcome to the third edition of *The Annals of Command and Control*.

When we published the first issue two years ago it was an ambition. With this third edition I believe we can say it has become a platform. For this issue we received a very good number of submissions, and, what matters even more, the quality of them was high. This did not happen by itself. It is the direct result of the work of our earlier authors, who gave the journal its first credibility, and of the reviewers who took the time to read carefully and to comment honestly. To all of them I would like to say thank you, and the same thank you goes to the authors and reviewers of the present edition, who worked closely with our editorial team and respected our deadlines. I would also like to thank my editorial team for their hard work and commitment in bringing this edition together.

The security environment around us is changing, and it is changing fast. The change does not stay in one place: it reflects on technology, which develops quicker than our procedures can absorb it; on people, who are asked to adapt to new tasks, to work with far more information than before, and to understand a new art of war and a new kind of threat, and on processes, from how we acquire our systems to how we select and assess our leaders. Command and Control sits exactly in the middle of these three, and it too has to shift: away from the idea of one leader directing the fight, toward agile decision making where no single actor is explicitly in the lead, but where actions are coordinated across the multinational forces and synchronised across all instruments of power multinationally, to achieve converging effects. The eight articles of this edition each touch a different aspect of this picture.

The edition opens with “Holonc Architectural Paradigm for Complex Modern C2” by Dr. Abder Rezak Benaskeur of Defence Research and Development Canada. The article argues that the multiplicity of domains, echelons, nations and systems has pushed legacy hierarchical models beyond their breaking point, and it proposes to look at C2 as a set of recursively

nested and intertwined decision loops with a fractal, self-similar character. The holarchy, system of autonomous but cooperating holons, is offered as the architectural mechanism to balance decentralised execution at the tactical edge with centralised unity of command.

The same author continues this line in “A Control-Theoretic Unification of Defensive and Offensive C2”. Military operations are mapped onto control theory, and the dichotomy between defence and offence is shown to be isomorphic to the control-theoretic duality between disturbance rejection and setpoint tracking. From this the author derives the primacy of defence, offensive postures being anticipatory loops that are derivative of it, and argues analytically that defence is the harder problem, because the attacker only needs to find one path while the defender must hold against all of them. The paper concludes that future advantage lies in minimising C2 loop latency and maximising the fidelity of our predictive models.

“Clausewitzian Friction in Multi-Domain C2” by Mr. Antonio Legaz Sanz of Indra Sistemas takes the discussion from the abstract to the operational. Technology has not removed the fog of war; in multi-domain operations it is amplified by the cyber and social domains, deep fakes, denial-of-service attacks and algorithmic disinformation into what the author calls a hybrid fog. The article then proposes integrated palliatives: dispersed headquarters supported by Federated Mission Networking, predictive AI filtering, immersive visualisation, and a revitalised Mission Command supported by intercultural training.

This friction is not always accidental, and that is the subject of “Russian Command and Control in Information Warfare and Reflexive Influence: Implications for NATO and Counter-Reflexive Control” by Dr. Ivana Stradner and Rear Admiral (Ret.) Mark Montgomery of the Foundation for Defense of Democracies. The authors describe how the Kremlin has institutionalised reflexive control inside its own C2, and how it treats the Alliance’s decision-making process, and not only its technical architecture, as the target.

They recommend the creation of a Perception Management Warfare Assessment Unit and the adoption of counter-reflexive control as a practical answer.

Technology brings its own dilemmas, and “Accountability under Autonomy: Human Oversight in AI-Enabled Command and Control” by Kjell Schipper, MSc, and Prof. Dr. Arnoud Lagendijk of Radboud University Nijmegen looks at one of the most difficult of them. Based on a qualitative case study of the Royal Netherlands Army, it applies the Comprehensive Human Oversight Framework to the question of how accountability is maintained when elements of target selection and engagement are delegated to autonomous systems. The conclusion deserves attention: accountability depends less on direct human control at the moment of engagement, and more on the institutional production of reliability through continuous validation, transparency and lifecycle-wide integration.

From technology we move to people. In “A Cognitive Framework of Assessing Leadership in the Context of Multi-Domain Operations”, Col (Ret.) Mehmet Akif Morkaya identifies a gap that is easy to overlook: NATO doctrine and procedures for assessing leadership within C2 structures are sparse, possibly because formal leader development remains a national responsibility. His Cognitive Framework of Assessing Leadership offers an institutional mechanism to help the Alliance identify the leaders and the staff who can cope with the speed, complexity and uncertainty of the future environment.

People also have to understand each other, and this very practical dimension is examined in “Masked Speech Recognition among Multinational Talker and Listener Cohorts” by Dr. Rebecca E. Bieber, Dr. Matthew J. Makashay, Mr. Benjamin M. Sheffield and Dr. Douglas S. Brungart. Working with personnel at the NATO Trident Jupiter exercise, the study measures how the STANAG 6001 language proficiency of both the talker and the listener affects the intelligibility of English sentences in noise,

and finds that the effect is more pronounced in multitalker babble than in speech-shaped noise. It is a useful reminder that even the best architecture ends with a person who must hear the message correctly.

The edition closes with the process dimension. “Set a Goal, Allocate Resources, Don’t Get in the Way: 7 Principles for Effective Acquisition of C2 Support Systems” by Mr. Johan Nordström, Dr. Susanna Nilsson, Mr. Jonas Herkevall and Ms. Moamia Gideskog of the Swedish Defence Research Agency (FOI) treats acquisition as a wicked problem, in which organisational, human and technical factors cannot be separated from each other. Supported by a literature survey, interviews and a case study, the authors present seven principles and challenge one of our most persistent habits, namely the belief that requirements must be comprehensive before the work can start.

Eight articles, eight different entrances to the same building. Some of them are theoretical, some are empirical, and some can be used directly next Monday morning. This mixture is exactly what we hoped for when the journal was founded.

To our readers: if this edition gives you something to think about, or something to disagree with, then please consider contributing to the next one. Submissions do not need to be purely academic. A reaction to an earlier article, a case study, a lesson learned or a well-argued opposite view are all welcome.

The Annals of C2 lives from your participation.



NATO
C2COE

HOLONIC ARCHITECTURAL PARADIGM FOR COMPLEX MODERN C2

Dr. Abder Rezak Benaskeur
Defence Research and Development Canada
Department of National Defence



Modern Command and Control (C2) is conducted within a highly multi-dimensional operational landscape, characterized by a complex multiplicity of domains, echelons, partner nations, organizations, systems, and operational tempos. The resulting complexity, data saturation, and cognitive overload have pushed legacy C2 paradigms beyond their breaking point, rendering traditional hierarchical models insufficient and, in many cases, obsolete. While newly defined concepts, such as Multi-Domain Operations (MDO), Joint All-Domain C2 (JADC2), and Pan-Domain C2 (PDC2), accurately identify the critical need for integration and convergence, they lack a comprehensive theoretical framework to model the underlying complexity of modern C2 problems or to engineer the scalable systems required to manage them. This paper analyzes modern C2 through the theoretical lens of complex adaptive systems and introduces a novel framework that defines C2 as a set of recursively nested and concurrently intertwined decision loops. This structural configuration endows the entire C2 architecture with an emergent, fractal-like property, wherein every echelon and phase of the overarching C2 process operates as a self-similar decision loop. We argue that this property, alongside corresponding holonic control paradigms, offers an inherent and robust framework for engineering systems capable of navigating the chaos of modern warfare. A holarchy, defined as a system of autonomous yet cooperative holons, provides the necessary architectural mechanics to enable C2 systems to balance decentralized execution at the tactical edge with centralized unity of command. Ultimately, this approach can help resolve the central paradox of modern C2: achieving unified, synchronized action through a decentralized, scalable, and resilient system of systems.

Keywords:

Command & Control (C2); Complex Adaptive Systems; Holonic Architectures; Multi-Agent Systems; OODA Loop; JADC2; MDO; PDC2.

1 Introduction

The modern operating environment has fundamentally transformed in recent years and is no longer defined by linear, predictable, state-on-state confrontation but by an unprecedented convergence of complexity factors. The latter are overwhelming traditional, hierarchical C2 structures (Van Bommel et al., 2019) and calling for the rapid operationalization of transformative concepts and constructs like NATO's Multi-Domain Operations (MDO) (NATO Military Committee 2023; NATO Standardization Office 2025; Allied Command Transformation 2021), Canada's Pan-Domain C2 (PDC2) (Department of National Defence 2024, 2023), and United States' Combined Joint All-Domain Command & Control (CJADC2) (Department of Defense 2022; Hoehn 2022; Butchart et al. 2020). As stated by Kalloniatis in his historical review of C2 (Kalloniatis, 2025), moving beyond a domain-centric C2, to support these new concepts, fundamentally requires a very different

technological approach to fighting in and through domains, but also a willingness of militaries to step away from legacy C2 approaches.

It is important to highlight that these new C2 concepts and constructs are not simple extensions of traditional ones, such as jointness and network-centric warfare, which are mostly based on an augmentative¹ approach; but a fundamental re-imagining of the battlespace and the dynamics that drives it through a more holistic, integrative, and convergent perspective. An effective C2 must, under the new constructs, coordinate and synchronize actions, effects, activities, and operations across a wide spectrum of dimensions, discussed in Section 2, that are deeply intertwined and mutually dependent in an unprecedented way.

A central argument of this paper is that the core challenge C2 stems from a fundamental theoretical mismatch. Traditional C2 models are inherently reductionist, built on rigid, top-down hierarchies and

¹ In which services add resources and capabilities to the joint force, but seams and frictions remain.

processes intended for battlespaces that are clear or complicated and, crucially, largely predictable (Alberts and Hayes 2003). These conditions align well with the two simplest domains of the Cynefin framework (see Figure 1), where an ordered and predictable environment permits reliance on established best practices or expert analysis (Snowden and Boone, 2007).

Consequently, traditional C2 models, along with their derivative doctrines, processes, and systems, prioritize linear planning, procedural control, and sub-component optimization. This divide-and-conquer approach assumes a purely linear system where the whole equals the sum of its parts; a logic applicable only within the clear or complicated domains, under the hypothesis that by analyzing, understanding, and engineering the parts, one will have the assurance that the whole is covered. However, operational experience has repeatedly shown that this assumption fails in practice. Because these linear models are fundamentally brittle, they routinely collapse upon first contact with the modern operating environment. They cannot effectively represent, let alone manage, the non-linear, emergent, and massively parallel phenomena, spanning the complex, chaotic, or confused domains, inherent to contemporary operations.




Figure 1: Cynefin Model for Sense-Making

A C2 system designed for sequential planning and procedural deconfliction is incapable of mastering an environment defined by simultaneous and interdependent effects created by adversaries that are increasingly adaptive, learning and deceptive. What is required is not a simple series of technical upgrades to existing models and systems but a fundamental paradigm shift in how C2 itself is modelled, analyzed, understood, and designed.

This article presents a new theoretical framework that can help bridge this complexity gap. The main original contributions and theses of the paper are:

1. The modern C2 problem is best characterized not as a simple, singular decision cycle, but as a set of recursively nested and concurrently intertwined decision loops. For clarity and simplicity, this dynamic will be illustrated using Boyd's Observe-Orient-Decide-Act (OODA) loop (Boyd, 1987), as shown in Figure 2.
2. Individual branches of these loops, such as observe or orient, are themselves recursive processes, giving the entire C2 process or system an emergent fractal nature.
3. Due to its fractal structure C2 is not only clear or complicated; it is, at least complex,² exhibiting emergent behavior, non-linear interactions, and inherent unpredictability.
4. Challenges of complex adaptive systems can be addressed within the framework defined by the four dimensions: dimensionality, non-linearity, simultaneity, and heterogeneity.
5. Holonic architectures, a specialized class of multi-agent systems, provide a natural and highly effective paradigm for analyzing these challenges and engineering the solutions required to manage them.

² Chaotic and confused (disordered) domains of the Cynefin model fall outside the scope of this paper due to the inherent difficulties of designing a C2 architecture capable of addressing them in a satisfactory manner.




Figure 2: The Boyd's OODA Loop

1.1 Holon

Because the holon concept is central to both the underlying theory and the proposed C2 architecture, it is defined here at the outset.

The term holon, derived from the Greek holos,³ was coined by Arthur Koestler in (Koestler 1967). It was proposed to describe Janus-faced entities that are simultaneously autonomous wholes and dependent parts. To explain complex biological and social systems, Koestler made two key observations:

1. These systems evolve and grow to satisfy increasingly complex and changing needs by creating stable intermediate forms, which are self-reliant and more capable than the initial systems.
2. In living and organizational systems, it is generally difficult to distinguish between wholes and parts. Almost every distinguishable element is simultaneously a whole, i.e., an essentially autonomous body, and a part, that is an integrated constituent of a larger, more capable body.

For instance, a biological cell is a complete, functioning system, a whole, with its own internal regulation, yet it is also a subordinate part of a tissue, which is part of an organ, and so on.

In the context of C2, a holon represents any entity, such as a sensor, an effector, a platform, a soldier, a squad, a command post/centre, or a task force that possesses:

- **Autonomy:** The ability, even very basic, to process information and make decisions to achieve local goals without constant external direction. Holons also possess a certain degree of control over their own constituent holons.
- **Cooperativeness:** The ability and willingness to interact with other holons, as part of a higher-level holon to which it is subordinate, to achieve a wider shared objective. Such interaction can be as simple as offering service, such as sharing surveillance data feeds by a sensor holon.

1.2 Document Structure

This article is organized as follows. Section 2 discusses some of the main challenges of modern operating environments. Section 3 traces the evolution of C2, detailing how its logical progression has culminated in the current complexity crisis. Section 4 discusses the C2 problem as a fractal decision loop structure. Section 5 analyzes the complexity of this structure through the four key pillars: dimensionality, non-linearity, simultaneity, and heterogeneity. Section 6 introduces holonic architectures as the proposed framework, detailing how their features directly map, and contribute, to the management of the introduced complexity properties. Section 7 briefly discusses a few implementation pathways, enabling technologies, and key challenges, before Section 8 summarizes the contributions of the paper and suggests directions for future work.

³ Meaning whole and the suffix -on implying part or particle as in proton or neutron.

2. Modern Environment Challenges

The new reality of modern operating environments, we henceforth refer to as multi-multi, is characterized by simultaneous and interdependent operations and actions across:

- **Multiple Domains:** The operating environment is no longer defined by the three traditional, physical, domains of air, land, and sea. Most emergent C2 concepts and models include, now, space and cyber, with some adding the information environment and the electromagnetic spectrum as domains.⁴ C2 systems must now integrate and fuse heterogeneous data and information sources from various such domains (e.g., a satellite sensor feed or a social media narrative) and coordinate synergistic actions and effects across all of them (e.g., a cyber penetration tool or a forward infantry platoon). The interaction between these domains is not simply additive or subtractive but increasingly multiplicative. For instance, a cyber vulnerability can render a maritime asset defenseless, while a space outage can totally blind land forces. This considerable degree of inter-dependence requires that the set of operations and activities need to be considered holistically rather than individually as it has traditionally been done (NATO Science & Technology Organization, 2022).
- **Multiple Echelons:** The traditional delineation between the political, strategic, operational, and tactical echelons of command is rapidly dissolving. Today's operational environment is defined by a convergence of these levels, where a single tactical event, whether a split-second decision made by an operator at the edge or a localized incident amplified on social media, can trigger immediate, unforeseen, and catastrophic cascading effects at the strategic and political levels. This phenomenon is often referred to as the strategic corporal effect. Conversely, this compressed distance cuts both ways. The democratization of advanced communication networks and real-time data often induces a sub- or counter-optimal centralization of command. Armed with live feeds and pervasive data, senior leaders and headquarters staff thousands of miles removed from the physical theater can bypass the chain of command to dictate tactical adjustments on the ground, a clear destructive manifestation of micromanagement. Consequently, modern C2 systems can no longer afford to treat echelons as isolated silos or strata. They must act as a seamless, high-velocity, and coherent system that compresses the latency between the sensor at the tactical edge and the ultimate decision-maker, balancing the need for rapid top-down strategic alignment with the preservation of decentralized local initiative.
- **Multiple Tempos:** This property is a direct consequence of the multi-domain, multi-echelon nature of modern hybrid operations, which demand the simultaneous, coordinated employment of all instruments of national power: Diplomatic, Informational, Military, and Economic (DIME). Crucially, each of these elements dictates and operates within its own distinct temporal envelope. Operational actions span a massive spectrum: from hyper-velocity, machine-to-machine cyber engagements and hypersonic missile strikes measured in milliseconds, seconds, and minutes, to slow-paced, long-term strategic influence campaigns and diplomatic maneuvers unfolding over months or years. Coordinating and synchronizing effects across these vastly disparate time scales present a cognitive and computational challenge of the highest order. To establish and maintain decision advantage within such an environment, C2 systems must enable a seamless integration of effects planning, execution, and continuous assessment into a formalized battle rhythm. This cyclical orchestration serves as the architectural framework for the command's decision loop, dynamically accelerating or decelerating to dictate the overarching operational tempo and systematically disrupt the adversary's C2 cycle.

4 Although the latter two are not widely recognized as operational warfare domains.

- **Multiple Organizations:** Success in the modern operations relies on the coordination and synchronization of actions and effects within, and across, all available and relevant organizations, since for some domains, e.g., cyber, the capabilities and authorities often reside with non-military organizations. C2 systems must move beyond the pure application of (kinetic and non-kinetic) combat power to seamlessly integrate with local, regional, and national governments, as well as Non-Governmental Organizations (NGOs) and International Organizations (IOs). This introduces a requirement for soft, negotiation- and consensus-based C2 rather than purely hard, authority- and direction-based C2. In many cases, the military effort is not the primary instrument of power; instead, the military component is relegated to a supporting role.
- **Multiple Partners:** Operations are now, by default, multi-national. This is a core component of NATO and most allies' strategy, hence introducing many profound C2 challenges in technology (e.g., multi-level security, data sharing, interoperability) and in culture, such as disparate doctrines, mistrust and distrust, and differing or conflicting policies, interests, priorities, and values.

In response to the above challenges and changes in the context, most modern forces and organizations have been pursuing ambitious new operational C2 concepts and constructs, along with their enabling technological overhauls and organizational reforms. Major initiatives and projects such as MDO, JADC2, and PDC2, represent a logical and necessary evolution to achieve these goals. These new concepts and paradigms correctly and justifiably identify the objective, which is to achieve a seamless integration of sensors, decision-makers, and effectors across all relevant domains, echelons, and participants. This integration will allow the creation of convergence and decision advantage by enabling shared situational awareness and collective, coordinated, and synchronized action. One of the initiatives that aim at enabling such integration is the NATO's Federated Mission Networking (FMN) (Dechand et al., 2023). FMN is discussed in greater detail in Section 5.4.

However, as promising and as ambitious and needed as these concepts are, they all expose a profound complexity explosion. By seeking to mesh every sensor with every effector and every decision node, they generate an exponential increase in the number of potential lines of interaction, data streams, and, more importantly, authority, decision, and reporting pathways. Furthermore, the interactions along most of these lines are inherently very non-linear and unpredictable, with an emergent behavior that, if left unchecked and not controlled, can lead to unexpected and dire consequences.

The vision of total integration, while theoretically appealing and strategically necessary, has created a problem of scale and complexity that outpaces not only humans, but the very technological systems and digital infrastructures designed to manage it. The result is not decision advantage but the very real risk of data saturation, cognitive overload, and, ultimately, a total decision paralysis. To add to the problem, the modern operational environment is not just complex; it is also very ambiguous and actively contested. Intelligent, adaptive, and deceptive actors are very aggressive in exacerbating the complexity and amplifying the congestion to exploit them to their advantage and force their adversaries' C2 systems into unsustainable positions.

In such highly contested environments, forces relying on rigid, centralized C2 structures risk transforming their units and command centres into high-signature, vulnerable targets. To mitigate this systemic fragility, contemporary military architecture must transition toward decentralized, distributed, and modular C2 frameworks. The foundational work of Alberts and his colleagues on C2 agility and harmonization underscores this critical imperative, advocating for a paradigm shift from static hierarchies to dynamic, network-centric complex systems. By framing C2 not as a fixed structural artifact but as a multi-dimensional approach space, several authors (Alberts, 2003, 2011; Alberts & Hayes, 2006; NATO Science & Technology Organization, 2022) demonstrated that rigid hierarchies introduce severe cognitive bottlenecks and structural vulnerability under conditions of high complexity. To counter these limitations, they champion edge organizations capable of self-synchronization, wherein mutual adjustment, shared situational awareness, and

the deliberate harmonization of disparate entities effectively substitute for rigid, top-down control.

However, this choice will unavoidably lead to the eternal central paradox of C2, that is: how does a force achieve unity of command and synchronized action when its C2 structure must be decentralized, scalable, and modular to survive and operate at the speed of relevance?

Some doctrines, such as the mission command articulation by NATO and the US Air Force (NATO Standardization Office, 2025; U.S. Air Force, 2021), attempted to address this issue, from a military perspective, through the centralized decision, decentralized execution mechanisms. However, these doctrines are not without their own challenges and limitations.

3. Evolution of C2

The evolution of C2 (see Figure 3) reflects a continuous adaptation to the changing character of warfare, driven by shifting posture and doctrine of adversaries and technological advancements, and more importantly, the increasing complexity of the operational environment that results from them.

Historically, C2 has, for a very long time, been rooted in hierarchical structures designed for linear battlefields and prized for their stability and predictability. In the contemporary era, C2 has shifted towards more networked and agile paradigms, with the focus increasingly transitioning towards new doctrines and concepts, such as MDO, JADC2, and PDC2. The vision behind most of these concepts is to connect sensors and effectors from all available military services and systems into a single meshed capability. The objective is to enable making decisions at the speed of relevance (Department of Defense 2022).

This evolution, which on many aspects can be considered a revolution, represents for most new concepts (e.g., PDC2 and CJADC2) a move from platform- and domain-centric warfare to combat cloud, data-centric architectures, in which data is increasingly becoming the primary currency, and where the goal is to achieve decision advantage by integrating sources and effects across all (land, air, sea, space, and cyber) domains (Deptula et al. 2019, Department of National Defence 2024). The MDO concept of NATO expands beyond the military sphere to mandate the integration of all instruments of national power (NATO Military Committee, 2023).

To understand the necessity of a new C2 paradigm, one must first trace the evolutionary path that led




Figure 3: Evolution of C2

to the current situation. Historically, the design of C2 systems has been a direct reflection of the technological and operational realities of its era, intended to impose order and predictability and synchronize action in each battlespace, despite increasing scale, velocity, and complexity. This evolution, driven by technological and doctrinal shifts, has been logical and necessary, but it has now reached a critical inflection point due to unsustainable complexity.

At its core, C2 is formally defined by doctrine as the exercise of authority and direction by a properly designated commander over assigned and attached forces in the accomplishment of the mission (Joint Chiefs of Staff 2017; NATO Standardization Office 2021; Alberts and Hayes 2006). This definition bifurcates the concept into two distinct functions (McCann & Pigeau 1996):

1. Command entails the legal and moral authority to make decisions, assess risk, and assign missions. This is a fundamentally human, cognitive, and leadership-centric function.
2. Control encompasses the systems, processes, and mechanisms used to direct, monitor, and adjust the execution of the missions assigned by the commander. This is a procedural, technical, and technology-heavy function.

The structure of C2 has evolved in response to the changing nature and the expanding scale of warfare.

3.1 Single-Domain C2

For a very long time, C2 was primarily seen and treated as a single-domain endeavor. Land, air, and maritime components operated in distinct, physically separated environments with little to no interaction between them. When interactions or interferences occurred, they were very slow-paced relative to the speed at which the operations are conducted to be of any relevance or require any form or level of coordination. For instance, an anti-U-Boat mission in North Atlantic during World War II had no direct inference with a land attack taking place, at the same time, on the French German border. C2 in this era was hierarchical, with as a main purpose

deconfliction (Alberts et al., 2010), which boils down to establishing boundaries and procedures to ensure forces from the three traditional services did not interfere with one another in any harmful way. The focus was on keeping out of each other's way to maximize the independent efficiency of each service.

3.2 Joint C2

The increased interference between the three physical environments, driven mainly by the extended reach of sensors, effectors, communications, and platforms, made joint C2 necessary. The purpose of joint C2 has been to achieve coordination and synchronization of effects across component boundaries, e.g., joint air support for ground operations. The advent of precision-guided munitions and longer-range communications meant that a ship could influence a land battle, and an aircraft could influence a naval engagement. These technological advancements necessitated a transition from deconfliction to coordination (Alberts et al., 2010). The latter was often achieved through rigid Air Tasking Orders (ATOs) and pre-planned engagement zones, requiring, in most cases, days of planning cycles. Despite all the best efforts, major seams between the domains remained and significant frictions persisted.

3.3 Pan-Domain C2

The contemporary operating environment, with the full addition of cyber and space as maneuver domains, rendered necessary the transition to pan-domain C2. This is a fundamentally different paradigm that moves beyond simple coordination to true interdependence and convergence, i.e., unity of thought and unity of action. The goal is no longer just to coordinate activities and synchronize effects but to achieve integration of all effects from all available and relevant domains to create complex, composite effects that overwhelm the most capable and most prepared of adversaries. An action in the cyber domain, for example, may now be the primary setup for a kinetic effect, which in turn relies on space-based assets, all while being shaped and justified by concurrent information operations, all these to deliver a massive kinetic

effect on an enemy position. This push for pan-domain interdependence is being institutionalized through several ambitious initiatives. Central to most of these concepts are technological efforts to build sensor-to-shooter meshes that connect all sensors and effectors across all services via cloud-like data fabrics. For instance, initiatives like CJADC2 and PDC2 leverage automation and Artificial Intelligence (AI) to process sensor data and recommend Courses of Action (COA). This rapidly assigns the optimal effector, regardless of its domain or service, for the most appropriate target.

This is a very ambitious undertaking. However, if not well understood and carefully managed, this concept can potentially, and will most likely, precipitate the conduct of C2, and the development of technology-heavy systems to support it, into a crisis. By seeking to connect every sensor to every effector, the much sought-after MDO, PDC2, or CJADC2 systems risk resulting in an unmanageable amount of data, shifting the problem from one of scarcity to one of saturation. If not properly managed, the latter risks creating two primary failure points:

- **Cognitive Overload at the Command Level:** The human decision-makers become the bottlenecks. The sheer volume, velocity, and variety of data from a pan-domain sensor grid will most likely far exceed cognitive capacity of the most capable humans (Natter et al. 2010; Johnson 2023). Using John Boyd's OODA loop framework, this saturation will lead to an inability to orient and decide faster than the environment is changing. This will potentially result in decision paralysis or, worse, catastrophic decisions.
- **System Brittleness at the Control Level:** Centralized data processing and decision-making models create latency, are vulnerable to single-point failures, and cannot scale effectively in contested, high-tempo environments where communication links are Denied, Disrupted, Intermittent, and Limited (DDIL). The reliance on hyper-connectivity can quickly become the primary vulnerability, since the adversary does not need to attack the entire force, but only to strategically target a few critical links that bind it, i.e., its Achilles' heel.

The preceding discussion highlights the most prominent challenge associated with each level. For instance, cognitive overload is not exclusive to the command level and occurs at the control level as well, albeit to a lesser extent. Similarly, system brittleness also affects the command level, though its impact is generally less critical there than at the control level.

Based on the above discussion, if not designed properly, MDC2, PDC2, and CJADC2 systems and architectures risk generating an unfavorable paradox of being both hyper-connected and yet agonizingly slow, ultimately failing to deliver decision advantage at the speed of relevance, their primary purpose. This is the so-called complexity gap (IBM Global Business Services 2010), that is the point where the complexity of the environment, amplified by the very C2 system designed to manage it, exceeds the cognitive capacity of its intended end-user and the technological capacity of the system itself.

The evolutionary and ambitious path to model and engineer C2 systems has reached a point that demands new and innovative approaches. We are attempting to solve a 21st-century complexity problem with 20th-century reductionist tools. The belief that more data equates to better control is the fallacy of centralized omniscience. This is a dangerous assumption in an era where the volume of data is virtually infinite, yet the cognitive, computational, and communication capacities required to process it remain strictly finite.

The holonic C2 paradigm proposed in Section 6 is intended primarily to enable the development of an architectural framework rather than to replace existing doctrines. The intent is to provide a structured way to distribute sensing, sensemaking, and decision-making across a collection of nested, cooperating entities, allowing operationalization of established concepts, such as mission command, in a scalable and resilient way, compatible with the requirement of the complex adaptive nature of modern operations.

4. Characteristics of Modern C2

The complexity and paradox detailed in Section 3 stems mainly from a fundamental mismatch between the multi-multi reality of modern C2 and the simple, single-domain, and often single-arm era models used to characterize and engineer the technological systems meant to manage them. To bridge this gap, new models are required, which must embrace, rather than ignore, the inherent complexity of the modern operating environment and underlying technological systems. In this section, we discuss the necessity for the C2 problem, in modern context, to be best understood not as a single, linear process but as a fractal system composed of many interacting decision loops.

4.1 Revisiting the C2 Loop

C2 is executed through a series of complex decision-making processes, often conceptualized via simple models such as Boyd's OODA loop, the Joint Directors of Laboratories (JDL) fusion model (Steinberg et al., 1999), and the NATO Allied Command Transformation (ACT) C2 cycle (NATO Command & Control Centre of Excellence, 2019). While these frameworks provide a useful baseline for discussion, most C2 processes are rarely as simple or strictly sequential as their linear structures suggest. Rather, they function as an integrated, non-linear system of continuous feedback and adaptation.

Since the clarity and completeness of the discussion below are not impacted by the framework chosen, the OODA loop, illustrated in Figure 2, will be used throughout the remainder of this document.

The observation phase of the OODA equates to the sensing function, that is the gathering of cues and intelligence from the environment of interest, such as a battlespace or an area of operation. The orientation phase corresponds to the situation of interest state estimation, filtering, inference, projection, and contextual interpretation. At this stage, data is fused to create a coherent situational awareness that represents the internal model presenting the state of the controlled system, e.g., the battlespace. This is the most critical phase of

the OODA loop, as it is where analysis and synthesis occur to form mental models and situation awareness for the decision-makers. The decision segment is where the selection of options takes place. The current state of the system/battlespace is compared to its desired state, as expressed, for instance, by the Commander's intent, to generate courses of action and select the preferred option. The action phase represents the execution of the selected option as orders are given to create the desired effects within the theater of operations through the actual application of combat power.

The goal of any commander or decision-maker is to get inside the adversary/opponent's OODA loop to gain advantage by being able to go through the four phases faster and more effectively, and to saturate the opponent with more entropy than they can handle; a situation that, in practice, manifests as ambiguity, uncertainty, or dilemma. Faster, low-entropy loops of one side inside slower, higher-entropy loops of the opponent create for the former a combination of the order and tempo dominance that leads to the much sought-after decision superiority.

The simplicity of the OODA loop represents both its main strength and its major drawback. On the upside, the OODA loop offers a model, that is both elegant and accessible to non-experts, allowing them to represent and discuss the core decision cycles of (military) organizations. On the downside, it oversimplifies the complexity of (modern) C2 systems and the challenges they face, especially when people misguidedly start using the OODA loop as an analysis and design tool for C2 systems. Boyd's model was primarily conceived as a single-agent, and by a logical extension, single-domain, single-echelon, framework, describing the cognitive cycle of an individual pilot, or, similarly by extension, a single, cohesive service, or organization. It was not designed or intended to explicitly model and manage all the challenges of the C2 problem in the modern multi-multi environment, which involves dozens, if not hundreds, of different agents running their own OODA loops simultaneously across multiple echelons and domains, all of which must somehow be coordinated, synchronized, and orchestrated.

4.2 C2 as Nested & Intertwined Loops

The critical insight derived from the above discussion is that this is not a single loop but a fractal hierarchy of nested, interacting loops across all echelons, from the strategic command center to the individual soldier at the tactical edge, all operating asynchronously (see Figure 4 and Figure 5). This observation leads directly to the first thesis of this paper about the characteristics of modern C2.

Thesis 1. The modern C2 is a system of recursively nested and concurrently intertwined decision loops.

C2 functions across (strategic, operational, tactical, etc.) echelons create a natural hierarchy of decision cycles (see Figure 4), like a set of Russian dolls.

For instance, a strategic-level Joint Task Force (JTF) Headquarters (HQ) runs a slow OODA loop, with an orient phase focused on political goals and a decide phase that may take days. The output of this strategic decide, e.g., establish and maintain

maritime and air supremacy in a specific theater, such as the Strait of Hormuz, to secure merchant transit, does not translate into a direct tactical act, e.g., conduct localized anti-submarine warfare patrols and suppression of enemy air defenses missions. Instead, it provides the orienting context for the OODA loops of subordinate Component Commands, such as a Naval task force. This component-level loop, in turn, provides the orienting context for the tactical OODA loops of individual ships or air units. Each higher-echelon loop contains and directs the loops of the echelons below it.

Simultaneously, these loops, regardless of their respective echelon, run in parallel across domains, e.g., sea, air, ground, land, cyber, space, electromagnetic spectrum, and partners, such as allied, civilian authorities, etc., as illustrated in Figure 5.




Figure 4: The Nested Loop Model of C2




Figure 5: The Intertwined Loop Model

A cyber component, an army unit, and a navy coalition partner are all executing their own OODA loops. In pan-domain operations, these loops are no longer inherently deconflicted but intertwined on multiple scales and dimensions. The act from a cyber loop, such as disable enemy air defense radar, may be the primary observe input for an air component's loop, enabling its own act, e.g., strike target. The entire C2 structure is a vast, coupled mesh of concurrent, interdependent decision cycles, of which not all actors are necessarily aware.

4.3 C2 is Fractal Organization and Process

The two properties, nested and intertwined loops, discussed in the previous section, lead to the second thesis about the main characteristics of modern C2.

Thesis 2. The modern C2 is fundamentally a fractal organization and process.

In mathematics and system theory, a fractal is defined as a structure that displays self-similarity across scales. In the C2 model discussed above, this self-similarity is evident because the individual branches of the OODA loop are, themselves,

complete OODA loops.

To illustrate, consider the observe phase of a unit commander's main C2 loop (the highest level, for instance). This phase is not a passive or simple atomic action.⁵ It is an active process that is a complete OODA sub-loop. To observe their battlespace in the most effective way, the operators in charge of this task (e.g., Intelligence, Surveillance, and Reconnaissance (ISR) operators) must execute their own OODA sub-loop, all as part of the higher-level loop's observe phase:

- **Observe Phase of the Sub-Loop:** For instance, a new track with ambiguous identity appears on the operator's screen. The observe phase must then detect the informational gap and generate the requirement for additional information to fill it.
- **Orient Phase of the Sub-Loop:** The orient phase will seek to fill this informational gap by identifying available assets and selecting the best method and asset(s) to collect the missing information, e.g., a UAS patrol is my best option.

⁵ Atomic is used here in the sense that it cannot be broken further down into smaller sub-actions.

- **Decide Phase of the Sub-Loop:** The decide phase will task the selected asset(s) to gather the missing information, e.g., task UAS-4 to scan grid AA.
- **Act Phase of the Sub-Loop:** The act phase will send the tasking order to the selected asset(s). Newly collected data and information are received and fed into the orient segment of the upper loop to update its model of the situation and inform its decisions.

The output of this entire sub-loop will provide the processes and products that fulfill the role of the observe segment of the upper OODA loop. In this example, the observe phase of the main loop is conducted in a closed-loop manner, based on dynamic information requirements, rather than in an open-loop system, using static tasking process. The mechanism for closing the loop on ISR processes and asset(s) is often referred to as ISR C2.

The same fractal pattern applies to the orient phase, which involves a complex, recursive sub-loop of information-gathering, sensemaking, and analysis. An example would be a close-loop, adaptive threat evaluation process, by opposition to the open-loop, static rule-based threat assessment process.

The C2 system therefore defines a recursion of cascading OODA loops going all the way down, from the highly complex and very deliberative political and strategic levels to the simplest procedural sub-tactical level, such as the one involved in the guidance system of a missile.

In addition to the decision processes, the fractal nature extends to the organizational structure of military forces and their command teams. A modern military force can be viewed as a fractal organization composed of structurally similar decision-capable entities recurring across multiple echelons, e.g., squad, platoon, company, battalion, brigade, division, joint task force. At each level, a commander or C2 node executes an OODA loop to control and shape an assigned area of responsibility, while remaining part of a recursion of higher-level loops that provide intent, guidance, and directives.

4.4 C2 is a Complex Adaptive System

In addition to being defined by nested and intertwined loops, the C2 system is fractal, in that it is self-similar at different scales of magnification. Within this Fractal-OODA concept, each of the observe, orient, decide, and act phases is seen as being itself a complete, self-similar OODA loop operating at a different timescale. The timescale naturally decreases with the level of abstraction of the loop.

The complexity of a C2 problem, and the associated C2 system, can be measured by the depth of the corresponding fractal OODA loop. A single task, and to certain extent, a single domain operation, can be seen as low-complexity, shallow fractal loop, while a multi-domain operation, or a complex political organization, can be seen as high-complexity, high-depth fractal loop, involving thousands of these nested, parallel, and recursive loops over multiple echelons.

These observations lead to the third thesis about of modern C2, a logical conclusion of the first two.

Thesis 3. Modern C2 systems are complex adaptive systems.

A system defined by a large structure composed of autonomous, or even semi-autonomous, adaptive agents, with their corresponding C2 loops, that are nested, intertwined, and recursively structured, is the quintessence of a complex adaptive system. The multi-multi C2 environment exhibits the hallmark properties of such systems, that is many interconnected, adaptive agents whose collective, emergent behavior cannot be predicted by analyzing the individual agents in isolation (Holland 1995; Gell-Mann 1995; Waldrop 1992):

- **Adaptive Agents:** Each OODA loop, at every scale, defines an adaptive agent, trying to achieve a goal in response to its (local) world state.
- **Non-linear Interactions:** The OODA loops are tightly coupled. A small, fast tactical loop, such as a cyber-attack, can create a non-linear, cascading effect that completely changes the

environment for a slow, strategic loop, e.g., a political decision.

- Feedback: The system is dominated by feedback through the observe phase that closes every loop at every scale.
- Emergence: The state of the overall system, e.g., the battlespace, is an emergent property of all the intertwined interactions of these OODA loops. The emergent state cannot be predicted by analyzing any single loop in isolation, nor can it be controlled in a top-down, linear and deterministic fashion.

Modeling C2 as a complex adaptive system should not be seen as a simple conceptual metaphor, but a structural description that allows a move away from the counterproductive pursuit of deterministic, linear control. This perspective highlights the importance of adaptive, emergent, and non-linear processes in modern C2 and allows focusing on designing C2 systems and solutions that enable turning their inherent complexity from a catastrophic fatality into a manageable variable. This move requires architectural frameworks that, themselves, present similar properties to those of complex adaptive systems. An example of such frameworks is given by holonic control architectures. The latter, also referred to as holarchies, are composed of nested holons, or agents, that execute OODA loops at various levels of the architecture. This definition makes these architectures a natural organizational embodiment of complex adaptive systems, and by extension of fractal C2.

Before delving into the details of holonic control architectures in Section 6, Section 5 will formally analyze the C2 problem through four key and challenging properties of complexity.

5. C2 as a Complex Adaptive System

The complexity of modern military operations increasingly drove experts to view C2 through the lens of complex adaptive systems theory. Unlike linear systems where outputs are proportional to inputs,⁶ complex adaptive systems are characterized by non-linearity, emergence, and the interaction of numerous independent artifacts and agents.

Alberts (2011) addresses C2 complexity by conceptualizing modern, multi-domain operations not as a series of complicated tasks, but as a complex endeavor defined by high uncertainty, dynamic volatility, and deep interconnectedness. To systematically analyze and classify these challenges, Alberts introduces the Endeavor Space (E-space) framework, which maps a mission's characteristics across key dimensions, including dynamics, dependencies, and coupling. Because complex endeavors resist traditional hierarchies and deconflicted operational models, organizational survival depends heavily on C2 Agility, i.e., the capacity to dynamically shift its C2 approach to match the specific region of the E-space currently occupied. Other authors (Grisogono 2006; Ryan 2008; Ilachinski 2004) have further explored the implications of complex adaptive systems for C2, suggesting that military organizations must embrace adaptability and resilience rather than seeking deterministic control through rigid hierarchies. This perspective acknowledges that the fog of war is an inherent property of the system's complexity, not just a consequence of a lack of information.

This leads us to the fourth thesis of this paper.

Thesis 4. When modern C2 is treated as a complex adaptive system, its fundamental challenges can be systematically characterized and addressed along four key dimensions: dimensionality, non-linearity, simultaneity, and heterogeneity.

⁶ More rigorously, linearity is defined by the superposition principle, that is, (1) the system response to a some of inputs is the sum of the responses to individual inputs; and (2) the system response to a scaled input is the scaled output to the original input.

The following sections analyze the modern C2 problem through these four pillars that we have identified as the key sources of complexity. Understanding and characterizing these sources can help account for the above discussed properties of C2 as a complex adaptive system and avoid failure of traditional, deterministic C2 models.

5.1 Dimensionality or the Scale Challenge

Dimensionality refers to the sheer scale and multiplicity of variables, entities, and relationships operating within a system. In a multi-multi C2 environment, this complexity manifests as:

- A vast network of actors, platforms, sensors, effectors, individual warfighters, command nodes, and increasingly, AI-powered autonomous agents.
- Terabytes to petabytes of disparate data that must be continuously ingested, fused, and exploited to generate actionable understanding.
- A dense web of reporting structures, authorities, and shifting relationships across joint, inter-agency, coalition, and civil partners.
- An exponential number of potential actions, cross-domain effects, and unintended consequences that can be generated at any given moment.
- Highly convoluted lines of dependency and cascading effects between interconnected actors and systems.

This high dimensionality inevitably leads to computational intractability. No single entity, human or artificial, regardless of its processing power, can achieve full observability of the system, model its complete state, or generate an optimal, deterministic control solution. The omniscience and omnipotence sought by legacy C2 frameworks are theoretically and practically unachievable.

The challenge of dimensionality, and the resulting requirements for modern C2 systems, can be effectively framed through Ashby's Law of Requisite

Variety. This fundamental principle dictates that the amount of variety in a system must match the amount of variety in its environment to achieve control (Ashby 1956). If the dimensionality defining the environmental variety exceeds the internal capacity (or variety) of the system attempting to control it, the problem will overwhelm the system, ultimately leading to its systemic failure or destruction. To operationalize this concept in defense contexts, Niven and Capewell (2019) proposed a conceptual framework anchored in variety calculus, specifically tailored to address C2 challenges within highly complex operational environments.

Consequently, for C2 systems to remain stable and effective, their internal variety must at least match the dimensionality of the operational environment they seek to control. The NATO STO SAS-143 report (Science & Technology Organization, 2022) underscores this requirement, noting that the imperative to engage in complex endeavors: "... is driven by the nature of 21st Century missions which require the orchestration of disparate effects in multiple domains which in turn requires a set of partners with the capabilities to operate in a variety of domains and create a variety of effects."

Let V_e denote the environmental variety and V_c the variety available to the C2 system through its decision-making structures and processes. Ashby's Law then requires

$$V_c > V_e \quad (1)$$

If this inequality does not hold, the environment will overwhelm the C2 system. In other terms, the problem will overwhelm the solution, leading to system instability, failure, and collapse.

Traditional, centralized C2 systems function as low-dimensionality controllers attempting to manage high-dimensionality environments (i.e., $V_c \ll V_e$). To make the situation computationally and cognitively tractable for a centralized commander, these frameworks artificially collapse, or, worse, entirely ignore, multiple operational dimensions. This reductionist approach compresses a highly complex reality into a falsely simplified problem

space. However, in most military applications, effectively reducing environmental variety (V_e) is virtually impossible. By forcing this artificial simplification, legacy architectures inevitably fall into the reduction trap, forfeiting the systemic fidelity required to respond effectively to the actual environment.

Any viable C2 architecture must be engineered to perform despite this reality, mitigating the dimensionality challenge by decomposing the problem space and efficiently distributing the computational and decision-making burden, thereby allowing local nodes to absorb local complexity. As discussed in Section 6, we argue that holonic architectures provide an optimal framework to operationalize this approach by:

1. **Generating and Maintaining Requisite Variety:** Providing the C2 system with the structural flexibility to dynamically scale and sustain the internal variety necessary to match the environmental complexity.
2. **Distributing Control Variety:** Dispersing this control variety across multiple echelons and units to manage the cognitive and computational burden effectively.

5.2 Non-linearity or the Predictability Challenge

Non-linearity is defined by opposition to linearity. A system is considered non-linear if it does not obey the superposition principle of linear systems, as introduced previously. In a non-linear system, small, localized actions can, and frequently do, produce disproportionately large, cascading, and unpredictable effects.

Modern conflicts, particularly asymmetric, hybrid, and grey-zone warfare, are explicitly planned, shaped, and conducted in a way that aims at maximizing the exploitation of this non-linearity. An adversary's act is no longer just a kinetic strike; but may also be a subtle cyber-attack on the opponent's power grid, a targeted disinformation campaign on

social media, or the use of rogue groups and nations as proxies. These actions are low-cost, covert, and deniable, but can trigger cascading major failures in infrastructure, trust, or public opinion, achieving strategic effects without crossing a traditional threshold of armed conflict. These are commonly referred to as below the threshold operations. Non-linearity encompasses the dynamics dimension of the endeavor space (E-space), the conceptual model proposed by Alberts (2011) to analyze problem characteristics and identify an appropriate C2 approach.

This phenomenon affects the very nature of cause and effect in C2 planning. In a linear system, effects are additive and reversible. If Plan A fails, one can stop and try Plan B. In a non-linear complex adaptive system, actions are often irreversible and path dependent. A decision made in the decide phase changes the system state in a way that may remove entirely any possibilities for future options. This makes predictive, linear, and plan-centric C2 models obsolete. The focus must shift from predict and plan to sense, learn, and adapt.

The Cynefin framework (see Figure 1) offers an elegant model for analyzing cause-and-effect relationships, illustrating how these dynamics dictate specific decision-making approaches based on the applicable domain. The clear and complicated domains are primarily characterized by linear relationships, whereas the complex and chaotic quadrants inherently deal with non-linearities.

5.3 Simultaneity or the Coordination Challenge

Simultaneity refers to the co-existence and interdependence of actors, actions, and effects across all domains and echelons at the same time. The traditional phased model of operations, such as the United States' shape, deter, seize initiative, dominate, stabilize, and enable civil authority (Joint Chiefs of Staff 2022, 2020), is a linear simplification that no longer holds. Simultaneity encompasses the dependencies and coupling dimensions of Dr. Alberts' endeavor space (Alberts, 2011).

In a complex pan-domain environment, the system is

characterized by massive, concurrent parallelism. A tactical cyber effect, an operational-level information campaign, and a strategic-level space-based surveillance operation are not sequential; they are all happening at the same time, and their effects are intertwined and interdependent. The act segment of one OODA loop contributes to and influences the observe segment of another, in real-time. This high degree of simultaneity collapses decision timelines and demands a C2 architecture that can manage, through automation and appropriate computing and communication infrastructures, a vast, parallel and distributed processing problem, rather than a sequential one. This implies that for actions decoupled from human effectors or immediate cognitive judgment, modern C2 architectures must pivot away from synchronization by scheduling, the rigid sequencing and allocation of distinct time blocks for pre-planned actions. Instead, they must transition to synchronization by convergence, wherein autonomous intelligent agents dynamically negotiate, bargain, and coordinate in real time to generate immediate, synergistic effects.

5.4 Heterogeneity or the Integration Challenge

Heterogeneity results from the diversity of the artifacts and agents that co-exist and interact within the system. The modern C2 challenge is not to build a single, monolithic system but to orchestrate a coalition or federation of independent systems, or a system of systems (Maier 1998). As well-stated by Groth (1999), "...monolithic, functionally specialized organizations do not scale well. They may work quite perfectly up to a limit, but then gradually crumble under the sheer weight of the required coordination."

At their most fundamental level, all systems are collections of components. Systems, as wholes, can achieve what none of their individual components can alone. Systems of systems⁷ form a particular class of systems gaining increasing interest from the military procurement and operations communities. Several qualifiers have been used to characterize systems of systems, also referred to as polythitic systems or metasystems. These qualifiers include complex, adaptive, dynamic, and networked. The

main characteristic that distinguishes systems of systems from other classes of systems is the combination of independent governance and independent operations.

These systems and the organizations that use them are heterogeneous in every respect:

- **Cultures:** Different service branches, coalition partners, and inter-agency bodies operate with different command styles, risk tolerances, values, and worldviews.
- **Doctrines & Processes:** Each organization brings its own unique Standard Operating Procedures (SOPs) and Tactics, Techniques, and Procedures (TTPs).
- **Technologies:** The environment is a mix of state-of-the-art platforms and, most of the cases, legacy systems, with varying capabilities, data standards, communication protocols, and security levels.

This heterogeneity creates seams and friction at every interface. Each actor/system retains managerial and operational independence, meaning their collaboration must be negotiated, not just commanded. The design of a system-of-systems C2 architecture is all about smoothing sharp edges and managing friction to prevent the unavoidable seams from hindering the exploitation of diversity within the organization or the force.

A system-of-systems solution for modern C2 would encompass various autonomous, heterogeneous, distributed systems that must be able to communicate despite differences in language, design, context, interface, format, content and execution platforms. This reality makes managing heterogeneity an even more critical requirement for systems of systems than monolithic systems.

A typical embodiment of the system-of-systems concept is the NATO Federated Mission Networking (FMN) framework (Dechand et al., 2023). FMN is a comprehensive conceptual and operational architecture engineered to ensure

⁷ Note that system of system is used here in the broad sense that includes family of systems. No distinction is made here between the two definitions as the subtle nuance does not add any value to the discussion.

rapid, secure, and resilient information sharing across heterogeneous systems in multinational coalition operations. Developed from critical lessons learned within the Afghanistan Mission Network (AMN),⁸ FMN shifts the paradigm away from ad-hoc, post-deployment network integration. Instead, it mandates interoperability by design, an architectural philosophy enabling allied and partner forces to seamlessly plug into a standardized, federated environment on day zero of an operation. Crucially, it achieves this unified capability without requiring contributing nations to sacrifice national sovereignty or managerial independence over their organic systems.

5.5 Discussion

The combination of high dimensionality, pervasive non-linearity, intense simultaneity, and extreme heterogeneity confirms that the modern C2 problem is a complex adaptive system. This conclusion proves that any C2 system based on traditional, linear, and reductionist models and approaches, and seeking to achieve complete, centralized knowledge and deterministic, optimized control is inadequate and doomed to fail. The only viable path forward is to embrace new paradigms that tame complexity and turn it into a manageable variable, leading to our fifth thesis in this paper:

Thesis 5. Holonic architectures provide a natural and very effective framework for modelling and managing the challenges of modern C2 complexity.

6 Holonic Architectures for C2

We believe that the design of C2 systems must be based on the principles of distribution, agility, and adaptation to match the properties of the environment itself. This realization serves as the critical bridge to our proposed approach for addressing the challenges of complexity discussed above. Consequently, holonic architectures are explored throughout the remainder of this paper as a structural concept to guide the development of C2

systems and solutions.

6.1 Holarchies

A system composed of holons, introduced in Section 1.1, is called a holarchy, which differs from a hierarchy, the well-known traditional rigid structure inherent to the military culture. Unlike hierarchies, where command flows one way (downwards), holarchies are collections of nested holons, where the complexity of the constituent holons is encapsulated and abstracted. This property offers a dynamic structure where control and information flow both ways, with lateral connections being just as important as vertical ones. This structure is a fundamental architecture of biological and social systems, from cells to organs to organisms, or from individuals to teams to departments (Simon 1962).

6.2 C2 Holarchies

To make an analogy, one can observe that military organizations may be seen as composed of holons, which are wholes and parts. These organizations are not atomistic (only parts), nor are they holistic (only wholes). Rather, they are wholes that are part of larger wholes. An example is a frigate that is part of a Naval Task Group, which itself is part of a Joint Task Force. There are also parts that are relatively independent, while similar parts may join and form larger wholes, such as in the case of coalitions. Holarchies, flexible hierarchies of holons, are created through an optimized balance of:

- Top-down Process: Similar to how hierarchies are created in traditional military organizations.
- Bottom-up Process: Via the aggregation of holons to satisfy emergent requirements and needs for handling rapidly evolving situations.

This observation makes the holonic concept well suited to address the challenges of C2 in the modern operating environment by offering an elegant and effective way to balance autonomy and self-control with coordination that can be seen as quasi-hierarchical and quasi-directed. In a military context, a holonic C2 system allows lower-level units (holons) to operate autonomously within their

⁸ Where fragmented systems severely hindered real-time collaboration.

sphere of responsibility while remaining aligned with the higher-level intent of the larger force (holarchy). This construct, which exemplifies the principle of bounded autonomy, strongly aligns with the doctrine of mission command (NATO Standardization Office, 2025 para 3.7 to 3.12). By balancing individual agency with organizational intent, it provides a structural framework for decentralized execution.

Figure 6 shows a conceptual model of a holarchic C2 command structure that maps C2 echelons onto a holarchy. Each level operates as an autonomous holon, managing its internal complexity (i.e., its sub-holons) while acting as a cooperative part of its parent holon (i.e., its whole). This structure enables decentralized execution and adaptability by distributing high-speed OODA loops at the tactical edge, while maintaining unity of command and global synchronization and adjustment through the top-down propagation of commander's intent.




Figure 6: Holarchic C2 Command Structure

This C2 holarchy is a concrete manifestation of the fractal property discussed in Section 4.3. Each node in Figure 6 is simultaneously a whole, running its own OODA loop over its internal resources and local environment, and a part of a larger loop at the next higher echelon. The same pattern repeats itself from the smallest unit up to the highest headquarters.

Some of the key properties of holonically organized,

controlled, and coordinated systems are discussed in the following sections. A C2 system designed as a holarchy would be built upon these key features, which would provide it with the mechanisms to address the above identified pillars of complexity.

6.2.1 Functional decomposition

The functional decomposition feature of holarchies allows for the management of the complexity of modern C2 systems by breaking down a large-scale system into smaller, self-contained functional units, i.e., holons and nested holarchies. This divide-and-conquer approach allows each holon to focus on a specific, manageable set of responsibilities or capabilities. In a holarchy, this decomposition is recursive, as discussed above. For instance, a complex high-level function, such as air superiority, is decomposed into sub-functions, e.g., surveillance, interception, and refueling, each of which is autonomously managed by a subordinate holon. This ensures that control and coordination are achieved locally, within manageable scopes, rather than overwhelming a single central decision holon with the entire system's complexity (Christensen 1994; Simon 1962).

6.2.2 Abstraction

Abstraction is defined as many-to-one mapping (Van Brussel et al. 1998). It refers to the holarchy's ability to manage complexity by hiding the internal details of lower-level (e.g., tactical) holons from higher-level (e.g., strategic) ones. A higher-level holon views its constituent holons as simplified, opaque functional units that offer specific capabilities or services, without needing to view and comprehend their internal state and processes. Holons interact through defined interfaces and services.

This encapsulation reduces data and cognitive overload and allows the entire system to scale, as decision-makers (i.e., higher-level holons) only need to see and process the aggregated or abstracted status of their subordinates (i.e., lower-level holons) rather than the raw data of every individual element. A JTF HQ holon would manage a handful of (air, land, sea, cyber, space) component holons. It does not need to know about the internal sub-holons,

such as individual wings, brigades, etc., that the component holons are managing. This structure mirrors the fractal OODA loop model proposed in Section 4.3 and makes holonic C2 inherently scalable.

6.2.3 Autonomy

Autonomy is the capability of an agent to decide and act on its own, without constant intervention from external or higher-echelon agents.

In a C2 holarchy, this implies that each C2 holon, e.g., a JTF HQ, a component command, or even a tactical unit would possess sufficient local intelligence, delegated authority, and resources to manage its own internal affairs and execute its own C2 loops to solve local problems and respond to local emerging needs and disturbances without constant top-down intervention from higher echelons. This autonomy is not absolute but bounded, as it operates within the boundaries of defined mission constraints and goals set by the parent holon, e.g., the commander's intent. This property allows for a balance between distributed responsiveness, at the edge, and centralized strategic alignment.

It is essential to note that the concept of holon is agnostic regarding the type, scope, and degree of autonomy, as these factors are inherently dependent on the specific problem, mission, or context. Autonomy of holon is not a fixed attribute; rather, certain of its parameters are established during the design phase, while others are dynamically determined during the operational phase.

6.2.4 Ad hoc Cooperation

Ad hoc cooperation describes the dynamic ability of holons to spontaneously form temporary alliances or virtual clusters to solve unanticipated problems or respond to emergent needs. Unlike a rigid hierarchical cooperation, which follows pre-defined lines of authority, ad hoc cooperation property allows holons to identify peers with aligned goals and the required skills and capabilities, negotiate mutual plans, and execute joint actions on the fly (Christensen 1994; Ulieru et al. 2006). Coordination, negotiation, bargaining, and other cooperation techniques allow holons to flexibly interact in a

spontaneous way and harmonize efforts, not just in response to prescriptive top-down commands. This property does not, however, make holons rogue agents. Their autonomy is bounded by a requirement to cooperate with peers and serve the objectives of their parent holon.

6.2.5 Self-Organization

Self-organization is a fundamental characteristic of complex adaptive systems that is closely related to ad hoc cooperation and autonomy. Self-organization in a holarchy refers to the spontaneous emergence, through immediately local interaction among holons, of coherent global order without the need for external, top-down direction (Koestler 1967; Holland 1995).

Driven by the Janus effect (Koestler 1967), where each holon seeks to preserve its autonomy while integrating into the larger whole, the system continuously adjusts its own control structures and power balances to maintain the overall stability and cohesion. In a C2 context, this means that the command structure itself can evolve organically. For instance, specific task forces can assemble, define their own leadership protocols, and dissolve based on the immediate requirements of the battlespace, ensuring the organization composition, structure, and control mechanisms match the complexity of the problem it faces (Leitão 2009).

6.2.6 Adaptiveness

Adaptiveness refers to the agent's or system's ability to alter, in a timely and effective manner, its own structure in response to newly defined goals, internal failures, or significant changes in the context of operations. This encompasses many aspects of the previous properties.

Because the holarchy is a fluid and modular organization, composed of self-reliant holons rather than rigid functional blocks, it can absorb disturbances by re-tasking or replacing its individual holons without collapsing the entire system (Tharumarajah et al. 1996; Leitão 2009). This allows for soft and hard reconfiguration to occur seamlessly during operations, ensuring that the C2 architecture remains stable and effective,

even when the mission or the force structure evolves (Babić 2012). An illustrative example of soft reconfiguration would be the dynamically changing communication topology between the force's C2 nodes, while an example of hard reconfiguration would consist of integrating new sensors or new coalition partners. This property enables the holarchy to self-repair and to adapt to non-linear threats by continually reconfiguring its internal relationships to match the changing environment. This concept is referred to as coalition formation or dynamic teaming in multi-agent systems literature (Shoham 1994).

6.3 Discussion

Traditional systems force a false choice between centralized and decentralized control. The former is unified but slow and brittle, presenting a single point of failure; while the latter is fast and resilient but can lead to total chaos.

With the feature set discussed above, holonic control and coordination concepts can directly and elegantly avoid this binary choice and help resolve the central paradox of modern C2, that is, how does a force achieve unity of command and synchronized action when its C2 architecture must be decentralized, scalable, and modular to survive and operate at the speed of relevance?

Moreover, holonic C2 will preserve the hierarchical and recursive nature of military organizations. It offers a hybrid structure that takes advantage of the capacities of distributed control and coordination while incorporating the benefits of the hierarchical control that allows for strong goal orientation. Such an organizational structure has the potential to preserve the chain of command and provide sufficient autonomy and robustness, while increasing adaptability and agility:

1. **Chain of Command:** Is preserved through the propagation of the commander's intent down the holarchy. A superior holon, such as a JTF HQ, passes a mission-type order containing the 'what' and the 'why' to its subordinate holons, for instance the Air Component. This flexible chain of command serves as the mechanism for bounding holon autonomy. The requirement to maintain a structural hierarchy (albeit one that moves beyond legacy, rigid definitions)

is primarily driven by the need to ensure a mandatory level of predictability within the collective system. Furthermore, the high stakes inherent in military operations necessitate a clear chain of command to maintain essential lines of accountability.

2. **Autonomy and Robustness:** Autonomy and robustness are provided through the freedom given to the subordinate holon, acting as an autonomous whole, to manage its internal resources and execute its own full C2 loop (the "how") to accomplish its part of the mission. It, in turn, passes its own intent to its subordinates. The knowledge of intent does not have to be limited to one level up; the number of levels for communicating intent can be mission- and context-dependent. This flexibility accommodates current practices, such as within NATO, where subordinate commanders are aware of intent two levels up.
3. **Adaptability and Agility:** Are enabled by the directed emergence or bounded autonomy, where the force's actions are locally adaptive but globally coherent.

Based on the above discussion, the concept of holonic architecture seems uniquely suited to manage the four pillars of complexity identified in Section 5:

- **Dimensionality and Heterogeneity:** Holonic architectures can help manage scale and heterogeneity through abstraction and encapsulation. The whole, e.g., a JTF HQ holon, does not need to receive and process all the data, potentially petabytes coming from every constituent holon, such as sensors, aircraft, etc., but has only access to a fused and abstracted report, e.g., I can provide 10 surveillance sorties and 5 strike sorties in Area X, from a holon, just one level below it, for instance, the Air Component holon. Both dimensionality and the heterogeneity of specific aircraft, sensors, or coalition partners are encapsulated and managed within the Air Component holon and remain invisible to the level above it.
- **Non-linearity and Simultaneity:** Holonic autonomy is the key to managing unpredictable, parallel operating environments. When a

local, non-linear event occurs, such as a pop-up threat, the holon at the tactical edge can traverse its own OODA loop at the speed dictated by the emerging situation to act. This helps avoid any unnecessary delays that would result in a loss of opportunity, if it had to send a request up the entire command chain and wait for a top-down direction. Holarchies avoid bottlenecks that result from centralized control, allowing the entire force to handle thousands of simultaneous events in parallel, as each holon addresses its local challenges in alignment with the shared, distributed intent.

The implications of Ashby's Law, formalized in Equation 1, above, provide a principled justification for the adoption of holonic architectures in distributed C2 systems. Having

$$V_c < V_e$$

Indicates that disturbances are overwhelming the control structure and leading to a loss of stability. Holonic architectures address this situation by distributing the control variety across multiple echelons of holons, all while encapsulating local complexity at each level. As a result, each holon absorbs a portion of the environmental variety locally, reducing the effective variety exposed upward in the holarchy and allowing the system to satisfy Equation 1 through distributed and multi-level variety absorption rather than centralized control. This collective absorption of variety across the C2 holarchy ensures that the overall system meets the requisite variety condition, thereby maintaining stability and coherence in complex operational environments.

7. Implementation Considerations

Proposing a new theoretical framework for C2 is a necessary first step, but its value is contingent upon its feasibility and practical applicability to real-world problems and situations. Transitioning from the current hierarchical paradigm to true holonic, fractal system-of-systems architectures and solutions is a significant undertaking, presenting multiple cognitive, technological, and doctrinal challenges. A decade ago, the holonic C2 paradigm would have

remained highly theoretically sound but deemed technologically impossible. Today, the foundational technologies for enabling this shift are already, not only emerging, but effectively in use. Many similarly complex domains have started demonstrating the viability of these advanced approaches and concepts.

7.1 Enabling Technologies

A holonic C2 architecture can leverage the significant advances in, and the convergence of many enabling technologies and tools, making its implementation and deployment objectively viable. This is particularly the case of computational tools and paradigms from the field of (distributed) AI, such as Generative AI and Agentic AI.

Some of the most promising technologies are discussed in the following sections. Rather than treating these technologies as a flat list, the document structured them as an architectural stack (see Figure 7) that mirrors the dual nature of holonic systems: the internal autonomy of a single entity and its collaborative integration into a larger system.

The structural layout offers a clear narrative progression:

1. Layers 1 and 2 establish the physical connectivity, compute survivability, data virtualization, and security paradigm. Without this solid base, cognitive models cannot run reliably, notably at the tactical edge.
2. Layer 3 constructs the internal cognitive mechanics of a single autonomous holon (the individual wholes).
3. Layer 4 scales the capability defined in Layer 3 by introducing the communication protocols and multi-agent dynamics required to bind individual entities into a collective capability (the cooperative parts).
4. Layer 5 acts as an overarching technological multiplier, showing how the entire C2 system can scale exponentially in speed and optimization capability using quantum computing advancements.




Figure 7: Architectural Stack for Holonic C2

7.1.1 Layer 1: Foundational Infrastructure & Fabric

Layer 1 establishes the physical and logical substrate required to support a highly distributed, multi-domain military C2 architecture. The technologies within this layer integrate programmable Software-Defined Networking (SDN), a dynamic edge-to-cloud compute continuum, and comprehensive digital twin virtualization into a resilient, adaptive operational grid. By ensuring seamless data transport and localized processing power, this foundational layer provides the essential connectivity and environmental state representation required for autonomous holons to function, adapt, and self-organize within contested, degraded, or operationally limited environments.

Engineering components and networking paradigms that contribute to this resilient physical and logical framework are examined in detail below.

7.1.1.1 Software-Defined Networking

Software-Defined Networking is an architectural approach that fundamentally changes how networks are designed and managed. It separates

the network's control from hardware. In traditional networks, control logic is distributed within each router or switch. SDN centralizes this into a software-based controller, enabling administrators to initialize, control, change, and manage network behavior dynamically via open interfaces, rather than manually configuring individual devices (Open Networking Foundation 2012).

Holarchies need dynamic connectivity. While SDN dynamically wires the network paths, the data processing and compute resources traveling over those paths must be managed with equal fluidity. An SDN-based approach allows the network topology shaping the holarchy to reconfigure itself, and reconfigure routing policies and bandwidth allocation, in real-time, in order to support the evolving needs of the mission. Examples include a sudden need for high-bandwidth drone video, or the creation of a new sub-holarchy. SDN is therefore a key enabler of holarchies, ensuring that the nervous system of the C2 architecture remains agile and adaptive.

7.1.1.2 Elastic Computing

In a Denied, Disrupted, Intermittent, and Limited environment, computing must stretch dynamically from a centralized cloud down to tactical edge nodes. A holarchy cannot function if the brain (or AI) resides in a central location, e.g., cloud. Such a monolithic brain would create the very latency and vulnerability that the concept seeks to avoid, as discussed above. This is why a holonic model requires AI-powered decision-making, the underlying reasoning capabilities, and the supporting computing power to be everywhere, from cloud servers down to processing units on autonomous uncrewed aerial vehicles (UAVs):

1. Cloud: Allows strategic holons to leverage massive, centralized cloud computational and storage resources for long-term situation analysis and decision-making processing.
2. Edge: Allows tactical holons, such as drones or soldiers carrying onboard computer (i.e., Edge AI), to execute their high-speed OODA loops autonomously, even in a DDIL communication environment, sustaining holon autonomy during network disconnects.
3. Fog: An intermediate layer that allows local holarchies, or clusters of holons, to share their computing power and storage capacity dynamically. This may consist, for instance, in a swarm of drones pooling their processors to analyze captured imagery.

Elastic architectures within the cloud-fog-edge continuum are essential for orchestrating processing and resources dynamically across heterogeneous environments, ensuring the overall system remains robust and resilient under fluctuating operational conditions. Unlike static deployment models, these architectures adaptively balance the workload by exploiting the distinct capabilities of each layer, i.e., the scalable massive storage and computational capacity of the cloud, the low-latency intermediary processing of fog nodes, and the proximity and immediacy of edge devices (Al-Dhuraibi et al. 2017). This adaptive distribution is governed by real-time analysis of critical contextual factors,

including the severe computational and energy limitations at the edge, communication constraints such as bandwidth volatility and intermittency, and the privacy and security requirements of all the data and information collected, processed, and shared.

In the context of a holonic C2 system, this elasticity ensures that tactical holons, at the edge, can maintain their autonomy and high-speed OODA loops by processing data locally when disconnected or when privacy demands it, while dynamically offloading resource-intensive tasks to fog or cloud layers when conditions allow (Shi et al. 2016). By intelligently shifting tasks in response to these factors, elastic architectures can help optimize the trade-off between performance, cost, and reliability across the entire holarchy.

7.1.1.3 Digital Twins

High-fidelity and fast simulations are becoming increasingly common. Computational power has taken great leaps with significantly reduced costs so embedding complex simulations, as digital twins, into C2 systems and getting adequate predictive performance, is becoming a technically achievable endeavor. As a reminder, a digital twin (Willcox et al. 2023) is a high-fidelity digital replica or simulation of physical systems and entities. This has the potential to provide the rich semantic data foundation. Every physical asset, platform, and environmental factor is encapsulated as a digital twin, which functions as the persistent state representation that individual holons query, manipulate, and synchronize during localized decision-making.

As an example, a holon's orient phase, which gives it the ability to understand itself and its environment, could be implemented as a digital twin. The latter has the potential to act as a crystal ball that opens a window into the immediate future, and potentially beyond, for a blue holon, providing a significant advantage over a red holon. This is much like having access to the opponent's playbook in a football game. For instance, an air component C2 holon could maintain a digital twin of its forces, allowing it to model, predict, and wargame the second- and third-order effects of potential courses of action, as part of the decide segment of its OODA loop, before

committing resources in the real-world through the act phase.

How to best exploit digital twins, in real time, in real world operations is a lucrative area of research and development with potential major scientific and operational payoffs.

7.1.2 Layer 2: Decentralized Trust & Resilience

Layer 2 establishes the cryptographic framework and security paradigms necessary to guarantee operational integrity and data verifiability across a distributed, multi-domain network. Operating under the assumption of an actively contested cyber environment, the technologies within this layer replace vulnerable, centralized authority nodes with decentralized trust architectures, such as distributed ledgers and continuous zero trust network meshes. By embedding peer-to-peer consensus, immutable transaction logging, and secure cross-domain sharing directly into the communication fabric, Layer 2 ensures that autonomous holons can safely self-organize, exchange high-value operational intent, and execute mission commands with absolute assurance, even when isolated from higher headquarters.

Examples of cryptographic mechanisms and security architectures that safeguard this decentralized trust are analyzed in the following subsections.

7.1.2.1 Zero Trust Architecture & Cyber Security Mesh

Zero trust architecture is a security framework premised on the strict operational mandate of “never trust, always verify” requiring continuous authentication, authorization, and cryptographic validation of every asset, user, and data transaction, regardless of its location within or outside a network. Complementing this approach, a cyber security mesh deconstructs traditional, centralized network perimeters by establishing a distributed, modular security posture that places a distinct, personalized security envelope around every individual node, data point, or access interface.

For a holonic C2 system, the convergence of these two paradigms is not a simple security enhancement but a fundamental operational necessity. Because holarchies are fluid, ephemeral, and frequently deployed in DDIL environments, legacy perimeter defenses fail when centralized authentication nodes become unreachable. By implementing a cyber security mesh, security policy enforcement is decoupled from central infrastructure and embedded directly into each independent holon, mirroring the holon’s inherent autonomy. As these heterogeneous holons dynamically self-organize into multi-domain task groups, zero trust architecture enables them to establish peer-to-peer trust matrices and execute localized cryptographic consensus mechanisms on the fly. Furthermore, by enforcing strict micro-segmentation at the node level, this paradigm drastically limits the blast radius of a cyber breach. For instance, if an individual edge holon, such as a tactical UAV or field sensor, is physically captured or digitally compromised, the rest of the holarchy can instantly isolate that asset, preserving the operational integrity of the broader distributed network.

7.1.2.2 Blockchain and Distributed Ledger Technology

Blockchain offers a distributed, tamper-resistant, and shared ledger to keep record of digital transactions that are cryptographically signed, represented as a chain of blocks replicated across a decentralized network of nodes. This technology was developed as a response to the erosion of trust in the financial system after the 2008 crisis. Blockchain has the potential to help increase confidence in socio-technical systems, and the data and information they produce and exchange, since it does not require any central trusted authority. Confidence in blockchain is induced by its underlying governance structure.

Collaboration, and the underlying coordination processes, requires trust, accountability, and a shared source of truth. In a holonic C2 system, blockchain technology can help solve the coordination with limited trust problem. It allows a federation of holons, more importantly if they are from different coalition partners, to maintain

a shared, immutable record of orders, rules of engagement, and resource allocations without relying on a single central server. This ensures that even if the top-level holon, such as the combined force headquarter, is disconnected, local lower-level holons can verify the authenticity of an old order or record a new engagement action that will eventually synchronize with the rest of the holarchy, when connectivity is reestablished, ensuring data integrity and preventing fog of war discrepancies.

7.1.3 Layer 3: Individual Holon Cognition

Layer 3 defines the internal logic, decision-making mechanics, and intelligence processing capabilities of a single autonomous entity or holon within the C2 system. The technologies within this layer combine probabilistic generative architectures, goal-oriented agentic execution loops, and deterministic neuro-symbolic reasoning to transform passive data streams into proactive operational intent and verifiable action. By fusing the rapid pattern-recognition proficiencies of advanced AI models with the strict, rule-based logic required to enforce military doctrine and Rules of Engagement (ROE), Layer 3 provides individual holons with the localized cognitive autonomy needed to decompose complex objectives, dynamically utilize local tools, and reason reliably without requiring continuous human intervention.

Some of the technological frameworks and cognitive execution loops that drive this localized holonic intelligence are discussed below.

7.1.3.1 AI & Language Models

The practical implementation of holonic architectures is becoming increasingly feasible through the advancements in modern AI and language models, both Large Language Models (LLMs) (Vaswani et al., 2017, Brown et al. 2020) and Small Language Models (SLMs) (Javaheripi & Bubeck 2023), which would represent the brain of the holons. For instance:

- Machine Learning: For the observe phase of the OODA loop, advanced machine learning algorithms can be used for pattern recognition, for instance, as part of an automatic threat evaluation process.

- LLMs: Would act as a general-purpose cognitive engine for high-level holons, e.g., JTF HQ, providing the reasoning capabilities required for the complex sense-making, i.e., orient and strategic planning, i.e., decide segments. LLMs can also enable semantic interoperability by allowing holons, at all echelons, to communicate via natural language interfaces, effectively translating commander's intent into executable sub-tasks across the holarchy. Another objective would be to avoid misinterpretation between coalition partners not speaking the same language. LLMs can ingest unstructured text (reports, chat, doctrine) and synthesize it into structured understanding. They can act as the universal translator between disparate semantic domains.
- SLMs: Are more adapted for the tactical edge. These lightweight, efficient models can run locally on resource-constrained hardware, such as drones or handheld devices, endowing lower-level holons with sufficient intelligence and autonomy to execute tasks and adapt to local conditions even in DDIL environments.

Together, these classes of models realize the fractal nature of the holarchy by offering a robust, cloud-enabled reasoning at the core, and fast, resilient intelligence at the edge.

7.1.3.2 Generative AI

Generative AI refers to a class of AI models designed to create new content, including text, code, imagery, and simulations, rather than simply analyzing or classifying existing data. These systems learn patterns, structures, and probability distributions from vast datasets to synthesize complex information.

In a holonic C2 context, this capability would be at the core of holons brains. For instance:

- An ISR holon could use generative AI to synthesize diverse intelligence streams into coherent reports.
- An advanced command post can use this technology to generate multiple distinct Courses of Action for commander review.

Generative AI allows implementing autonomy at all echelons of the C2 holarchy.

7.1.3.3 Agentic AI

Agentic AI is defined as a class of AI systems capable of autonomous, goal-directed behavior. Unlike passive AI, which responds only to prompts, agentic systems possess the ability to

1. Perceive their environment.
2. Reason through complex problems.
3. Decompose these problems into simpler tasks.
4. Actively use tools and interfaces to accomplish the tasks and achieve the high-level and more complex objectives.

These agents operate with a continuous feedback loop, planning and executing actions, and correcting their plans without requiring much human intervention; or at least much less than does the non-agentic AI.

These properties make agentic AI a natural fit for holonic C2 architectures, allowing for the practical realization of the recursive autonomy central to the holonic paradigm. High-level (agentic) holons can autonomously decompose complex strategic objectives into a set of simpler sub-objectives or tasks, negotiate with peer holons how to allocate and use available resources to create and synchronize effects, effectively functioning as a holonic C2 system.

7.1.3.4 Neuro-Symbolic AI

Neuro-symbolic AI is a hybrid approach that integrates the learning, data-driven capabilities of neural networks with the reasoning and explanation power of symbolic, knowledge-based AI approaches. Neuro-symbolic systems combine the best of the two worlds, using neural components to perceive and structure data, which is then processed by symbolic engines to perform logical deduction and planning. This synergy creates systems that are both robust in learning and capable of transparent,

explainable, and verifiable/auditable reasoning.

In the context of C2 holarchies, neuro-symbolic AI provides the mechanism for the creation of calibrated trust. The commander can understand the rationale behind an autonomous (low-level) holon's decision thanks to the intelligible symbolic approach used by the reasoning layer. At the same time, the system retains the ability to adapt to raw, messy battlefield data through its neural layer.

7.1.4 Layer 4: Collective Intelligence

Layer 4 governs the macro-architectural framework, communication protocols, and organizational dynamics that bind isolated autonomous entities into a cohesive, cooperative holarchy. Moving beyond individual holon capabilities, this layer leverages multi-agent system dynamics and standardized interaction protocols to enable dynamic self-organization, emergent task allocation, and real-time collaborative problem-solving across disparate operational domains and organizational echelons. By establishing a unified mechanism for exchanging intent, negotiating roles, and managing localized context, the technologies within Layer 4 transform separate, heterogeneous units into a synchronized collective network capable of executing complex, pan-domain operations without relying on vulnerable, centralized command hierarchies.

The multi-agent organizational structures and standardized communication protocols enabling this systemic, holonic synchronization are detailed below.

7.1.4.1 Multi-Agent Systems

Multi-agent systems are composed of multiple intelligent agents interacting within a shared environment (Wooldridge 2009; Shoham 1994). Each agent is a discrete software entity capable of autonomous action to achieve its objectives. Agents are operating with only partial knowledge of the overall environment and system. In a multi-agent system, the agents interact through communication, cooperation, coordination, and negotiation to solve complex problems that, most of the time, exceed

the capabilities of any single agent.

From the above definition, a holarchy is fundamentally a multi-agent system. Each holon is an autonomous software agent with its own goals, state, and decision-making logic. Agent-based modeling provides the essential toolkit for simulating holonic C2 architectures. This allows to implement and test how different holonic structures, coordination and communication protocols, as well as autonomy levels respond to various situations, helping to de-risk development and validate doctrinal concepts.

Modern frameworks (e.g., LangGraph, AutoGen) allow developers to build stateful, multi-agent systems with complex execution topologies. For instance, LangGraph (LangChain AI., 2024) excels at defining cyclical graphs and nested agent relationships, mimicking a holarchy.

7.1.4.2 Agent Interaction & Interoperability

The transition from Generative AI to Agentic AI marks a pivotal shift in the design and functionality of software systems. Agents are no longer passive text generators but active entities capable of planning, reasoning, and executing tasks with minimal human intervention. However, a significant barrier remains: the $N \times M$ integration problem, where N agents must connect to M tools and other agents, leading to a combinatorial explosion of point-to-point configurations. To address this challenge, the industry has moved toward standardization, developing distinct protocols to serve as the connective tissue of the agentic ecosystem. These include frameworks such as the Model Context Protocol (MCP) (Anthropic, 2024), the Agent-to-Agent (A2A) protocol (Google, 2025), the Secure Low-latency Interactive Messaging (SLIM) framework (Cisco, 2025), the Agent Communication Protocol (ACP) (IBM, 2024), and the classic Agent Communication Language (ACL) (FIPA, 2002). Crucially, these protocols are not mutually exclusive alternatives; they serve fundamentally different structural purposes. Achieving complete operational interoperability within a multi-agent system requires an architecture that layers these protocols cooperatively.

In a holonic C2 architecture, this multi-tiered

protocol fabric allows heterogeneous holons, potentially built on entirely different software frameworks, to seamlessly understand each other's capabilities, negotiate tasks, and collaborate on distributed problems. For instance, A2A or ACP acts as the horizontal coordination tier, enabling a high-level strategic holon to discover, authenticate, and delegate operational tasks to specialized edge holons without knowing their internal software specifications. Once tasked, an individual edge holon utilizes MCP as its local vertical adapter, safely pulling data or triggering local sensors within its isolated edge environment. This entire distributed interaction is wrapped in SLIM, which provides the secure, end-to-end encrypted, low-latency transport layer necessary to pass messages safely across contested or degraded communication environments. Fusing these standards eliminates legacy data silos and ensures that the shared world model used for high-stakes command decisions remains grounded in real-time, verifiable operational reality.

7.1.5 Layer 5: Strategic Frontiers

Layer 5 serves as the ultimate computational accelerator for hyper-complex, multi-domain military operations. This layer introduces advanced, next-generation paradigms, such as Quantum AI, designed to process massive, high-dimensional datasets simultaneously and solve complex combinatorial problems that far exceed the limits of classical computing. Layer 5 has the potential to drastically compress the command decision cycle, granting the collective operational network unprecedented cognitive overmatch and execution speed in highly contested future environments.

An example of the emergent technologies driving this computational velocity is introduced in the following subsection.

7.1.5.1 Quantum Artificial Intelligence

Quantum computing represents a paradigm shift from classical algorithmic constraints by leveraging quantum mechanical phenomena, specifically superposition, entanglement, and interference, to process the hyper-complex data structures encountered in multi-domain battlespaces at unprecedented speeds. Quantum

AI (QAI) represents the convergence of quantum computing and machine learning, aiming to achieve a computational advantage over tasks that are fundamentally intractable for classical binary architectures (Biamonte et al., 2017). QAI utilizes quantum bits (qubits) existing in a simultaneous superposition of states to map a high-dimensional mathematical space, exploiting quantum interference to isolate optimal variables far more efficiently than sequential classical processing. To bridge the gap between theoretical quantum advantages and current hardware limits, variational quantum circuits have emerged as a framework for near-term deployment (Cerezo et al., 2021). These function as hybrid quantum-classical architectures where a quantum processor handles high-dimensional state space mapping while a classical optimizer iteratively updates the circuit's parameters.

In the context of holonic C2, QAI acts as a critical force multiplier within the orient and decide phases of the OODA loop:

- **Orient:** QAI can identify subtle, non-local correlations across massive, heterogeneous datasets that remain invisible to classical algorithms, enhancing the system's capacity to anticipate emergent threat vectors.
- **Decide:** QAI can efficiently navigate NP-hard combinatorial optimization problems, such as dynamic resource allocation and multi-domain effector-to-target matchmaking, in real time (Biercuk and Fontaine, 2017). For example, calculating the mathematically optimal weapon-target assignment of N defense missiles against M incoming threats can be resolved in milliseconds, bypassing the combinatorial explosion that would stall classical supercomputers for minutes or hours.

This operational velocity provides the processing power required for a holonic system to match the complexity of a multi-multi environment, turning data saturation into decision advantage.

7.2 Pathways to Implementation

Transitioning from current rigid, hierarchical C2 structures to fluid, holonic architectures is not a simple plug-and-play upgrade. This transition represents a fundamental paradigm shift. It calls for an evolutionary approach to de-risk and control the move from today's interoperable system-of-systems to a truly holonic one, with the required autonomy, cooperation, and recursion properties, through a series of phased milestones, as discussed in the following sections.

7.2.1 Structural Interoperability & Abstraction

In this foundational step, the disconnected capabilities are transformed into holon-ready entities. This requires moving beyond simple connectivity to true semantic interoperability. Implementation focuses on defining standardized interfaces and ontologies⁹ that allow diverse units, be they legacy frigates or next-gen cyber teams, to expose their capabilities as abstract services, e.g., abstracted status, capabilities.

In this phase, a brigade holon, for instance, defines what it can do (its capabilities) and what it needs (its requirements) without exposing the complexity of how it achieves them. This specifically addresses the heterogeneity challenge by creating a common transactional wrapper around diverse legacy and modern systems. The success of this approach clearly hinges on the willingness of participants to implement and adhere to common standards. This is one of the most difficult requirements to fulfill, particularly within multi-national coalitions.

7.2.2 Bounded Autonomy

Once holons, viewed as abstracted entities, can communicate semantically, they must learn how to reason, decide, and act autonomously. This phase introduces bounded autonomy to specific functional areas. For instance, high-frequency, high-complexity functions such as ISR C2, logistics resupply, or airspace deconfliction are encapsulated into support holons. These entities are powered by Agentic AI and SLMs, allowing them to execute OODA loops for specific tasks without frequent

⁹ Leveraging standards like the Model Context Protocol.

direct intervention from higher echelon holons. For instance, a logistics holon would autonomously negotiate route adjustments based on real-time threats and only escalate to higher echelon holons when finding a solution breaks its pre-defined bounds, as defined by the Commander's Intent.

7.2.3 Cooperation and Self-Synchronization

The third phase shifts the focus from vertical relationships to horizontal collaboration, directly addressing the simultaneity challenge. By deploying agent interaction protocols and distributed ledger technology, holons begin to form temporary, ad-hoc coalitions to solve immediate problems. For instance, a kinetic strike holon and a cyber effect holon could autonomously coordinate time-on-target synchronization via peer-to-peer negotiation, bypassing the bottleneck of a centralized Air Operations Center (AOC). This creates a system capable of synchronization by emergence and convergence rather than by directed scheduling.

7.2.4 Full Holarchy

The final state is the realization of a full fractal C2 system. Here, the entire force functions as a cohesive complex adaptive system. Top-level holons, e.g., human commanders, operate primarily at the intent level, using generative AI interfaces to articulate the mission objectives and boundaries. The holarchy then recursively decomposes this intent and provides assessments and options to achieve it. The system naturally adapts and reconfigures itself (self-organization) to match the changing environment, achieving the goal: a force that is adaptive and resilient to overcome the complexity it faces.

7.3 Critical Challenges and Limitations

Adopting a holonic paradigm for C2 cannot be reduced to a simple technological upgrade. This significant evolution requires overcoming many barriers that are as much human and doctrinal as they are technical. The transition from centralized hierarchies to distributed, autonomous holarchies introduces novel risks that must be proactively managed to prevent infinitesimal events at the edge from becoming triggers of a total collapse from which the entire holarchy cannot recover. This is the well-known butterfly effect.

7.3.1 Technical Challenges

Some of the most important technical challenges include:

- **Verification & Validation:** The primary technical barrier is the verification and validation (V&V) of non-deterministic systems. Traditional V&V relies on exhaustive testing of all possible states to ensure safety. However, a holonic system powered by agentic AI and generative AI operates in an effectively infinite state space, making exhaustive testing impossible. For instance, how to certify an Integrated Air & Missile Defence (IAMD) holon or a cyber offense holon when their own behavior is emergent and path-dependent? The main challenge resides here in how to move from certifying the code to certifying the boundaries, or sandboxing. The objective is to achieve runtime assurance by checking the behavior of individual holons and the whole holarchy against a set of rigorous rules of holonic engagement, instantly overriding any action that does not abide by these rules.
- **Resilience in Disconnected Environments:** While holons are defined as cooperative entities, true cooperation, and its enabling coordination processes, require information exchange. In a DDIL communication environment, the nervous system of the holarchy is vulnerable. A fractured holarchy faces the split-brain problem found in distributed databases. This problem refers to situations where, because of a network failure, a distributed system is partitioned into clusters, causing different parts to act independently, all assuming they are the sole leaders. If a brigade holon loses contact with its division holon, and both make contradictory decisions based on divergent local data, how is consistency restored upon reconnection?
- **The Cognitive Gap:** Integrating legacy platforms, with limited capabilities, into a smart, technology-heavy holarchy remains a persistent friction. An old tank, using technologies from the 1970s, will never be natively able to act as a smart autonomous holon.

7.3.2 Human and Doctrinal Challenges

From the human and doctrinal perspectives, the transition to a holonic C2 framework is fraught with challenges; some of the most significant include:

- **Trust:** The single greatest barrier to adoption of a holonic C2 formalism would be trust. Military operations demand predictability, yet complex adaptive systems, such as holarchies, are inherently unpredictable. Can a commander trust an air component holon to act autonomously in accordance with their intent? Can that holon trust the abstracted data from its peers? Trust presents a dual risk: one of over-trust, that is blindly relying on the AI and automation technologies, and one of under-trust or dis-trust, which refers to natural aversion to the risk of AI and autonomous systems. Hitting the sweet spot between these two extremes is one of the biggest challenges of AI-powered holonic C2. Tackling this trust calibration problem will require a lot of transparency, traceability, and explainability.¹⁰ Without such transparency, human holons (commanders) will never agree to relinquish control to software, AI-powered holons.
- **Leadership Philosophy:** The implementation of holonic C2 represents a paradigm shift away from the rigid hierarchies of traditional military doctrine. While legacy frameworks are optimized for vertical predictability and formal accountability, a holonic architecture necessitates a degree of distributed autonomy that fundamentally disrupts these foundational pillars. Implementing this concept demands a doctrinal shift to true mission command or command by influence, where leaders define intent (the what and why) and abstain from dictating the how. For a generation of leaders raised on the illusion of total control, this loss of granular control will feel unnatural, counterintuitive, and counterproductive.
- **Accountability:** Last, but not least, accountability might represent a challenge among the hardest to overcome. If an autonomous drone holon, acting within its valid boundaries (as defined by

its rules of holon engagement), commits a war crime while reacting to emergent complexity, who would be responsible? The software engineer? The commander who set the intent? The system itself? Answering this question will require new legal frameworks that define liability for autonomous holons and ensure that meaningful human control, especially for actions involving kinetic effects and potential loss of human lives.

7.4 Case Studies and Analogous Domains

While the implementation of a full-scale, autonomous holonic C2 architecture in a military context is novel, the underlying principles are neither experimental nor untested. The structural patterns of holarchy, that is bounded autonomy, cooperative negotiation, and recursive resilience, have been validated in other complex, distributed domains where the cost of failure is unacceptable and centralized control is infeasible. Examining these analogue areas can help derive critical lessons for the military domain.

7.4.1 Industrial Manufacturing

This application domain represents the origin of the holon concept. The concept of holonic manufacturing systems emerged in the early 1990s from the Intelligent Manufacturing Systems consortium as a response to the rigidity of mass production lines (Van Brussel et al. 1998). In traditional manufacturing, a central computer dictated every movement of every robot. If one machine failed or a rush order arrived, the entire schedule had to be re-optimized globally, causing massive downtime. Holonic Manufacturing Systems, where machines, workpieces, and transport units act as autonomous holons, provided a framework to decentralize this process.

With such an architecture, a task holon, which defines what needs to be built, enters the holarchy and negotiates directly with resource holons, which correspond to machines that can build what is required. This is the direct equivalent of dynamic resource allocation in military operations. A strike mission holon in a C2 holarchy must autonomously

¹⁰ For instance, "I am recommending Course of Action B because the adversary's logistics pattern suggests a feint at Sector A."

find new shooter holons if its primary assets are neutralized or unavailable, ensuring mission completion despite constraints and contingencies.

A prominent real-world example of a holonic approach in manufacturing is the Siemens Electronics Works Amberg (EWA) plant in Germany (Gilchrist, 2016). The EWA is a benchmark for smart factories and serves as a prime application of holonic manufacturing systems. In this environment, the production system is not a rigid, top-down hierarchy, but rather a collection of autonomous, cooperative building blocks, or holons.

7.4.2 Smart Grids and Microgrids

Smart grids provide an excellent practical example of a holonic system, offering a compelling blueprint for modern C2 architectures. Traditional power grids are highly centralized. A single node, the primary power plant, pushes energy downstream to passive consumers, directly mirroring legacy, top-down C2 structures. This model suffers from a critical vulnerability: a single point of failure. Because the architecture is rigid, a minor disruption at the central hub can easily cascade into a widespread regional blackout.

In contrast, modern smart grids are evolving into holarchies, nested networks made up of semi-autonomous microgrids. This architecture balances collective efficiency with local independence:

- **Normal Operations:** These microgrids cooperate dynamically with the macro-grid to optimize global energy flow.
- **Crisis Operations:** During a disruption, such as a cyberattack or a natural disaster, a microgrid can instantly island (isolate) itself from the main network. By pivoting immediately to autonomous operations powered by local generation (solar, wind, battery storage, or traditional generators), the microgrid ensures that local critical assets remain online.

This concept of islanding is identical to the military requirement for operations in DDIL environments. A brigade holon must function as an integrated part of its division holon under normal conditions (i.e.,

connected mode) but must be capable of seamless islanding (i.e., autonomous mode) if connectivity is being denied or compromised, continuing to fight based on the last received Commander's Intent before synchronizing when connectivity is reestablished. This self-contained resilience is precisely how modern, decentralized C2 functions to maintain mission assurance when the broader network is compromised.

7.4.3 Mosaic Warfare

Within the defense sector, the most significant step toward holonic C2 can be found in DARPA's Mosaic Warfare concept (Deptula et al. 2019; Rokvić 2025). Traditional force design relies on monolithic, multi-function platforms, such as an F-35 that assumes all the roles, from sensing to shooting, and everything in-between. These are powerful but brittle capabilities, in that, the loss of a few platforms can significantly degrade the entire force.

Mosaic warfare consists in decomposing these complex, composite capabilities into a set of smaller, single-function, low-cost systems, or tiles of the mosaic, that can be recomposed on the fly based on the mission requirements. For instance, a sensor from a satellite is temporarily coupled with a jammer on a drone, and a missile launcher from a frigate or a destroyer to dynamically form an effective kill web. This definition aligns mosaic warfare closely with holonic architectures. Both paradigms operate on the principle that the collective kill web constitutes an emergent capability greater than the sum of its individual tiles. However, while mosaic represents the closest contemporary approximation of a holonic framework, the two concepts remain distinct and should not be considered identical.

The complexity in mosaic warfare of managing thousands of disaggregated tiles obviously exceeds by far human cognitive capacity. This is where holonic C2 goes a step further by providing the holonic control layer where the tiles (holons) can autonomously negotiate their roles to form a cohesive whole in real-time.

7.4.4 Holonic Multi-Agent Systems

Academic research has provided a theoretical backbone for these applications. Gaud et al. (Gaud et al. 2007) formally defined the Holonic Multi-Agent Systems (HMAS) framework, proving that holarchies provide the optimal balance between the stability of hierarchy and the flexibility of heterarchy (flat networks). Similarly, we demonstrated (Benaskeur and Irandoust 2005) the application of these principles to distributed sensor networks, showing that holonic coordination reduced bandwidth usage and improved tracking accuracy compared to centralized surveillance approaches. These studies confirm that the holonic architectural paradigm is not simply a philosophical construct but a viable engineering solution for the complexity of the multi-multi battlespace.

autonomy is bounded by a shared, systemic objective. Holons do retain local decision-making authority to handle local disturbances, but they are hardwired for cooperation.

7.4.5 Systems of Systems

While the concepts of systems of systems (Maier, 1998) and holonic systems might initially appear close due to their shared focus on engineering complex, large-scale architectures from independent components, they are fundamentally distinct in their operational philosophy and primary objectives.

System of systems are fundamentally designed to address interoperability challenges among disparate, legacy, or heterogeneous constituent systems, essentially working at the seams to enable integration and shared information across systems. The foundational premise of system of systems is the preservation of managerial and operational independence. The constituent systems (e.g., individual satellites, radars, and command centers in a defense framework) were often built at various times, by different vendors, for different primary missions. They choose to cooperate to achieve an emergent capability, but they can decouple and revert to their isolated functions at any time. These systems of systems operate primarily on a horizontal and peer-to-peer axis.

In contrast, holonic systems are explicitly concerned with the internal dynamics of collaboration, coordination, and orchestration. They rely on holarchies to achieve collective goals, optimize resource allocation, and dynamically reconfigure in response to disruptions. In a holonic framework,

8. Conclusion

In this paper, we have argued that the fundamental problem facing modern C2 is not technological, but deeply theoretical. We have demonstrated that the accelerating complexity of the modern, multi-dimensional operating environment has rendered traditional C2 models not just suboptimal, but structurally obsolete. The linear, hierarchical, and deterministic frameworks that have served military organizations for centuries are fundamentally incapable of coping with the non-linearity, simultaneity, and hyper-dimensionality of today's battlespace.

While contemporary C2 concepts correctly identify the need for integration, they overlook a critical systemic vulnerability by pursuing the flawed ideal of total hyper-connectivity. There is a persistent, techno-centric belief that maximizing connectivity, by linking every sensor to every effector, indiscriminately improves control. However, viewed through the lens of complexity science, fully interconnected systems are inherently brittle and lack the structural modularity required to absorb operational shocks. As nodes become more tightly coupled, the architecture suffers from the fragility of interdependence, wherein localized errors propagate instantaneously across the network, triggering catastrophic cascading failures. This realization dictates a fundamental requirement for modern C2 architectures: a deliberate shift from maximizing connectivity to optimizing it, intentionally privileging localized responsiveness and loose coupling over centralized global control and tight integration.

To provide an architectural foundation for this shift, this paper introduced a novel theoretical framework that redefines the C2 problem itself. By modeling C2 not as a singular, linear workflow, but as a fractal system of recursively nested and concurrently intertwined OODA loops, we have shown that modern C2 inherently functions as a complex adaptive system. We subsequently proposed holonic architecture as the most theoretically sound and practically viable solution to navigate these complexity challenges. By defining C2 nodes as holons, the holarchic model elegantly resolves the central paradox of modern command: achieving comprehensive unity of effort without relying

on rigid, centralized control. Holonic C2 offers a structural blueprint for directed emergence, yielding an operational system that is:

1. Fractal: Exhibiting self-similarity and structural consistency across all strategic, operational, and tactical scales.
2. Distributed: Capable of bounded autonomy, where local holons execute high-speed OODA loops to seize fleeting opportunities within the commander's intent without waiting for centralized approval.
3. Adaptive: Dynamically capable of self-organization, self-reconfiguration, and self-healing in direct response to evolving mission demands, physical attrition, or compromised communications.

Moving beyond abstract theory, this paper outlined the engineering feasibility of this paradigm. We argued that the convergence of several technologies brings this theoretical model within operational reach. We are witnessing the maturation of the exact computational substrate required to sustain a holarchy: intelligent AI agents capable of autonomous localized reasoning, distributed ledgers for decentralized synchronization without single points of failure, and software-defined networks capable of reshaping architectural topology in real time. However, realizing this vision demands a profound cultural transformation alongside these technical advancements. The most formidable barriers to holonic C2 remain political, doctrinal, and human in general.

In conclusion, the future of decision advantage does not lie in building a massive, monolithic centralized brain, but rather in optimizing and coordinating a resilient, distributed nervous system composed of specialized cognitive nodes of varying scales. Embracing the holonic paradigm allows us to build C2 architectures that mirror the very complexity they are engineered to master. The transition from rigid command and control to agile command and coordination via autonomous holarchies represents the necessary evolutionary leap to prevent systemic collapse when a resilient command architecture is needed most.

References

- Alberts, D. 2011. *The Agility Advantage: A Survival Guide for Complex Enterprises and Endeavors*. Department of Defense Command and Control Research Program.
- Alberts, D., & Hayes, R. 2006. *Understanding Command and Control. The Future of Command and Control*. CCRP Publications.
- Alberts, D., & Hayes, R. 2003. *Power to the Edge: Command and Control in the Information Age*. CCRP Publication Series.
- Alberts, D., Huber, R. & Moffat, J., 2010. *NATO Net Enabled Capability Command and Control Maturity Model*. In Technical report, Office of the Assistant Secretary of Defense. Command and Control Research Program (CCRP).
- Al-Dhuraibi, Y., Paraiso, F., Djarallah, N., & Merle, P. 2017. Elasticity in Cloud Computing: State of the Art and Research Challenges. *IEEE Transactions on Services Computing* 11 (2): 430–47.
- Allied Command Transformation. 2021. *NATO Warfighting Capstone Concept (NWCC)*. North Atlantic Treaty Organization.
- Anthropic. 2024. *Model Context Protocol*.
- Ashby, W. 1956. *An Introduction to Cybernetics*. Chapman & Hall.
- Babić, B. 2012. Axiomatic Design of Flexible Manufacturing Systems. *Journal of Engineering Design* 23 (3): 237–59.
- Benaskeur, A., & Irandoust, H. 2005. *Holonic Approach for Control and Coordination of Distributed Sensors*. TR 2008-015. Defence R&D Canada – Valcartier.
- Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. 2017. Quantum Machine Learning. *Nature* 549 (7671): 195–202.
- Biercuk, M., & Fontaine, R. 2017. *The Leap into Quantum Technology: A Primer for National Security Professionals*. War on the Rocks.
- Boyd, J. 1987. *A Discourse on Winning and Losing*. Air University Press.
- Brown, Tom et al. *Language Models are Few-Shot Learners*. 2020. *Language models are few-shot learners*. *Advances in neural information processing systems*, 33, 1877-1901.
- Butchart, R., Schaus, J., & Hunter, A. 2020. *Pathways to JADC2: Expert Perspectives on the Future of Command and Control*. Center for Strategic; International Studies.
- Cerezo, M., Arrasmith, A., Babbush, R., et al., 2021. Variational quantum algorithms. *Nature Reviews Physics*, 3, 625–644.
- Christensen, J. 1994. *Holonic Manufacturing Systems: Initial Architecture and Standards Directions*. *Proceedings of the First European Conference on Holonic Manufacturing Systems (Hanover, Germany) 1: 1–20*.
- Cisco. 2025. *Slim: Secure Low-Latency Interactive Messaging (SLIM) for the Internet of Agents*.
- Dechand, M., Galvin, K., Jense, H., Pullen, J. M., & van den Berg, T. 2023. *Advancing Modelling and Simulation in NATO Federated Mission Networking*. NATO Science and Technology Organization (STO) Meeting Proceedings (STO-MP-MSG-207). doi:10.14339/MP-MSG-207-07-PDF
- Department of Defense. 2022. *Summary of the Joint All-Domain Command and Control (JADC2) Strategy*. Department of Defense.
- Department of National Defence. 2023. *Pan-Domain Force Employment Concept: Prevailing in a Dangerous World*. Concept Paper. Canadian Armed Forces.
- Department of National Defence. 2024. *Pan-Domain Command and Control (PDC2) Concept Paper: Decision Advantage in a Dangerous World*. Canadian Armed Forces.
- Deptula, D., Penney, H., Stutzriem, L., & Gunzinger, M. 2019. *Restoring America's Military Competitiveness: Mosaic Warfare*. Mitchell Institute for Airpower Studies.

- FIPA. 2002. FIPA Agent Communication Language Specifications. Foundation for Intelligent Physical Agents.
- Gaud, N., Hilaire, V., & Koukam, A. 2007. Holonic Multi-Agent Systems. In *Self-Organizing Software*. Springer.
- Gell-Mann, M. 1995. *The Quark and the Jaguar: Adventures in the Simple and the Complex*. Macmillan.
- Gilchrist, A. 2016. *Industry 4.0: The Industrial Internet of Things*. Apress.
- Google. 2025. Agent2agent.info
- Grisogono, A. 2006. The Implications of Complex Adaptive Systems Theory for C2. Proceedings of the 11th International Command and Control Research and Technology Symposium.
- Groth, L. 1999. *Future Organizational Design*. John Wiley & Sons, Inc.
- Hoehn, J. 2022. Joint All-Domain Command and Control: Background and Issues for Congress. CRS Report No. R46725. Congressional Research Service.
- Holland, J. 1995. *Hidden Order: How Adaptation Builds Complexity*. Helix Books. Addison-Wesley.
- IBM Global Business Services. 2010. *Capitalizing on Complexity: Insights from the Global Chief Executive Officer Study*. Executive Report. IBM.
- IBM. 2024. Agent Communication Protocol (ACP): Only One Protocol for AI Agents.
- Ilachinski, A. 2004. *Artificial War: Multiagent-Based Simulation of Combat*. World Scientific.
- Javaheripi, M., & Bubeck, S. 2023. Phi-2: The Surprising Power of Small Language Models. Microsoft Research Blog.
- Johnson, J. 2023. Automating the OODA Loop in the Age of Intelligent Machines: Reaffirming the Role of Humans in Command-and-Control Decision-Making in the Digital Age. *Defence Studies* 23 (1): 43–67.
- Joint Chiefs of Staff. 2017. Joint Publication 1: Doctrine for the Armed Forces of the United States. U.S. Department of Defense.
- Joint Chiefs of Staff. 2020. Joint Publication 5-0: Joint Planning. U.S. Department of Defense.
- Joint Chiefs of Staff. 2022. Joint Publication 3-0: Joint Campaigns and Operations. U.S. Department of Defense.
- Kalloniatis, A. 2025. Brief History of C2. In *Annals of C2* (Vol. II). NATO C2COE.
- Koestler, Arthur. 1967. *The Ghost in the Machine*. Macmillan.
- LangChain AI. 2024. LangGraph: A low-level orchestration framework for building stateful, multi-actor applications with LLMs. GitHub Repository. <https://github.com/langchain-ai/langgraph>
- Leitão, P. 2009. Holonic Rationale and Self-Organization on Design of Complex Evolvable Systems. *Holonic and Multi-Agent Systems for Manufacturing*, 1–12.
- Maier, M. 1998. Architecting Principles for Systems of Systems. *Systems Engineering: The Journal of the International Council on Systems Engineering* 1 (4): 267–84.
- McCann, C., & Pigeau, R. 1996. *Redefining Command and Control* (Chapter 12). *Human In Command: Exploring the Operational Environment*. Kluwer Academic / Plenum Publishers
- NATO Command & Control Centre of Excellence. 2019. *The Future of the Command Post: Part 1*. NATO C2COE.
- NATO Military Committee. 2023. MCM-0004-2023: Alliance Concept for Multi-Domain Operations. Memorandum. North Atlantic Treaty Organization.
- NATO Standardization Office. 2021. AAP-06: NATO Glossary of Terms and Definitions (English and French). 2021 Edition. North Atlantic Treaty Organization.

- NATO Standardization Office. 2025. Allied Joint Publication-3: Allied Joint Doctrine for the Conduct of Operations. Edition D, Version 1. North Atlantic Treaty Organization.
- NATO Science & Technology Organization. 2022. Agile Multi-Domain Command and Control. Report No. STO-TR-SAS-143. NATO STO.
- Natter, M., Ockerman, J., & Baumgart, L. 2010. Review of Cognitive Metrics for C2. *ITEA Journal of Test & Evaluation* 31 (2): 179.
- Niven, G., & Capewell, D. 2019. The Variety Calculus—an alternative proposition for command & control in a complex world. In *24th International Command and Control Research and Technology Symposium (ICCRTS)*.
- Open Networking Foundation. 2012. OpenFlow Switch Specification, Version 1.3.0. ONF.
- Rokvić, V. 2025. Mosaic Warfare as a New Art of War? *The Review of International Affairs* 76 (1195): 421–48.
- Ryan, A. 2008. Military Applications of Complex Systems. In *Philosophy of Complex Systems*. Elsevier.
- Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. 2016. Edge Computing: Vision and Challenges. *IEEE Internet of Things Journal* 3 (5): 637–46.
- Shoham, Y. 1994. Designing Efficient Multi-Agent Systems. STAN-CS-TR-94-1501. Stanford University, Computer Science Department.
- Simon, H. 1962. The Architecture of Complexity. *Proceedings of the American Philosophical Society* 106 (6): 467–82.
- Snowden, D., & Boone, M. 2007. A leader's framework for decision making. *Harvard Business Review*, 85(11), 68–76.
- Steinberg, A., Bowman, C., & White, F. 1999. Revisions to the JDL data fusion model. *Proceedings of SPIE*, 3719, 430–441.
- Tharumarajah, A., Wells, A., & Nemes, L. 1996. Comparison of the Bionic, Fractal and Holonic Manufacturing System Concepts. *International Journal of Computer Integrated Manufacturing* 9 (3): 217–26.
- Ulieru, M., Brennan, R., & Walker, S. 2006. The Holonic Enterprise as a Collaborative Information Ecosystem. *International Journal of Production Economics* 100 (1): 15–37.
- U.S. Air Force. 2021. The Air Force (Air Force Doctrine Publication 1).
- Van Bommel, I., van Hekken, M., Paulissen, R. Eikelboom, A., van de Kuijt, J., Benda, R., & Ziekenoppasser, L. 2019. NetForce Command – An alternative for hierarchical command and control in complex, dynamic and networked mission environments. TNO (in collaboration with the Land Warfare Centre of the Royal Netherlands Army)
- Van Brussel, H., Wyns, J., Valckenaers, P., Bongaerts, L., & Peeters, P. 1998. Reference Architecture for Holonic Manufacturing Systems: PROSA. *Computers in Industry* 37 (3): 255–74.
- Vaswani, A., et al., Attention Is All You Need. 2017. *Advances in Neural Information Processing Systems*.
- Waldrop, M. 1992. Complexity: The Emerging Science at the Edge of Order and Chaos. Simon & Schuster.
- Willcox, K., Bingham, D., Chung, C., et al. 2023. Foundational Research Gaps and Future Directions for Digital Twins. National Academies Press.
- Wooldridge, M. 2009. An Introduction to Multiagent Systems. John Wiley & Sons.

Author biography

Dr. Abder Rezak Benaskeur (Abderrezak.benaskeur@forces.gc.ca)

Abder Rezak Benaskeur has been a member of the Canadian Defence Team with Defence Research and Development Canada (DRDC) since 2000. He holds a Ph.D. in Control and Automation and a master's degree in Artificial Intelligence (AI). Throughout his tenure, he has contributed to the definition, development, and delivery of a series of high-impact, large-scale Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR) programs for the Canadian Armed Forces, the North American Aerospace Defence Command (NORAD), and international coalition partners. He also serves as the Canadian principal member for the NATO Command, Control, Communications, Computers, Cyber, and Intelligence (C5I) Scientific and Technical Committee.

His research and development efforts focus primarily on integrating advanced AI, automation, and decision-support technologies into Pan-Domain Command and Control (PDC2) frameworks for joint domestic and expeditionary missions, with a specialized emphasis on Integrated Air and Missile Defence (IAMD) operations. His interdisciplinary expertise spans control theory, human-autonomy teaming, distributed systems architectures, data fusion, adversarial plan recognition, and dynamic threat evaluation.



NATO
C2COE

A CONTROL-THEORETIC UNIFICATION OF DEFENSIVE AND OFFENSIVE C2

Dr. Abder Rezak Benaskeur
Defence Research and Development Canada
Department of National Defence



This article establishes a formal framework for Command and Control (C2) by mapping military operations to control theory, utilizing this mathematical lens to analyze the strategic relationship and underlying dichotomy between defensive and offensive postures. This operational dichotomy is shown to be isomorphic to the foundational control-theoretic duality between disturbance rejection (feedback control) and setpoint tracking (feedforward control). Expanding on this isomorphism, the paper advances three additional core theses. First, it establishes the fundamental primacy of defence by demonstrating that offensive postures are not independent objectives, but rather derivative, anticipatory control loops engineered to neutralize disturbances, and thus reduce entropy, based on predictions generated at the same echelon or received from higher command. Second, the paper analytically proves that defence is the more challenging posture because, while the attacker solves a simple maximization problem of finding a single path to high entropy, the defender must solve a computationally harder min-max problem to maintain low entropy against all possible disturbance vectors. Finally, the paper redefines the interaction of opposing forces, proposing that defence is governed by vulnerability minimization (robustness) whereas offence is driven by capability maximization (control authority). Ultimately, this framework provides a rigorous syntax for analyzing modern warfare, enabling a deeper understanding of strategic confrontations and the prediction of their outcomes, while concluding that future operational advantage lies in minimizing C2 loop latency and maximizing predictive model fidelity.

Keywords:

Warfare, Command & Control, OODA Loop, Offence, Defence, Control Theory, System Dynamics, Feedback Control, Disturbance Rejection, Feedforward Control, Setpoint Tracking.

1 Introduction

The study of C2 has long existed at the intersection of multiple disciplines, including military strategy, and doctrine, organizational psychology, and systems theory. In the military doctrine (Joint Chiefs of Staff 2017a; NATO Standardization Office 2021), C2 is often defined as follows:

C2 is the exercise of authority and direction by a properly designated commander over assigned and attached forces in the accomplishment of the mission.

This definition bifurcates the concept, which at least in theory symbolizes unity, into two distinct functions:

- Command entails the legal and moral authority to make decisions, assess risk, and assign missions. It is a fundamentally human, cognitive, and leadership-centric function.
- Control encompasses the systems, processes, and mechanisms used to direct, monitor, and adjust the execution of the missions assigned by the command function. It is procedural, technical, and increasingly a technology-heavy function.

1.1 The OODA Loop

C2 is executed through a series of complex decision-making processes, often conceptualized via the Observe-Orient-Decide-Act (OODA) loop (Boyd 1987), as illustrated in Figure 1. While this framework provides a useful baseline for discussion, it is important to note that most C2 processes are rarely as simple or strictly sequential as the OODA model's linear structure might suggest; rather, they function as an integrated, non-linear system of feedback and adaptation.




Figure 1: The Boyd's OODA Loop

The observation phase equates to the sensing function, that is the gathering of cues and intelligence from the environment of interest, such as a battlespace or an area of operation. The orientation phase corresponds to the situation of interest state estimation, filtering, inference, projection, and contextual interpretation. At this stage, data is fused to create a coherent situational awareness that represents the internal model presenting the state of the controlled system (e.g., the battlespace). This is the most critical phase of the OODA loop, as it is where analysis and synthesis occur to form mental models and situational awareness for the decision-makers. The decision segment is where the selection of options takes place. The current state of the system/battlespace is compared to its desired state, as expressed, for instance, by the Commander's intent, to generate courses of action and select the preferred option. The action phase represents the execution of the selected option as orders are given to create the desired effects within the theater of operations through the actual application of combat power and other available means.

The goal of any commander or decision-maker is to get inside the adversary/opponent's OODA loop to gain advantage by being able to go through the four phases faster and more effectively, and to saturate the opponent with more entropy (ambiguity, uncertainty, and dilemmas) than they can handle. Faster, low-entropy loops of one side inside slower, higher-entropy loops of the opponent create for the former a combination of the order and tempo dominance that leads to the much sought-after decision superiority.

1.2 Control-Theoretic Perspective of C2

The concepts and formalisms of control theory offer a rigorous language to describe, understand, and predict the dynamics of conflicts, and the underlying C2 processes, that is often absent in doctrinal literature. Warfare is fundamentally a dynamic process primarily involving opposing, red and blue, sides, actors, or simply agents, each attempting to impose their will upon the opponent and shape the theater of operations, subject to constraints and entropy, a.k.a. the fog of war. The latter is mostly due to disturbances created by

the combined actions of the opponent and the uncontrolled entities and phenomena active in the environment, as well as, in some cases, the uncertainty inherent to one side's own imperfect processes, behaviors, information, and systems.

From a control-theoretic perspective, a controller manages a dynamic system, often referred to as the plant in control theory, by directly or indirectly observing manifestations of its state and applying control actions to drive its future states toward more desirable conditions or equilibria. This functional analogy between control theory and military C2 is precise, instructive, and practical in many ways. This abstraction maps onto C2 when analyzed through its three irreducible elements:

- Doctrine is the body of fundamental principles that guides how military forces orchestrate their actions in pursuit of national objectives. It establishes a shared professional language and a baseline of expected behavior, enabling commanders to communicate their intent and desired end-state or set-point with precision. By providing this common framework, doctrine functions as the reference signal for the military's C2 system, defining the ways (how we fight) to achieve the ends (what we wish to achieve).
- Process embodies the cognitive and procedural control law of the organization. Exemplified by the OODA loop (see Section 1.1), it serves as the operational logic that transforms raw environmental inputs into actionable intelligence and decisive outputs. Process is the dynamic application of doctrine, ensuring that actions remain aligned with the commander's intent despite a changing environment.
- System constitutes the tangible implementation of military power. It is an integrated socio-technical structure comprising sensors, communication networks, decision-makers, and effectors. This structure serves as the interface between the command logic (i.e., the process) and the operational environment, providing the means or the with what required to translate decisions into physical effects.

Converting C2 within this control-theoretic framework reveals a more natural link and fundamental duality between offensive and defensive postures that represent the purpose of armed forces and the essence of most of their operations. A defensive system aims to maintain its state within a bounded performance or operational envelope despite external disturbances. Conversely, the purpose of an offensive activity or operation is to drive the system state away from its current equilibrium toward a new one. These postures are two complementary modes of a unified control process, whereas traditional military doctrine often treats them as distinct, opposing phases, each with its own set of objectives, rules, processes, systems, and strategies.

Fundamentally, control problems resolve to either driving a system toward a desired state or preventing active disturbances from driving it away. Offence and defence are pure embodiments of this reality, and thus structurally linked manifestations of a single, continuous control process.

1.3 Document Structure

The remainder of the paper is organized as follows. Section 2 reviews some of the offensive and defensive processes commonly used in modern warfare. Section 3 models the battlespace as a dynamic system. Section 4 proposes a control-theoretic unification framework to highlight the duality between offence and defence using the standard feedforward and feedback-based control loops. Section 5 explores the structural primacy of defence by framing offensive operations as anticipatory control loops serving higher-order, or projected, defensive objectives. Section 6 argues that defensive postures are more challenging than offensive postures, a phenomenon we will refer to as the attacker's advantage. Section 7 introduces the concepts of capability and vulnerability to offer a better understanding of the relationship between defence and offence, and how they shape the interaction between opposing forces. Section 8 concludes the paper.

2 Offensive and Defensive Processes

Modern warfare is governed by a set of structured processes that operationalize C2. All these processes are, to a large extent, instantiations of the OODA loop (see Section 1.1), created, in quasi siloes, by different operational communities, resulting in slight variations.

Defensive processes, which include the detect-to-engage (DTE) sequence and the joint engagement sequence (JES), are primarily reactive, triggered by the detection, or the expectation thereof, of a threat, which they aim to negate. Offensive processes, which include the kill chain and the targeting cycle, are proactive, driven by the commander's intent to reshape the battlespace and gain advantage. They focus primarily on generating own effects rather than negating those intended by the opponent.

2.1 Detect-to-Engage Sequence

The DTE sequence is a fundamental air defence process used, mainly in the naval domain, to neutralize incoming threats before they penetrate the defended area or asset's vulnerability radius, also referred to as the keep-out zone. This process is often characterized by high-speed execution and automation, given its tactical context. The sequence typically unfolds as follows:

- **Detect:** This consists in tasking of Intelligence, Surveillance, and Reconnaissance (ISR) assets to scan the volume of interest looking for presence of potential threats.
- **Localize:** This is the determination of the precise location of the objects. This is often combined with the Track phase.
- **Track:** This equates to monitoring of the objects' movement through the continuous measurement of their state (e.g., position, velocity) to establish and maintain tracks.
- **Identify:** This is the classification of tracked objects as friend, foe, neutral, etc.
- **Recognize:** This determines of the objects' class and class amplification based on their type, size, and shape.

- **Evaluate:** This corresponds to the assessment of the threat potential and severity for the established tracks based on their identity, type, as well as past and current behavior.
- **Plan/Decide:** This is the allocation of one or more effectors (missiles, guns, etc.) to counter the assessed threats based on the applicable engagement criteria.
- **Engage:** This equates to the execution of the planned course of action to neutralize threats.
- **Assess:** This where the verification of the outcome of the executed response plan, including kill assessment, is conducted.

2.2 Joint Engagement Sequence

The JES expands the DTE concept to the joint force level, particularly in Integrated Air and Missile Defence (IAMD). It involves the coordination of sensors and effectors across multiple domains. The JES emphasizes networked control, where the different functional steps of the DTE sequence are distributed across a web of dispersed nodes to maximize the probability of successful neutralization of the threat (Joint Chiefs of Staff 2017b). Most phases of the JES are identical to those of the DTE sequence, but with the added complexity of the requirement to coordinate sensors and effectors across different domains. The JES phases are:

- **Surveil:** This consists in monitoring the environment to detect the presence of potential threats.
- **Detect:** This is identical to the DTE sequence.
- **Track/ID:** This equates to the combination of localize, track, identify, and recognize of the DTE sequence.
- **TEWA:** This stands for threat evaluation and weapon assignment and combines evaluate and plan/decide of the DTE sequence.
- **Engage:** This is identical to the DTE sequence.
- **Assess:** This is identical to the DTE sequence.

2.3 Kill Chain

The kill chain is the offensive counterpart to the DTE, often referred to as F2T2EA, which stands for find, fix, track, target, engage, and assess:

- **Find:** This consists in locating a target within the area of interest. This is the equivalent of detect in the DTE sequence.
- **Fix:** This is equivalent of localize in the DTE sequence.
- **Track:** This is identical to the DTE sequence.
- **Target:** This selects the appropriate effector for the create the desired effect. This is similar to plan/decide in the DTE sequence.
- **Engage:** This is identical to the DTE sequence.
- **Assess:** This is identical to the DTE sequence.

Unlike the DTE process, which waits for a threat to enter the area of responsibility and the sensor's field of view, the kill chain actively searches for targets to engage as stated by the overarching mission objective (Brose 2020).

2.4 Targeting Cycle

The targeting cycle, also referred to as D3A, or decide, detect, deliver, and assess (Joint Chiefs of Staff 2013):

- **Decide:** This is the most critical phase, where priorities are set and high payoff targets are selected based on the Commander's intent and the mission objectives.
- **Detect:** This corresponds to the combination of find, fix, track phases of the kill chain process.
- **Deliver:** This is similar the engage phase of the other processes.
- **Assess:** This is identical to the other processes.

The main difference between the kill chain and the targeting process is that the latter puts the decide function, which is the equivalent to the kill chain's

target, first. The C2 system in the targeting process decides on the objectives, or what it wants to achieve, before starting to look for the target.

Overall, defensive processes are optimized for speed and robustness to minimize losses. Offensive processes are optimized for precision and effect, giving enough margin for much longer planning horizons. The next section shows that the distinction between offence and defence goes beyond these differences (in purpose) to include a more structural topology within the C2 loop.

Note that, in this context, the term loop is used broadly to describe a continuous feedback mechanism intended to influence or regulate the environment. It denotes a dynamic, adaptive system rather than a rigid, sequential iteration of discrete tasks.

3 Battlespace as a Dynamic System

To apply control theory to military strategy, we consider the theater of operations as a dynamic system Σ , whose state vector $x(t) \in \mathbb{R}^n$ represents the (physical, cognitive, informational, political, etc.) status of the conflict. This may correspond to territory held, force attrition, political will, population morale, etc. The dynamics of such a conflict can then be modeled by the state equation:

$$\dot{x}(t) = f[x(t), u(t), d(t)]$$

Where:

- $x(t)$ is the system state, i.e., the reality of the operational environment.
- $u(t)$ is the control input, i.e., the actions taken by the friendly forces. These actions are defined by the relevant elements of power at the disposal of the friendly forces, which may include military, economic, political, etc., measures.

- $d(t)$ is the disturbance vector, i.e., the actions taken deliberately by the enemy (as part of elements of power) or the uncontrolled environmental factors, which may also include interferences generated by noncombatant activities in the vicinity.

The objective of an effective C2 system is then to minimize the error function:

$$e(t) = r(t) - y(t)$$

defined as the difference between the desired state of the operational environment, representing the commander's intent or mission objective, $r(t)$ and the actual output $y(t)$, which represents the current or projected state of the operational environment as provided by ISR assets and processes. In this framework, the commander, or the entire command structure, acts as the controller that processes the error signal $e(t)$ and generates the control input $u(t)$ to drive the error to zero.

4 Duality of Postures

The first thesis of this paper concerns the relationship between defence and offence in the context of warfare.

Thesis 1. Defence and offence are dual constructs.

This duality leads us to reframe defence not as a passive waiting, but as a disturbance rejection, or feedback, control problem, where the goal is to maintain system stability and performance against exogenous shocks or stressors that seek to take the system away from its current state and toward a less desirable state. Conversely, offence is reframed as a setpoint tracking, or feedforward, control problem, where the controller seeks to drive the system state to a new equilibrium that represents a more desirable state. This duality is further illustrated in Figure 2.




Figure 2: Defence vs. offence

Feedforward and feedback control mechanisms are illustrated in Figure 3.

The main loop is a feedback loop that aims at maintaining the status quo by injecting energy into the area of responsibility, through the utilization of the elements of power, in reaction to enemy actions that present a disturbance to the system.




Figure 3: Duality of Warfare Postures

The feedforward loop is an offensive control mechanism used to gain advantage over enemy, as directed by higher command echelons, or to deny the enemy the freedom of action before the enemy can disturb the situation within the area of responsibility. The purpose of this proactive action is to avoid falling into a reactive/defensive posture, which is structurally weaker, as we will argue in Section 6.

4.1 Defence as Disturbance Rejection

Defensive operations align more with the problem of disturbance rejection, where the primary objective would be the preservation of the status quo that corresponds to the system performance as defined by $y(t)$. Examples of defensive objectives may include the preservation of territorial integrity, protection of force, or guarantee of operational freedom of maneuver. Mathematically, the reference signal is constant:

$$\dot{r}(t) = 0$$

or zero, if representing deviation from a norm:

$$r(t) = 0$$

The system is in an equilibrium state x_{eq} . Note that, in many cases, the internal state of the system is directly accessible to the ISR assets and only manifests itself through the output $y(t)$.

The enemy's attacks, certain non-combatant activities, and/or simple changes in the environment all may result in an exogenous disturbance $d(t)$ that pushes the system state away from the current state x_{eq} toward a less desirable state (see Figure 2). In this mode, the C2 system acts as a stabilizing controller that is continuously observing adversarial activities or environmental disturbances and applying counteractions to maintain the system within acceptable operational or performance bounds. Mathematically, the C2 problem can then be formulated as finding a course of action or plan defined by the control law:

$$u^* = \min_u \| r(t) - y(t) \|$$

in the presence of $d(t) \neq 0$. This is the classic regulator or disturbance rejection problem (Ogata 2010). The commander of the defending side monitors the battlespace, or area of responsibility, for deviations caused by the enemy actions. When a deviation, e.g. an incursion or attack, is detected, the commander applies combat power, or other non-military means made available to him, to restore the operational environment to its original state. This problem is also referred to as the maintenance control.

A concrete example is an Integrated Air & Missile Defence (IAMD) system, like Israel's Iron Dome, where, using the available interceptors that serve as the control law $u(t)$, the objective is to nullify the incoming rocket raids that represent the disturbance $d(t)$ and restore the area of responsibility to its original state that defines the reference signal $r(t)$, where the sky is free of hostile projectiles. The defence system does not seek to change the area of responsibility. Its sole purpose is to nullify the disturbance by absorbing the shock energy and return the area of responsibility to its original state.

The quality and effectiveness of the defence strategy is measured by the system's robustness, that is its ability to maintain stability despite large

entropy introduced by the stressors or disturbances $d(t)$. Another defining feature is the gain margin, that is the ratio of the system's gain to the gain of the disturbance. A high-gain defensive controller, such as a very strict not-one-inch doctrine, reacts aggressively to small disturbances. An illustrative example of such a doctrine is when the probability of raid annihilation (PRA) is used as the main metric by air and missile defence C2 systems.

Theoretically, a successful defence, in this case, will be the one that maximizes PRA. In practice, a single leaker will result in a failure, as PRA will instantly drop to zero. To avoid such a situation, a C2 system will tend to throw all it has at the incoming threats, leading to multiple overkills, the succession of which will unavoidably result in mission failure in the long run, under the assumption that the attacker has resources to sustain the offensive long enough to exhaust the defender.

4.2 Offence as Setpoint Tracking

In an offensive posture, the goal is to change the status quo as specified by the commander's intent, e.g., seize a new territory, destroy enemy capacity, or alter the political landscape. This intent defines the reference signal $r(t)$ for the force and its C2 system. Here, the reference signal $r(t)$ is time-varying and distinct from the current state/equilibrium.

$$r(t) \neq r(t_0) \text{ and } \dot{r}(t) \neq 0$$

Under the assumption of commander's rationality, this new target state must be more valuable/desirable than the current state (see Figure 2). The C2 problem can then be formulated as computing and executing the best plan (course of action or control law) $u(t)$ that minimizes the error:

$$e(t) = r(t) - y(t)$$

between the observed battlespace state and the commander's intent, driving the battlespace state $y(t)$ toward the defined objective, i.e., desirable state.

$$e(t) = r(t) - y(t) \rightarrow 0 \quad \text{as } t \rightarrow \infty$$

This is the servo or setpoint tracking problem. The commander of the attacking force is not reacting to disturbances or stressors to the current state of the battlespace but rather generating an error signal intentionally by setting a new goal, e.g., Capture Area A.

The control law $u(t)$ is applied to overcome the battlespace's inertia and the opposing friction, or enemy resistance, and move the state toward the target value $r(t)$. This problem is also referred to as the feedforward or transition control.

Offensive, or setpoint tracking, control requires high-fidelity and high-confidence models of the situation. Such situation models would encompass the current and projected red picture (enemy), blue picture (own and allied forces), and white picture (neutral entities and environmental conditions), all derived from the orientation phase of the loop (see Section 1.1). This heavy reliance on models makes the offensive mode very vulnerable to model mismatch. If the attacking side's model of the situation, and more importantly, the disturbances therein, is wrong or flawed, feedforward control will fail to drive the system toward the desired state, potentially leading to a significant residual error or even the instability of the whole system. This aligns well with the risk inherent in offensive operations, that is, if the intelligence is wrong, the attack fails, and the attacking side would be left overextended and overexposed.

4.3 Phase Transitions

Most control problems are characterized by fluid phase transitions between setpoint tracking and disturbance rejection; with warfare operations being no exception. In the latter, defensive and offensive postures are not static or sequential, but rather fluid, dynamic, and often intertwined phases. Every offensive, or setpoint tracking, action introduces a destabilizing disturbance into the adversary's defensive, or disturbance rejection, control loop. On the other hand, a successful defensive reaction may give rise to a new, emergent opportunity for a potential counter-offensive, or a new setpoint tracking situation.

This dynamic interplay between the two postures resonates with the core tenet of OODA loop, that is, to be inside the OODA loop of the opponent. In other words, decision advantage is not gained with superior resources, but with the ability to orchestrate the mode transitions more rapidly and coherently than the opponent. These transitions are not instantaneous, as, at the macro level, the entire OODA loop can be seen as a time delay τ . As τ increases, the system's phase margin decreases and the system risks instability. As a reminder, in control theory, the phase margin is a measure of a system's stability and robustness. It acts as a safety buffer against time delays. If a system has a large phase margin, it can handle significant delays or uncertainties in the feedback loop without oscillating or losing control. If the phase margin is zero or negative, the system is unstable.

The analysis of these time delays highlights a certain level of asymmetry between offence and defence. Defensive postures, or feedback control loops, generally have shorter τ due, for instance, to interior lines, known terrain, etc. Offensive postures, or feedforward control loops, are characterized by longer τ due, for instance, to complicated supply lines, unknown terrain, etc. The goal of an effective C2 will then be to artificially increase the enemy's τ , e.g., via deception or direct effects/shocks, until their (OODA) control loop becomes unstable.

5 Primacy of Defence

Most traditional defence relies on feedback control, which is inherently reactive. The fundamental limitation of such a mechanism remains causality, since the disturbance $d(t)$ must affect the system before the controller can react. In military terms, this means, for instance, that the territory is already under fire, and damage has occurred, at least to some extent. To improve performance and achieve perfect disturbance rejection, control theory introduced feedforward control. If the disturbance $d(t)$ can be measured or predicted before it impacts the system, the controller can apply a compensatory input to cancel it out. This leads us to the second thesis of this paper.

Thesis 2. Defence is the primary posture.

This thesis challenges the traditional view that offence and defence are independent, equal postures and argues that pure offensive postures, that is, aggression for its own sake rarely benefit a rational agent. As reminder, from a decision-theoretic perspective, a rational agent is an entity that selects actions to maximize its expected utility, given its current knowledge and beliefs. Rationality in this context is a mathematical definition of consistency between an agent's preferences, beliefs, and actions.

In both military doctrine and biological systems, offensive actions often emerge not as a desire for conquest or growth, but as a mechanism to neutralize disturbances that are projected from a model of the situation (i.e., based on one's beliefs) and/or derived from higher-echelon's defensive objectives.

- **Projected Disturbances:** In this case, the offensive actions aim at canceling out projected losses before they occur. For instance, when a military force launches a pre-emptive strike into enemy territory, they are acting on a predicted disturbance d_p . They are effectively neutralizing the threat, or disturbance vector, at its source while it is still forming and before it turns into an actual nuisance or disturbance $d(t)$ on the system. This is consistent with the conclusion of recent joint force studies (Joint Chiefs of Staff 2023), stating that the goal is to "*fight the enemy's plan, not just their forces*".
- **Derived Disturbances:** In this case, the offensive actions aim at securing a broader defensive objective, where offence becomes a subroutine of grand defensive strategy. For instance, upon a direction from higher echelon, a tactical force takes a local offensive action, by defining a new local setpoint, to neutralize a likely enemy offensive action that may create a global disturbance d_g and compromise the strategic position or equilibrium, hence contributing to a global disturbance rejection control loop.

The main observation here is that offensive postures are concrete and direct contributions to multi-echelon and/or anticipatory defence

strategies. These two situations are not necessarily mutually exclusive, as they can coexist and interact. This reality further complicates the C2 loop by blurring the distinction, not only between defence and offence, as discussed above, but also between pre-emptive offence and preventive offence. The latter are discussed in the next subsection.

5.1 Preventive vs. Pre-emptive Offence

Offensive military actions are often based on possibilistic and probabilistic intelligence rather than hard evidence provided by confirmed signals, such as from ISR assets. The key parameters in this distinction, or rather the confusion, between pre-emptive offence and preventive offence becomes the confidence level in the prediction model of the situation and enemy actions, as well as the thresholds applied to the disturbances d_p and/or d_g to justify, authorize, and trigger the offensive course of action u .

- Pre-emptive offence: Applies when the disturbance is expected with a high level of confidence and is imminent, i.e., the question is then only on 'when,' not 'if.' The disturbance rejection loop, or the mitigation strategy, has already started, but damage has not yet been sustained.
- Preventive offence: Applies when the disturbance is either anticipated, and still gathering or forming, or suspected with a certain level of likelihood but not confirmed. In control terms, this pushes the prediction horizon of the controller further into the future, resulting in a reduced confidence level of the predicted disturbance. In other words, there is a large uncertainty on 'when,' combined with some uncertainty on 'if.'

Both strategies place an immense burden on the processes taking place left of the launch, i.e., the Indicators and Warnings (I&W) and the Intelligence Preparation of the Battlefield (IPB). If the model used to generate d_p or d_g implies a threat that does not materialize, or not when and as expected, the system will have exerted a disproportionate and unjustified control energy u to cancel a phantom

disturbance, e.g., a falsely perceived threat, thus introducing a significant amount of instability into the system, or geopolitical situation, e.g., by starting an unnecessary, and otherwise avoidable, war.

The following subsections discuss examples, two from the warfare domain and one from the biological domain, that support the above thesis and discussion.

5.2 Anticipatory Defence Doctrine

The evolution of the United States defence strategy provides a definitive example of this shift from feedback to feedforward control architecture. Throughout the Cold War, the strategy was largely based on feedback; a strategy of containment and deterrence relying on retaliation after an attack. The post 9/11 era necessitated a shift toward what can be termed anticipatory defence, as codified in the 2002 National Security Strategy, often referred to as the Bush doctrine. This doctrine argued that the United States cannot let the enemies strike first.

This is a rejection of pure feedback or reactive control, that is, waiting for the disturbance or attack, in favor of feedforward or proactive control. The doctrine states (The White House 2002):

...adapt the concept of imminent threat to the capabilities and objectives of today's adversaries... The greater the threat, the greater is the risk of inaction... To forestall or prevent such hostile acts by our adversaries, the United States will, if necessary, act pre-emptively.

This example, and the associated doctrinal logic, remain consistent with Thesis 2, that is, the offensive strikes are not for the purpose of conquest, but as a perceived defensive necessity to protect the security of the defending side (the United States) and its interests against a projected future loss.

Recent geopolitical developments provide further context for this doctrine. For example, the United States policy decisions regarding Venezuela, including economic sanctions and the early 2026 law-enforcement intervention known as Operation Absolute Resolve, which resulted in the apprehension

of President Maduro, have been articulated primarily through the lens of mitigating regional risks to national security. Similarly, diplomatic discussions surrounding strategic interest in Greenland are often framed by policymakers as long-term measures to maintain stability in the Arctic amid shifting global competition. In these contexts, state actions that outside observers might characterize as offensive, such as economic pressure, military intervention, or territorial interest, are often rationalized by state actors as essential defensive measures to safeguard vital national security interests against projected geopolitical disturbances, losses, or threats.

5.3 Air & Missile Defence

In air and missile defence, counter-surveillance and counter-targeting operations offer tangible examples of offensive actions employed and justified as preventive defensive measures. For instance, by proactively attacking the surveillance and communications capabilities of the red force, assumed as a potential attacker, the blue force, or presumably the defending side, disrupts the adversary's OODA loop during the observation and orientation phases, long before the main attack (e.g., kinetic effect) is initiated by the red force (the decision phase) and materializes (the action phase).

This anticipatory action effectively breaks the red force or attacker's OODA loop upstream and prevents the engagement sequence from reaching its terminal phase, which presents the highest risk for the blue force or defender. Reaching this stage (the action phase) would force the blue force/defender to enter a high-gain disturbance rejection mode with minimal margins for error. Thus, attacking the attacker's eyes and ears is an, if not the most, effective form of preventive defence

5.4 Biological Aggression

The above ascertainment is not unique to warfare, as it is deeply rooted in biological systems too. In nature, offence is almost invariably a function of defence. The primary driver of aggression in nature is the protection of key resources such as food, territory, and offspring. The concept of defensive

aggression is used in behavioral ecology to describe offensive behavior displayed by an entity to prevent a loss of fitness options. A defensive animal will launch a violent, offensive attack to protect its state. For example, a mother bear or a lioness, justifiably or unjustifiably, perceiving a threat of total genetic loss, will defend her cubs by exhibiting a very aggressive behavior. This aggression is an offensive control $u(t)$ driven by a defensive objective $r(t)$ that represents the safety of offspring and a predicted disturbance $d_p(t)$ that represents the intruder's proximity, perceived as an indicator of a potential future loss, pushing the animal to act pre-emptively to cancel the threat before the loss is incurred. Thus, even in nature, the offensive act is a high-gain control response designed to maintain the status quo/survival against a projected high-entropy event, such as predation or starvation.

It is important to note that aggression is a costly strategy, presenting risks such as injury or death, and is therefore utilized only when the cost of inaction and the potential resulting loss of resources is higher (Archer 1988). Animals do not spend most of their time in the aggressive mode, and do not expend energy on aggression without survival imperatives, real or perceived.

In nature, as in control theory, offence or setpoint tracking is always a transient, or temporary, regime and never a steady-state or permanent one. A system, or any organism or entity, that spends most of its time in setpoint tracking, offence, or aggression, will not achieve much, will run out of resources, and will be subject to significant fatigue and exhaustion. To illustrate this principle, consider a formidable hunter such as the cheetah. For the latter, the cost of tracking preys at 70 mph is so exorbitant, pushing the animal's body to its absolute limits. Consequently, cheetahs cannot maintain their offence mode for prolonged periods. If a cheetah attempted to make prolonged or successive aggressive pursuits, its metabolism would quickly succumb to fatal overheating. This explains why cheetahs, as do many other felines and predators, treat the high-intensity chase as a strictly transient regime, spending most of their time in a low-energy steady-state of rest and observation to allow for recovery.

The above observations, similarly, explain the downfall, throughout history, of most expansionist empires that can, in a way, be considered as predators on the hunt. The latter have all been characterized by an inherent, often institutionalized drive to increase their territory, wealth, and influence through military conquests, colonization, or political annexations; aggressions often stretched far beyond its transient nature to become a permanent *modus operandi*.

Prosperity happens in peacetime, which corresponds to stable equilibria, and not during transient modes, or in unstable equilibria typical of wartime and aggression periods.

6 Attacker's Advantage

The title of this section is inspired by, and a direct reference to Carl von Clausewitz's concept of Defender's Advantage in his seminal work *On War* (Clausewitz 1832). The striking opposition of the two titles and positions they respectively carry, may come as surprise to many, is yet another evidence of how much warfare has evolved since the 19th century.

Clausewitz argued that defence is the stronger form of warfare and easier than offence for the following reasons:

- The defender only needs to keep the enemy from winning.
- The attacker requires much greater effort and resources.
- The defender can afford to wait and let the attacker to exhaust themselves, make mistakes, or give up.

Basil Liddell-Hart (Liddell-Hart 1933) fully supported Clausewitz's postulate that defence is easier than offence and updated it for the 20th century, noting that weapons such as machine guns and anti-tank guns, considered modern for that period, had made the defensive posture mathematically superior; often citing a ratio of 3:1 required for a successful attack.

These arguments, which may have been valid for the 19th and 20th centuries' operational friction of moving forces, are challenged by the many arguments presented in this section. The latter revisits the comparison of offence and defence in the context of the 21st century's warfare.

In control theory, the design of a control law $u(t)$ typically balances two competing objectives: tracking a reference signal $r(t)$ and rejecting exogenous disturbances $d(t)$, as discussed in the previous sections. While both are duals (see Section 4), they are not structurally and computationally equivalent, with the defence being significantly more demanding and challenging. This leads us the third thesis of this work.

Thesis 3. Defence is harder than offence.

The primary sources of difficulty in defensive missions are discussed in the next sections.

6.1 Information

In an offensive posture, or setpoint tracking, the controller generates the reference trajectory $r(t)$. The attacker knows exactly where, when, and how the system state should evolve. They possess full information about the desired future or target state. Conversely, the defender, or a disturbance rejection controller, must react to a signal $d(t)$ that is, by definition, unknown or stochastic at best. If $d(t)$ were fully known a priori, it would not be a disturbance but a deterministic input that could be perfectly cancelled out via feedforward control.

The defender is thus perpetually solving an estimation problem, that is, an observer design, alongside the control problem. This reality is clearly visible through the review of C2 processes presented in Section 2. Defensive processes, such as the DTE or JES, require more phases than the offensive ones, as they must not only localize and track the objects of interest, but also identify them, recognize them, and evaluate their threat level before action can be taken against them. All these extra steps, not necessarily the easiest, correspond to the problem of observer design of the control loop.

6.2 Causality

Defensive posture is fundamentally bound by causality. A feedback controller cannot act until an error:

$$e(t) \neq 0$$

is detected. This implies that a deviation, such as a potential threat breaching the area of responsibility or causing actual damage to protected assets, must occur before triggering the mitigation mechanism. According to mission success criteria, these deviations represent a degree of failure or loss and introduce an inherent time delay τ , or phase lag, into the defensive loop.

The attacker, operating in a feedforward mode, suffers no such lag regarding their own intent, since they execute a pre-planned control law $u(t)$. In the high-speed dynamics of modern warfare, e.g., involving hypersonic missiles or cyberattacks, if the time constant of the attack τ_a is smaller than the defender's reaction time τ_d , i.e., the attacker traverses his OODA loop faster than the defender, the overwhelmed defence system quickly becomes uncontrollable, and the battle is inevitably lost.

6.3 Entropy

The greater difficulty of defence can also be framed through the second law of thermodynamics. It is generally easier to create disorder, or entropy, than to maintain order. The defender's goal is to maintain the system internal state $x(t)$, or in most cases its external manifestation state $y(t)$, within a tight bounded set corresponding to a low-entropy state, despite disturbances $d(t)$. The attacker's goal is simply to push the system outside this set toward to a high-entropy state, and ultimately chaos.

A prime illustration of this principle is guerrilla warfare. Although the guerrilla's strategic goal is often defensive, such as protecting a homeland or expelling an occupier, their tactical actions are never passive. A guerrilla force simply cannot achieve their strategic goal (i.e., defence) by waiting in a trench; it must seize the initiative. To succeed

against a superior enemy despite limited resources and training, the guerrilla must constantly inject entropy in the environment. By choosing when, where, and how to fight, for instance through attacks, harassment, sabotage, and raids, they force the superior power to expend immense effort just to negate these disturbances and maintain a stable, low-entropy state.

Another illustration considers the C2 brain of an Integrated Air & Missile defence (IAMD) system. To succeed in maintaining low entropy, the defence system must successfully annihilate the entirety of the raid composed of N incoming warheads. The failure of a single intercept, that is a single leaker, will lead to a transition toward a high-entropy state, where damage or destruction is incurred. The attacker, conversely, only needs to succeed once, i.e., have one leaker. This was exactly the case during the 2026 Iran-Israel-US conflict. Although the U.S. and Israel employed highly sophisticated, multi-layered air defense systems, Iran exploited the principles of saturation and asymmetric attrition by deploying waves of low-cost drones alongside ballistic missiles. Because a defensive shield must maintain a near-perfect interception rate to prevent lethal leakers, the resulting cognitive and logistical strain eventually induced systemic failure, leading to the penetration of radar sites and regional bases.

Ultimately, the conflict demonstrated that technological superiority cannot compensate for the requirement to defend every point of interest against an adversary who needs only a single gap to succeed. This reality is captured with precision by the statement of the Irish Republican Army (IRA) following the 1984 Brighton hotel bombing:

Today we were unlucky, but remember we have only to be lucky once, you will have to be lucky always.

6.4 Complexity

The difficulty of defence scales exponentially with the complexity of the system, while the difficulty of offence scales linearly with the discovery of new opportunities provided by vulnerabilities in adversary's system (see Section 7 for more details).

A defensive C2 architecture must be robust, meaning it must maintain stability across the entire range of possible disturbances $d(t)$. This requirement for total robustness forces the defender to adopt conservative strategies that present high-gain margins, and which often sacrifices performance and agility. For instance, a firewall that blocks all malware might also block legitimate traffic, leading to a Denial of Service (DoS).

The attacker faces no such requirement for covering the entire space. Offensive C2 practices local optimization, meaning it optimizes its vector of attack $u(t)$ against a specific, identified subset of the defender's state space, representing possible vulnerabilities where the localized robustness is deemed to be minimal, offering a maximal opportunity. In cyber-warfare, this asymmetry is often referred to as the zero-day problem. A software system with millions of lines of code represents a vast attack surface. The defender must theoretically patch every possible vulnerability. The attacker needs only to find a single open path, i.e., the existence of a single exploitable opportunity. The computational complexity of verification, that is, proving a system has no vulnerabilities, is significantly higher than the complexity of exploitation, that is, finding one vulnerability. To rigorously analyze the interaction between opposing forces, and more importantly defence and offence, the concept of vulnerability is further discussed in next section and compared to the notion of capability.

The above elements support our argument that defence is the harder posture. Theoretically, the defender is required to solve a min-max problem, that is, minimizing the maximum possible damage from any disturbance $d(t)$:

$$\begin{aligned} J_d^* &= \min_u \max_d \int \mathcal{L}(\mathbf{x}, \mathbf{u}, \mathbf{d}) dt \\ &= \min_u \max_d \int (\mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{u}^T \mathbf{B} \mathbf{u} - \mathbf{d}^T \mathbf{d}) dt \end{aligned} \quad (1)$$

whereas, the attacker is only required to solve a simple maximization problem, that is, maximizing damage given a specific chosen path.

$$\begin{aligned} J_a^* &= \max_{u,d} \int \mathcal{L}(\mathbf{x}, \mathbf{u}, \mathbf{d}) dt \\ &= \max_{u,d} \int (\mathbf{x}^T \mathbf{A} \mathbf{x} + \mathbf{u}^T \mathbf{B} \mathbf{u} - \mathbf{d}^T \mathbf{d}) dt \end{aligned} \quad (2)$$

Mathematically, the min-max problem will always be computationally harder than the max problem. This makes reactive defence technically and operationally more demanding and challenging (Thesis 3). Therefore, an effective C2 architecture cannot rely solely on reactive or feedback defence. It must also incorporate anticipatory feedforward or offence elements, of a certain type, to lower the difficulty of the control task (see Section 5.1).

7 Vulnerability vs. Capability

The concepts of vulnerability and capability are introduced, decoupled from vague qualitative descriptions, and grounded in system dynamics. The capability of a given side, such as an offensive agent O , is not an absolute magnitude but a relative measure against the vulnerability of the opposing side, such as a defensive agent D . From a control perspective, as the robustness of agent D increases, the required control authority, or capability, of agent O to induce a state change must scale, at least, by an equivalent order of magnitude. Note that robustness is defined as inversely proportional to vulnerability.

Below, we propose a taxonomy that distinguishes between structural, or system-inherent, and situational, or context-dependent, attributes.

7.1 Structural Vulnerability

Structural vulnerability V_{str} represents the inherent susceptibility of the system to disturbances or stress. In control theory, this relates to the location of the system's open-loop poles defined by the eigenvalues of the system matrix $\mathbf{A}$ in Equations 1 and 2 above. A system with poles near the imaginary axis is significantly more structurally vulnerable than one with deep left-half-plane poles. In C2, for instance, a structural vulnerability can be attributed to the organization's topology, creating single points of failure, e.g., centralized command nodes, or unencrypted communication backbones. This determines the stability margin of the whole defence system.

7.2 Situational Vulnerability

Situational vulnerability V_{sit} represents a transient state of exposure. It corresponds to a specific configuration of the state vector $x(t)$ of the defending force at time t that places it, for instance, within the region of attraction of an adversary's effectors. Unlike structural vulnerability, which is constant/time-invariant, situational vulnerability $V_{sit}(t)$ is dynamic. Examples include a radar system temporarily shutting down for maintenance or a naval unit moving through an uncovered straight, exposing its flank to a potential attack.

7.3 Structural Capability

Structural capability C_{str} refers to the maximum theoretical control authority available to an agent or a force. It encompasses the range (reach), lethality, and bandwidth (number of simultaneous actions) of the force. In control theory, structural capability defines the reachability set, that is the set of all states the controller can theoretically force the system into, given infinite time. It is defined by the input matrix B (see Equations 1 and 2 above) and the saturation limits of the actuators u_{max} .

7.4 Situational Capability

Situational capability C_{sit} is often referred to as opportunity. It represents the subset of the capability that is actionable at time t given the set of constraints and restraints standing in current context. An artillery battery has high structural capability, but if it has no ammunition or no line of sight, its situational capability is zero.

The above definitions paved the way for our next thesis.

Thesis 4. The determining factor for defence is vulnerability, while for offence is capability.

Based on the above definitions and discussion, a defensive posture is fundamentally an optimization problem focused on minimizing vulnerability. The primary objective of the maximum defence is to eliminate structural vulnerability through passive measures such as redundancy, diversification, and hardening. If

$$V_{str} \rightarrow 0$$

which corresponds to perfect structural robustness, the system becomes insensitive to disturbances $d(t)$ regardless of the opponent's capability, be it structural or situational. In such a case, no active defence is required. However, since the total absence of vulnerability is impossible to achieve, as no perfect passive measure exists, defence strategy must also incorporate mechanisms to manage situational vulnerability. Active defence measures or disturbance rejection control methods are required precisely because the system retains some structural vulnerability, which the offensive agent O may want to exploit if the opportunity presents itself. The defending agent D will have to monitor the state $x(t)$ to ensure it does not drift into a region of high V_{sit} .

Conversely, an offensive posture is driven by capability. The attacking agent O initiates an action if, and only if, it estimates that its:

- Structural Capability $C_{strO}(t)$ exceeds the structural robustness, defined as the inverse of structural vulnerability $V_{strD}(t)$, of the defending force or agent D .
- Situational Capability $C_{sitO}(t)$ exceeds the situational robustness, defined as the inverse of situational vulnerability $V_{sitD}(t)$, of the defending force or agent D .

In control theory, this refers to the control energy applied that must be greater than the control energy required to overcome the inertia of the system. Within our control-theoretic framework, the setpoint tracking or offence is valid only if the controller corresponding to the attacking side possesses sufficient control energy $\|u\|^2$ to overcome the system's natural damping, which corresponds to the defending side resistance. The offensive is mathematically doomed to fail, i.e., hit an uncontrollable target state, if

$$C_{strO} < [V_{strD}]^{-1} \text{ or } C_{sitO} < [V_{sitD}]^{-1}$$

8 Conclusion

This paper discussed the military concepts of offence and defence as isomorphic to the control concepts of setpoint tracking and disturbance rejection. Table 1 summarizes the key points of this discussion.

Table 1: Comparison of Warfare Postures & Processes

Posture	Defensive	Offensive
Processes	DTE & JES	F2T2EA & D3A
Type	Reactive	Proactive
Trigger	Threat Presence	Commander's Intent
Mode	Disturbance Rejection	Setpoint Tracking
Objective	Minimize Vulnerability	Maximize Capability
Metric	Probability of Survival	Probability of Kill

By mapping military postures to control theory, the paper provided a rigorous syntax for operational analysis. The proposed framework validated three fundamental theses:

1. Mathematical Duality: Defence is rigorously mapped to disturbance rejection, or maintaining stability against shocks, while offence is mapped to setpoint tracking, or driving the state to new coordinates.
2. Structural Primacy: Strategically, offence is not an independent end but a feedforward control mechanism serving global or projected defensive objectives. Defence is the primary posture, while offensive actions are derivative subroutines designed to reject predicted or projected disturbances.
3. Asymmetric Difficulty: Defense is structurally the harder posture. This asymmetry manifests in the operational effort of gauging readiness, the compounding cost of mitigating every vulnerability, and the disproportionate severity of potential losses. In optimization terms,

the defender must solve a complex minimax problem, that is preserving low entropy across all conceivable disturbance vectors, whereas the attacker faces a far simpler maximization search, needing only to identify a single path to high entropy.

Furthermore, the paper redefined the interaction between opposing sides through the lens of vulnerability versus capability, demonstrating that an offensive action is only viable when the attacker's capability exceeds the defender's robustness.

Looking forward, modern warfare will increasingly be characterized by the proliferation of smart and deceptive adversaries, enabled by advancements in the fields of Artificial Intelligence and autonomous systems, equally for blue and red sides. Future research must, therefore, focus on the application of robust and learning control methods to adversarial environments, where disturbances are not just stochastic, but intelligent and adaptive, mainly driven by a hostile intent to maximize entropy.

The tempo of operations accelerates, e.g., with the advent of hypersonic missiles, the phase margin, or the available reaction time, of C2 processes and the underlying systems shrinks. Thus, the outcome of future conflicts and confrontations will most likely be decided by which side minimizes the latency of its decision processes and maximizes the accuracy of the prediction and inference models.

References

Archer, John (1988). *The Behavioral Biology of Aggression*. Cambridge University Press.

Boyd, John R. (1987). *A Discourse on Winning and Losing*. Air University Press.

Brose, Christian (2020). *Kill Chain: Defending America in the Future of High-Tech Warfare*. Hachette Books.

Clausewitz, Carl von (1832). *On War*. Edited and translated by Micheal Howard and Peter Paret. Princeton University Press, 1989.

Liddell-Hart, Basil (1933). *The ghost of Napoleon*. Faber & Faber.

Joint Chiefs of Staff (2013). *Joint Publication 3-60: Joint Targeting*. U.S. Department of Defense.

Joint Chiefs of Staff (2017a). *Joint Publication 1: Doctrine for the Armed Forces of the United States*. U.S. Department of Defense.

Joint Chiefs of Staff (2017b). *Joint Publication 3-01: Countering Air and Missile Threats*. U.S. Department of Defense.

Joint Chiefs of Staff (2023). *Joint Concept for Competing*. Technical Report. Department of Defense.

NATO Standardization Office (2021). *AAP-06: NATO Glossary of Terms and Definitions (English and French)*. 2021 Edition. North Atlantic Treaty Organization.

NATO Standardization Office (2022). *Allied Joint Doctrine (AJP-01, Edition F, Version 1)*.

Ogata, Katsuhiko (2010). *Modern Control Engineering*. 5th ed. Prentice Hall.

The White House (2002). *The National Security Strategy of the United States of America*.

Author biography

Dr. Abder Rezak Benaskeur (Abderrezak.benaskeur@forces.gc.ca)

Abder Rezak Benaskeur has been a member of the Canadian Defence Team with Defence Research and Development Canada (DRDC) since 2000. He holds a Ph.D. in Control and Automation and a master's degree in Artificial Intelligence (AI). Throughout his tenure, he has contributed to the definition, development, and delivery of a series of high-impact, large-scale Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR) programs for the Canadian Armed Forces, the North American Aerospace Defence Command (NORAD), and international coalition partners. He also serves as the Canadian principal member for the NATO Command, Control, Communications, Computers, Cyber, and Intelligence (C5I) Scientific and Technical Committee.

His research and development efforts focus primarily on integrating advanced AI, automation, and decision-support technologies into Pan-Domain Command and Control (PDC2) frameworks for joint domestic and expeditionary missions, with a specialized emphasis on Integrated Air and Missile Defence (IAMD) operations. His interdisciplinary expertise spans control theory, human-autonomy teaming, distributed systems architectures, data fusion, adversarial plan recognition, and dynamic threat evaluation.



NATO
C2COE

CLAUSEWITZIAN FRICTION IN MULTI- DOMAIN C2

Clausewitz's fog of war persists and intensifies in multi-domain operations (MDO), where cyber and social domains saturate command and control systems with disinformation, distributed denial-of-service attacks, and algorithmic manipulation, eroding shared situational awareness in volatile, uncertain, complex, and ambiguous (VUCA) environments. This article traces fog's evolution from classical epistemological uncertainty to MDO's ontological contestation of reality itself, analyzing systemic degradations in NATO C2 architectures: Joint Common Operational Picture corruption, OODA loop delays, cultural friction amplified by Hofstede's dimensions in coalitions, and hierarchical vulnerabilities to decapitation. It proposes integrated palliatives, dispersed headquarters with Federated Mission Networking, predictive AI filtering hybrid fog, immersive VAR/AR visualization and Auftragstaktik (Mission Command) revitalized through intercultural training. These mitigations restore temporal superiority and coalition resilience against hybrid threats, reframing friction as navigable operational terrain. Phased implementation through 2035 offers NATO decisional primacy in peer competition, vindicating Moltke's decentralized legacy for cyber-social eras while acknowledging residual risks: technological dependency, AI ethics, and friction's ineradicable persistence demanding adaptive human-AI symbiosis.

Keywords:

Operational friction, Multi-domain operations, Decision-making, C2 systems

"Friction is the only conception which, in a general way, corresponds to that which distinguishes real war from war on paper."

- Carl von Clausewitz, 1889 -

1. Introduction

Carl von Clausewitz captured the essence of warfare's inherent unpredictability with his seminal metaphor of the "fog of war," describing it as the pervasive uncertainty that shrouds military operations: "War is the realm of uncertainty; three quarters of the factors on which action in war is based are wrapped in a fog of greater or lesser uncertainty" (Clausewitz 1989, p. 101). This fog, intertwined with the concept of *Reibung* or friction—the accumulation of physical, moral, and informational impediments that render even simple actions extraordinarily difficult—forms the theoretical bedrock for understanding command and control (C2) challenges. Friction manifests as "countless minor incidents—the kind you can never really foresee [that] combine to lower the general

level of performance, so that one always falls far short of the intended goal" (Clausewitz 1989, p. 119), distinguishing real war from abstract planning. In contemporary multi-domain operations (MDO), this classical formulation evolves dramatically, amplified by the cyber and social domains into a hybrid fog that overwhelms sensors with deepfakes, Distributed Denial of Service (DDoS) attacks, and algorithmic disinformation, engendering volatile, uncertain, complex, and ambiguous (VUCA) battlespaces.

The relevance of this evolution to NATO C2 cannot be overstated. Traditional domains—land, sea, air—already strain shared situational awareness (SA) through incomplete intelligence and nonlinear interactions, as Barry Watts taxonomized friction into human limitations, informational uncertainty and structural nonlinearity (Watts 1996, p. 77). MDO exacerbates these: cyber intrusions manipulate data streams, while social media floods the information environment with adversarial narratives, eroding the Joint Common Operational Picture (JCOP) essential for coalition coherence¹. In Ukraine (2022-2026), Russian hybrid tactics exemplify this, where cyber disruptions and social disinformation saturated

¹ The JCOP is prioritized because informational and social manipulation targets the shared operational truth, directly fracturing coalition interoperability and amplifying cultural frictions under duress.

Ukrainian C2, slowing the decision-making cycle despite superior intelligence, surveillance and reconnaissance (ISR) —demonstrating friction's persistence amid technological proliferation. NATO doctrines like AJP-01 emphasize “centralized intent and decentralized execution” to counter such fog, yet centralized HQ remain high-value targets (HVTs), vulnerable to precision strikes or electronic warfare. Cultural frictions, per Hofstede's dimensions (e.g., power distance varying across allies), further degrade multinational SA, fostering mistrust under duress.

This article's objective is to reexamine Clausewitzian friction through an MDO lens, proposing dispersed, resilient HQs as a doctrinal evolution. It advocates functional nodes linked by redundant networks and AI-driven filtering to mitigate cyber-social fog without sacrificing unity of effort. Such structures eliminate single points of failure, accelerate OODA via edge-processing and align with *Auftragstaktik*'s initiative emphasis.

The structure proceeds logically: Section 2 traces fog's conceptual evolution from Clausewitz to MDO

hybrids; Section 3 dissects C2 degradations (e.g., cognitive overload, targeting delays); Section 4 outlines palliatives (AI, VAR, intercultural training) and Section 5 synthesizes strategic implications for NATO decision-makers. A central question animates the analysis: How can C2 structures dissipate multi-domain fog while preserving coalition cohesion?

2. The Concept of Fog of War in Clausewitz and its Evolution

Clausewitz's original conception frames the fog of war as profound informational uncertainty: the “uncertainty of all data” permeating *On war* (1832, Book I, Ch. 3), where combatants grapple with veiled realities. Friction (*Reibung*), its mechanical counterpart, renders execution laborious: “This tremendous friction is everywhere in contact with chance, and thus incidents take place upon which it was impossible to calculate, their chief origin being chance” (Clausewitz 1989, p. 120). This duality—fog obscuring, friction resisting—distinguishes abstract theory from battlefield chaos, aggregating physical




Image 1: Bataille d’Jena, livrée le 16 Octobre 1806” (Battle of Jena, delivered on 16 October 1806), engraving after Jacques François Joseph Swebach (1769–1823), ca. 1806–1850. Depicts Napoleon's forces during the Battle of Jena-Auerstedt (14 October 1806).

(terrain delays), moral (fear-induced paralysis) and intellectual (biased intel) drags into inescapable general friction.

Prussian reforms post-Jena-Auerstedt² (1806) birthed *Auftragstaktik* as antidote. Prussian superiority evaporated against Napoleonic fluidity; rigid hierarchies froze initiative amid fog (Shamir 2011, p. 29). Scharnhorst and Gneisenau, via Kriegsakademie, instilled “correct concepts grounded in experience,” training officers for independent judgment and *Fingerspitzengefühl*—fingertip feel—to pierce uncertainty without micromanagement (Shamir 2011, p. 33). Clausewitz championed moral forces, arguing centralized control amplifies friction in dynamic war.

Napoleonic innovation accelerated this: autonomous corps—self-sustaining combined-arms—enabled decentralized execution within unified intent, slashing fog propagation. Davout’s isolated victory at Auerstedt exemplified: days without orders, exploiting gaps via initiative (Britt 2003, p. 33). Moltke the Elder refined: “No plan of operations extends with any certainty beyond the first contact with the main hostile force”; subordinates must “act without orders in accordance with [higher] intentions” (Hughes 1993, pp. 131-132). Austro-Prussian (1866) and Franco-Prussian (1870) wars proved it: decentralized corps outmaneuvered centralized foes, validating friction mitigation through calibrated trust.

The twentieth-century carnage reaffirmed fog’s tenacity. World War I trenches embodied physical-

moral friction—mud, gas, attrition defying plans; World War II blitzkrieg highlighted cognitive strains, Germans dominating OODA until Allied adaptation (e.g., Normandy air superiority). Boyd revolutionized friction as an interactive weapon—“Clausewitz concerned with reducing friction; failed to address magnifying adversary’s” (Boyd 2018, p. 58)—via OODA loops disrupting enemy orientation. Harmony (building trust via training) minimized internal friction; tempo imposed external chaos³.

The tenacity of Clausewitzian friction had already been validated during the Cold War; the joint ACEVAL/AIMVAL programs (1974-1978) demonstrated that technological parity washes out when tactical fog amplifies cognitive overload⁴. Despite these warnings, the subsequent Digital Revolution in Military Affairs (RMA) promised nothing short of fog’s obsolescence, positing omnipresent ISR coupled with network-centric warfare architectures as panaceas for uncertainty. Visionaries like Cebrowski and Garstka envisioned ‘dramatically better battlespace awareness,’ linking every sensor to every shooter in a web of information dominance to preempt friction through real-time transparency (Cebrowski & Garstka 1998, p. 28). The Gulf War of 1991 tantalized with proof-of-concept: precision-guided munitions decimated Iraqi forces under ISR superiority, suggesting fog’s retreat before technological inevitability. Yet reality tempered hubris. Situational awareness deficits endured profoundly—Watts documented that 80% of air-to-air losses stemmed from unawareness, underscoring how human perceptual limits persist despite hardware revolutions (Watts 1996, p. 78).

2 The twin battles of Jena and Auerstedt (October 14, 1806) resulted in the catastrophic defeat of the Prussian army by Napoleon’s forces. This military collapse exposed the obsolescence of Prussia’s rigid, un-reformed Frederician system, serving as the catalyst for comprehensive institutional and tactical overhauls led by Scharnhorst and Gneisenau.

3 Within John Boyd’s OODA loop framework, “harmony” is not an abstract ideal, but the operational prerequisite for fast-paced, decentralized command. Boyd argued that a shared implicit orientation, forged through common training and absolute trust, allows distributed units to act cohesively without constant top-down direction. When tactical or cultural fog destroys this harmony, internal friction spikes, paralyzing the organization’s collective capacity to out-pace the adversary’s decision cycle.

4 The Air Combat Evaluation (ACEVAL) and Air Intercept Missile Evaluation (AIMVAL) were joint U.S. Air Force and Navy test programs (1974–1978) designed to assess fighter tactics and missile concepts post-Vietnam. They validated their findings through hundreds of highly realistic, heavily instrumented live-fly simulations at Nellis AFB, pitting advanced F-14s and F-15s against simpler F-5 aircraft. Crucially, the tests proved that when technological advantages forced engagements into chaotic, multi-aircraft visual range dogfights, the resulting cognitive overload neutralized hardware superiority—demonstrating that situational awareness, rather than technological determinism, remained the dominant factor in combat outcomes.

Cyberspace and space domains then birthed a pernicious “black fog,” where fabricated data eclipses truth through hacks, spoofing, and algorithmic deception, inverting RMA’s transparency dream. Russia’s 2008 Georgia invasion epitomized this: cyber units flooded civilian phones and computers with panic-inducing messages, clogging evacuation roads and delaying Georgian defenses without a single kinetic shot fired (Arquilla 2021, p. 5). This non-kinetic friction—manipulated information as force multiplier—exposed RMA’s contradictions: technologies reducing physical and human frictions spawn informational blackouts of unprecedented scale. Gardner elucidates the irony: while UAVs and networks alleviate exertion, adversarial countermeasures render battlespaces more opaquer, nonlinearity magnifying minor deceptions into strategic paralysis (Gardner 2024). Thus, digital promise circles back to Clausewitzian essence, demanding C2 evolution beyond sensor proliferation.

Multi-domain operations represent the culmination of fog’s hybridization, seamlessly integrating cyber and social domains with traditional kinetic ones to create unprecedented layers of opacity in the battlespace. The Russian Gerasimov Doctrine provides a stark illustration of this shift. It envisions conflict not as linear clashes but as non-contact warfare, blending military force with non-military instruments like disinformation campaigns to saturate enemy C2 without ever engaging in direct confrontation (Gerasimov 2013). The ongoing Ukraine conflict from 2022 to 2026 vividly operationalizes this hybrid approach. Russian cyber operations, such as the Viasat satellite hack that blinded Ukrainian communications at invasion’s outset, merged with relentless social media barrages—including deepfakes and Telegram bot swarms—to fragment Ukrainian situational awareness. Despite influxes of Western ISR assets, these tactics slowed OODA loops, underscoring

fog’s enduring potency even amid technological superiority (Google Threat Analysis Group 2023). As NATO’s AJP-6 doctrine warns, adversaries now contest the very information environment, turning data into a weaponized medium.

Cyber fog has emerged as the most insidious variant. DDoS attacks overwhelm bandwidth, crippling multinational mission network pipelines⁵, while generative AI fabricates hyper-realistic adversarial inputs that infiltrate sensor feeds undetected. This marks a profound evolution from Clausewitz’s physical fog of obscured terrain or smoke. Cyber denial is ontological—it doesn’t merely hide information but corrupts the data streams themselves, as seen in the Stuxnet worm’s 2010 sabotage of Iranian centrifuges through invisible digital sabotage (Rid 2013, pp. 42-46). AI’s brittleness compounds the threat. Algorithms trained on specific datasets, like Russian T-90 tanks, falter against variants such as Chinese Type 99s. Subtle adversarial perturbations—mere pixel tweaks—can spoof recognition systems entirely, transforming reliable ISR into a liability (Scharre 2018, p. 231).

Parallel to cyber incursions, social fog elevates online networks to a de facto “fifth domain.” State and non-state actors deploy botnets and algorithmic amplification to shape narratives, targeting adversary morale and cohesion. ISIS’s Syria campaigns turned routine drone strikes into viral spectacles via warstreaming, converting tactical events into strategic setbacks through global outrage or recruitment booms (Borchert, Schütz, and Verbovszky 2021, p. 56). In NATO coalitions, Hofstede’s cultural dimensions exacerbate these frictions. Nations with high power distance, like France, resist subordinate initiative, while mismatches in individualism versus collectivism foster miscommunications under stress—as evidenced in ISAF Afghanistan, where cultural fog

5 These generic multinational networks find their modern institutional framework in NATO’s Federated Mission Networking (FMN) initiative. FMN is an evolving, capability-based framework that establishes interoperable mission networks for coalitions through synchronized sets of technical and operational standards, known as “Spirals.” Consequently, cyber denial and bandwidth throttling within these pipelines do not just disrupt local communications, they directly compromise the standardized infrastructure required to dynamically integrate allied Command and Control (C2) systems.

dissolved multinational C2 unity (Auerswald & Saideman 2014).

To clarify the transformation from classical to MDO fog, let's walk through key contrasts that bring Clausewitz into the digital age. Think back to his examples: "mud slowing cavalry or the unnerving whistle of cannonballs hiding enemy positions" (On War, p. 113). Now picture DDoS attacks choking your real-time JCOP feed—same delay, new domain. Fear warping a soldier's judgment evolves into AI adversaries tricking image recognition with invisible pixel tweaks. At its core, fog shifts from knowing less (epistemological gaps from bias or hiding) to questioning reality itself (ontological assault via deepfakes making truth optional) (Lindsay 2020). This isn't evolution; it's reinvention.

Barry Watts helps us see why fog never really goes away. His breakdown holds up perfectly today: humans drown in data floods from drones and social feeds; enemies hide via cyber tricks; small hacks snowball into network meltdowns (Watts 1996, pp. 76-77). Nikolas Gardner nails the twist—UAVs and AI ease physical strain but create fresh headaches with black-box decisions and surprise escalations (Gardner 2024). Technology promises light but often delivers denser shadows.

Pulling it together, fog has journeyed from Clausewitz's hazy uncertainty to MDO's reality-warping storm. NATO can't ignore it—decentralized setups like C2COE's dispersed HQ with FMN are the answer, blending human gut instinct and AI smarts to wrestle friction into submission. The grip loosens, but never breaks.

3. Consequences of Clausewitzian Friction in Modern C2 Environments

Multi-domain operations have supercharged Clausewitzian fog into a hybrid force that relentlessly undermines C2, turning friction's unpredictable drag into cascading systemic failures across every operational layer. Classical fog might have blinded isolated battles with smoke or terrain; today's cyber-social infusion strikes deeper, corrupting the digital backbone of C2—poisoning JCOP and stalling OODA loops until entire coalitions grind to a halt. NATO's traditional operational headquarters, with their centralized behemoth designs, only magnify these weaknesses, as C2COE studies repeatedly warn (C2COE, 2020)⁶. What follows is a clear breakdown of the main damage—situational awareness collapse, cognitive overload, cultural breakdowns, and more—making the case for urgent, resilient redesigns that can fight back.

The degradation of situational awareness stands as fog's primary C2 assault. Adversarial data falsification—deepfakes infiltrating feeds, spoofed sensor inputs—poisons JCOP, the shared battlespace visualization critical for coalition synchronization. In Ukraine, Russian cyber ops flooded Ukrainian networks with fabricated troop movements, eroding COP reliability and inducing paralysis (Google TAG 2023). This mirrors Watts' informational uncertainty: biases and concealment yield chronic intel deficits, now supercharged by AI-scale manipulation (Watts 1996). Resultant OODA delays prove fatal; Boyd's tempo advantage evaporates as observe-orient phases stretch amid contested data, allowing adversaries inside decision cycles. C2COE's "Future Command Post" warns centralized JCOP nodes become single points of failure, vulnerable to DDoS-induced blackouts that ripple fires and maneuver (C2COE 2019a)⁷.

6 This study explicitly warns that traditional, centralized physical headquarters act as high-value targets (HVTs) easily detected by multi-domain ISR and neutralized by precision strikes, necessitating an urgent shift toward physically dispersed and virtually federated command nodes.

7 While historical data (e.g., Sparks 1984) attributed tactical losses to a lack of visual awareness, contemporary NATO Science and Technology Organization (STO) research shows this vulnerability has changed. In modern network-centric warfare, SA deficits are driven by cognitive overload and data saturation. As Mica Endsley's human factors research demonstrates, when operators are flooded with multi-domain sensor feeds, their capacity to comprehend and project threats collapses, validating that human perceptual limits remain the primary bottleneck despite hardware revolutions.

Cognitive overload hits next, turning MDO's data explosion into a paralyzing flood for centralized headquarters staff already stretched thin. Imagine petabytes pouring in hourly—drone feeds, wearable biometrics, social media sentiment analysis—overwhelming human brains designed for far less. Endsley's situational awareness pyramid crumbles: basic perception drowns, comprehension fragments, projection becomes guesswork (Endsley 1987). Functionally siloed HQ teams waste 70% of time sifting noise for signals, per C2COE operator surveys, crippling timely sensemaking (C2COE 2019b). Social media's viral lies spike the chaos further, exploiting Hofstede's cultural differences—high-context teams miss low-context cues—that block shared pattern recognition. Cold War-era centralized designs worsen it: massive footprints scream “target me” to enemies, while rigid hierarchies choke data fusion at every layer.

Cultural friction exacerbates misunderstandings in multinational coalitions and lets them spread like wildfire under fog's pressure—Hofstede's framework makes this painfully clear. Consider power distance gaps, high scores in Latin/Asian partners (90+ for France or Spain) make subordinates hesitate for explicit orders, while low Germanic scores (35 for Germany) push bold initiative. Ambiguity fractures mission command; teams freeze awaiting clarification (Hofstede et al. 2010). The ISAF mission in Afghanistan illustrated this friction, where clashing national military cultures and differing operational tolerances eroded mutual trust and delayed joint execution (Auerswald and Saideman 2014). MDO's social fog weaponizes these rifts—Gerasimov-style misinformation tailors narratives to exploit divides, sowing discord. Army research pins it down: cultural mismatches double decision times in coalitions. Centralized C2 pours fuel on the fire, hierarchies hardening distance gaps.

Cyber vulnerabilities strike at the heart of hybrid fog's damage to C2, methodically dismantling multinational network architectures instantiated under FMN framework. Centralized hierarchies

create glaring single points of failure that hackers and electronic warriors exploit ruthlessly. DDoS floods or precision jamming shred the standardization agreements of active FMN Spirals, cutting sensor-to-shooter pipelines and triggering blackout cascades, Russia's NotPetya ransomware (2017) illustrated this operational paralysis, devastating critical Ukrainian infrastructure and military-linked logistics during ongoing hybrid conflict (Scharre 2018). C2COE studies flag static HQ as prime HVTs, where one compromised router freezes coalition JCOP entirely (C2COE 2019a). Watts' nonlinearity plays out viciously: tiny breaches spawn massive paralysis, flipping “resilient” networks fragile under fog (Watts 1996). AJP-6 (Edition A) warns of FMN fragility in a contested electromagnetic spectrum, screaming for redundancies legacy designs ignore.

Decision cycles grind to a crawl under relentless assault, targeting and fires chains strangled by distorted inputs. The F2T2EA process—find, fix, track, target, engage, assess—seizes as disinformation pumps false positives, warping Measures of Effectiveness (MoE) and Measures of Performance (MoP). Ukraine watched Russian decoys spoof HIMARS strikes, delays compounding per C2COE reviews (Google TAG 2023). OODA's orient-decide phases explode amid contested feeds, handing tempo to enemies. NATO C2COE command post research indicates that tactical fog can double fires timelines within complex joint environments⁸, while social psychological operations fabricate ‘victories’ to misdirect strategic leaders (C2COE 2019a). Top-down hierarchies worsen this paralysis, as rigid validation queues systematically eliminate agility.

Pulling these threads together reveals a vicious cycle—SA crumbling, cognition paralyzed, cultures clashing, networked data integrity compromised, and decisions stalled—that leaves legacy C2 systems brittle against hybrid threats.

Mission command crumbles next in this cascade, as fog systematically erodes the trust essential for decentralized execution, replacing initiative

⁸ While this study predates NATO's formal adoption of Multi-Domain Operations (MDO) terminology in 2021, its findings on high-intensity joint operations remain prescient. The research demonstrates that simulated informational saturation and degraded networks double targeting timelines by overwhelming human cognitive processing within the F2T2EA cycle.

with risk-averse micromanagement. ADP 6-0 and Auftragstaktik thrive on confidence amid uncertainty; fog obliterates it, sparking risk aversion as subordinates sense warped SA and freeze (U.S. Army 2022). Hofstede divides intensify—high-uncertainty cultures like Japan (score 92) craving explicit orders (Hofstede et al. 2010)—culminating in U.S. ISAF micromanagement overriding allies amid Afghanistan’s cultural-social fog, trust evaporating as Taliban misinformation widened gaps (Alwyn-Foster 2005; HDCDTF 201). Consequently, John Boyd’s core principle of organic harmony shatters, leaving the coalition’s internal friction to rival the external chaos imposed by the enemy.

Looking ahead to 2026, the risks posed by hybrid threats from State and non-State actors will embed persistent AI-driven fog into NATO’s operational calculus, positioning static headquarters as glaring HVT in contested multi-domain environments. Russia’s post-Gerasimov adaptations have already demonstrated this in Ukraine, where AI-orchestrated cyber intrusions sustained information denial from 2022 through 2026, fragmenting coalition responses despite Javelins and HIMARS influxes (Google Threat Analysis Group 2023). China’s PLA doctrine, emphasizing “system destruction warfare,” integrates cognitive AI for multi-domain saturation: swarms spoofing Link-16 while hypersonics eye static JTF HQ as first kills (Kania 2019). C2COE’s forward-looking studies paint a dire picture—centralized facilities, relics of layered Cold War defense, concentrate emissions and personnel, inviting anti-access/area-denial (A2/AD) decapitation (C2COE 2025). Gardner illuminates technology’s profound paradox: AI offers tantalizing fog mitigation through superhuman pattern recognition and data synthesis, yet engenders nonlinear instabilities wherein autonomous systems exceed human oversight capacities, precipitating unintended escalations or fratricidal engagements (Gardner 2024). Absent dispersed C2 architectures, NATO confronts pre-kinetic operational paralysis—adversarial cyber deceptions immobilizing decision loops while inbound threats multiply unchecked.

Fog tears through C2 functions relentlessly—here’s how it hits each one. Intelligence gets swamped with false positives from deepfakes and spoofed sensors; Ukraine analysts wasted time on 40% ghost

convoys, missing real dangers (Google TAG 2023). Logistics breaks under cyber attacks like NotPetya (2017), which froze Ukrainian rail and previewed multinational mission network vulnerabilities—IoT hacks fake inventories, starving troops mid-fight (Rid 2013). Targeting and fires warp completely: disinformation twists MoE/MoP, stretching F2T2EA cycles 2-3x as decoys multiply (C2COE 2019a). Maneuver stalls in SA black holes, friendly fire surging like digital Normandy. Sustainment snaps from supply hacks, bogus casualty reports. Force protection misses insiders, cultural fog hiding threats. Even OPSEC/CBRN crumbles in contested EMS. Watts explains why: small cyber jabs cascade into theater collapse (Watts 1996).

On the strategic level, fog steals NATO’s tempo advantage, where faster and better decision making wins wars. Hybrids seize initiative in the chaos they create. Coalitions fracture as Hofstede gaps (such as power distance clashes) intersect with adversarial psychological operations, severely degrading trust. This dynamic was visibly demonstrated within the multi-national ISAF mission in Afghanistan, where the inability to counter sophisticated Taliban propaganda regarding civilian casualties and operational friction deeply eroded mutual trust, exacerbated national ‘caveats,’ and delayed joint execution ** (Larson et al. 2012; Auerwald and Saideman 2014). Peer fights hand A2/AD victories; static HQ decapitation kills air/sea control (Brose 2020). Adversaries set the rhythm while NATO reacts from fogged bunkers. AJP-01 talks mission command but lacks structures—time for friction-proof evolution.

Ukraine proves hybrid fog outlasts tech wonders, grinding down will through distrust that drags into 2026 stalemates. Russia’s 2022 opening salvos—Viasat cyber blackout plus Telegram deepfakes faking Ukrainian routs—shattered SA despite NATO intel flood (Google TAG 2023). HIMARS inflicted pain, but fog’s tricks (spoofed convoys, “ghost army” psyops) kept the grind going without breakthroughs. Donbas 2026 deadlocks show tech parity (Bayraktars vs. Lancets) neutralized by cyber-social friction; OODA bogs in bad intel, allies split by ROE fights and “Zelenskyy corruption” lies. C2COE counts 25-30% decision delays from fog, trust halving coalitions (C2COE 2025).

It mirrors Prussia at Jena (1806): paper superiority—numbers, drills—collapsed under fog-amplified friction. Rigid chains froze as Napoleon’s tempo struck; subordinates waited while corps pounced, just like NATO’s cautious teams today. Static JTFs invite the same fate—tech-heavy forces crushed by fog-thriving hybrids. AJP-01 ideals ring hollow without structures; dispersed AI-boosted C2 is the fix, or history replays.

4. Mitigating Clausewitzian Fog in Modern C2

Taming Clausewitzian fog in today’s multi-domain operations calls for straightforward, battle-tested fixes that turn friction from enemy roadblock into something commanders can actually work with. Old-school centralized C2 crumbles under hybrid attacks; the answer lies in smart, resilient setups—spread-out headquarters, AI helpers, and 3D visualizations—that match C2COE roadmaps and NATO’s FMN upgrades. These blend human judgment with cutting-edge tech, echoing Boyd’s idea of team harmony and AJP-01’s mission command, to get decision cycles moving again despite the fog.

Dispersed headquarters stand out as the foundation of any fog-fighting strategy, taking the NATO Command and Control Centre of Excellence’s “Future Command Post” concepts and turning them into real-world mobile teams linked by tough FMN-based networks. Forget massive static Joint Task Force (JTF) bases; they are high-value targets for hypersonic missiles or cyber strikes. Instead, forces break into agile pods for operations, intelligence, and fires that run independently. While the original framework envisioned these nodes staying in sync via the comprehensive integration of Federated Mission Networking (FMN) Spiral 5+ standards (C2COE 2019a), subsequent implementations have faced hurdles. Specifically, since the publication of those 2019 studies, NATO Affiliates have lowered their collective technological ambitions for subsequent Spirals and delayed their deployment timelines.

Despite these nation-level rollout delays, the technical blueprint remains clear: single failure points vanish when automated radio meshes and blockchain-secured connections maintain data flows under

contested electronic environments. These local network architectures follow guidelines from the Federated Mission Networking Implementation Plan (FMNIP). Real-world validation emerged during the Ukraine conflict, where decentralized “spiderweb” command layouts successfully evaded adversarial kinetic strikes (Google TAG 2023). This modern decentralization directly mirrors Helmuth von Moltke’s classic doctrine: establish tactical goals centrally but empower localized execution to eliminate operational friction and communication delays. Simulated assessments predict that these highly mobile frameworks can yield 50% faster decision cycles and cut headquarters risk by 70% (C2COE 2019b). Furthermore, multinational coalitions directly benefit, as structural friction and cultural differences diminish when operational functions are neatly modularized.

Predictive AI cuts through hybrid fog by pooling massive data streams into lakes—cloud hubs blending JISR feeds—and using machine learning to sharpen Endsley’s SA Levels 2-4 (comprehension and projection). Humans drown in IoT/social overload; AI spots anomalies, predicts patterns, and forecasts threats like Gerasimov psyops via sentiment analysis (Endsley 1987). It starts descriptive (Level 2: correlating cyber-social signals), goes consensual (Level 3: suggesting COAs for human check), and monitored (Level 4: auto-tasking sensors, with veto on kills). Recent frameworks by the NATO C2COE on Federated Mission Networking underscore that this symbiosis dramatically accelerates data synthesis, enhancing decision advantage against complex disinformation campaigns without displacing human authority (NATO C2COE 2025). This approach aligns with Army Doctrine Publication (ADP) 6-0, positioning AI as a tool to bolster tactical initiative while establishing strict transparency and traceability to mitigate “black-box” algorithmic risks.

VAR/AR transforms info output, swapping flat screens for 3D JCOP that makes battlespace grasp instant and intuitive. Picture cyber threats glowing red or social sentiment as heatmaps overlaid on real maps—slashing overload confusion (C2COE 2019a). HoloLens-style AR holograms battle on tables; VR sims test COAs in fog-like chaos, turbo-charging Endsley Levels: spot threats (Level 1) in a flash, project ahead with live sims. Dispersed coalition teams collaborate in shared

virtual rooms, dodging Hofstede cultural snags via visuals over wordy slides. C2COE tests clock 30% faster sensemaking; Ukraine VR drills sharpened targeting (Microsoft 2018). While current FMN implementations still struggle with the bandwidth required for real-time 3D data rendering, evolving standards aim to hook it all together, team-piercing the fog.

Mission command puts humans back in charge through fog, NATO's go-to doctrine tweaking German Auftragstaktik for MDO with smart trust and bendy ROE. Rigid top-down orders flop; say the goal, let teams figure the how. AJP-01 (Edition F) now formalizes this by elevating mission command to the Alliance's overarching philosophy, mandating centralized command and decentralized execution to spark localized initiative (NATO 2022). ADP 6-0 agrees: bosses set "what," not "how," building that fingertip feel (U.S. Army 2022). Bandy ROE swaps stiff rules for on-the-fly tweaks, like Ukraine's fire teams outrunning Russia's stiff chain (Ochs 2020). Trust seals it—bosses take risks, teams grab chances; C2COE pushes pre-fog team-ups via drills. Tweak for cultures (Hofstede): high-context allies get unspoken cues easy. Operational models demonstrate that tightening this cognitive alignment accelerates Boyd's OODA loop synchronization under dense friction (Boyd 2018).

Resilient networks anchor the base, with ad-hoc meshes and blockchain securing cyber-social feeds amid the FMN Spiral 5 rollout. However, because its original, advanced capabilities have been progressively delayed and pushed to subsequent Spirals, current implementations remain limited to foundational services. Jammed EMS calls for self-healing Mobile Ad-hoc Networks (MANET)—tactical radios auto-reconfiguring post-disruption—keeping dispersed nodes synced without central chokepoints (FMNIP 2023). Blockchain locks provenance: tamper-proof ledgers timestamp data, blocking spoofs/deepfakes; Twitter sentiment gets trusted via crypto hashes (Rouse 2018). Spiral 5 fuses domains (Link-16 to 5G), RAdAC balancing risk/speed. Ukraine's Starlink meshes dodged jams, proving it works (Google TAG 2023). To formalize these advantages, current NATO C2 modernization frameworks aim to project up to a 60% increase in network availability over legacy systems under contested conditions, mitigating the severe operational swings of degraded electronic environments.

To reduce operational friction in multinational environments, command and control systems must integrate intercultural training and VUCA exercises. Cultural dimensions—such as power distance and uncertainty avoidance—frequently generate misunderstandings that worsen the fog of war. In NATO's Steadfast exercises, rotating command roles addresses this by forcing different cultures to interact directly, such as balancing U.S. individualistic approaches with German communication styles. Data from the HDCOTF (2021) shows that these cultural rotations slashed communication errors by 35%. Based on these metrics, the C2COE (2019b) now recommends pre-deployment cultural modules to build implicit trust and boost coalition situational awareness before operations begin.

AI-powered war-gaming simulates fog for COA vetting, accelerating planning sans real risk. Reinforcement learning generates adversarial scenarios—cyber-social hybrids testing dispersed HQ—projecting outcomes probabilistically (Stone et al. 2016). Direct JCOP-to-model feeds enable real-time iteration; Endsley Level 4 projection shines, commanders stress-testing ROE flex. C2COE trials cut planning 50%, fog internalized as training artifact (C2COE 2019a). Ukraine's AI sims honed drone defenses, paradigm for NATO.

Cross-cueing sensors across domains harness IoT proliferation with edge AI processing to pierce fog at source, fusing disparate feeds into coherent SA. Multi-domain IoT—drones eyeing kinetic, wearables tracking human/social, cyber probes scanning networks—generates exponential data; edge AI (on-device ML) filters locally, slashing latency vs. cloud reliance. Cross-cueing algorithms query orthogonally: UAV visuals cue cyber scans for spoof detection, social sentiment triggers kinetic retasking (C2COE 2019a). Endsley comprehension surges; Ukraine's Bayraktar-cyber loops exemplify, reducing false positives 45% (Google TAG 2023). While the FMN Spiral 5 rollout enables the conceptual framework for this seamless handoff, its advanced cross-domain features have been delayed to future Spirals, forcing current edge processing to rely on localized networks to counter jamming via AI-prioritized resilient paths. C2COE envisions 70% SA uplift, nonlinearity preempted via predictive fusion (C2COE 2019b).

A phased implementation roadmap ensures progressive feasibility and doctrinal alignment. Phase 1 (2026-2028) establishes FMN Spiral 5 infrastructure, with FMN Affiliates—NATO nations and partner organizations—investing individually in capability development, implementation, and validation on a national-account basis, synchronized only insofar as achieving the Spiral’s interoperability requirements for the portion of forces each nation assigns to NATO within a given timeframe (NATO ACT 2024; AJP-6 2023). This affiliate-driven model means Spiral adoption does not automatically extend across a nation’s entire force structure or all services, but rather targets the specific force elements designated for NATO-led operations. Phase 2 (2029-2032) incorporates AI capabilities—data lakes, edge machine learning, and Endsley Levels 2-3 automation—enhancing situational awareness and targeting precision within these federated, nationally-resourced networks. Phase 3 (2033-2035) achieves fully dispersed HQs with mobile nodes, Level 4 monitored AI, and immersive VAR ecosystems, realizing the C2COE 2035 command post horizon (C2COE 2019a). Iterative validation occurs through NATO exercises like Steadfast and CWIX interoperability testing; NATO DPP guidance aligns national capability targets with Alliance-wide Spiral milestones, though actual investment and implementation timelines remain sovereign national decisions.

These palliatives stand essential against Section 3’s fog ravages—from decapitation blackouts to hybrid-induced decision paralysis. Dispersed HQs and FMN/blockchain nets render cyber/EMS strikes futile, as Ukraine’s webs endured relentless jamming (C2COE 2019a). Predictive AI/ML outflanks social-cyber tsunamis, restoring foresight amid overload (Endsley 1987); VAR/AR forges 3D intuitive grasp, dissolving cross-domain haze even in scattered cells (C2COE 2019b).

Mission command reignites human fire via trust-forged ROE and fingertip instinct, alchemizing desinformation chaos into counterstrikes; Hofstede-grounded drills knit multinational sinews, erasing cultural fissures (AJP-01; Ochs 2020). Phased ascent—FMN bedrock, AI infusion, 2035 dispersed maturity—tames nonlinearity, channeling Ukraine precedents into NATO dominance (Google TAG 2023). Without them, fog cascades unchecked;

with them, consequences invert—vulnerability to velocity, paralysis to preeminence.

5. Conclusion

“Everything in war is very simple, but the simplest thing is difficult. The difficulties accumulate and end by producing a kind of friction that is inconceivable unless one has experienced war.” (On War, 1832, p. 119) Clausewitz’s enduring insight—that friction transforms simplicity into impossibility—finds its most acute expression in multi-domain operations, where cyber and social domains weaponize uncertainty itself. Yet this article demonstrates that friction’s intensification paradoxically validates Prussian solutions: decentralization, not technological omniscience, offers NATO’s path through hybrid fog. The palliatives explored converge into a synergistic architecture that transcends their individual merits, forging what might be termed adaptive C2 resilience—a system where dispersed nodes, AI cognition, immersive visualization, and mission command philosophy interlock to convert MDO’s chaos from existential threat into manageable operational friction.

The synthesis emerges most powerfully at the seams between technologies and doctrine. Dispersed headquarters eliminate decapitation vulnerabilities, but their true force multiplies when paired with predictive AI: autonomous nodes sense-make locally via edge ML, while an FMN-based network weaves them into a coalition organism designed to self-heal under jamming. VAR/AR’s 3D battlespace doesn’t merely visualize—it democratizes SA across Hofstede’s cultural chasms, letting high-context Germanic and low-context Anglo allies grasp shared intent intuitively, bypassing the linguistic friction that paralyzed ISAF. Mission command, revitalized through these tools, transforms from aspirational doctrine into operational reality: subordinates armed with Level 3-4 AI projections and flexible ROE execute Moltke’s vision at cyber speed, outpacing adversaries trapped in centralized hierarchies. The whole exceeds its parts—fog persists, but NATO’s C2 is designed to bend with it rather than break.

Ukraine’s 2022-2026 crucible offers early empirical signals. Russian hybrid saturation—Viasat



Image 2: Prussian cavalry charge during the Battle of Auerstedt (14 October 1806), painting by Georges Hyon (19th century). Part of the twin Battles of Jena-Auerstedt in the Napoleonic Wars, depicting Prussian forces against Napoleon's Grande Armée.

blackouts, deepfake floods—should have shattered Ukrainian cohesion per classical friction theory. Instead, Starlink meshes, AI sentiment filters, and decentralized fire authority sustained resistance, not through fog's elimination but through friction equalization: Russia's cyber advantages met Ukrainian adaptive grit, stalemating rather than steamrolling. NATO's phased 2035 trajectory aims to scale this lesson—FMN Spiral 5 networking (technology, processes, and trained people) as the foundation, AI-infused data lakes by 2032, full dispersed ecosystems completing Auftragstaktik's digital evolution—though this networking capability is not yet fully fielded across the Alliance and remains a work in progress. The payoff, once realized, is not omniscience but decisional primacy: coalitions that sense faster, decide clearer, and act bolder than fog-vulnerable rivals, turning Gerasimov's tactics into self-defeating noise. Yet humility tempers optimism. Friction, as Clausewitz warned, is ineradicable—a structural feature of war's nonlinearity, not a bug to debug. AI ethics loom: Level 4 autonomy risks unintended escalation, black-box decisions eroding trust faster than algorithms build it (Gardner 2024). Technological

dependency births fragility; adversaries evolving quantum decryption or adversarial ML could collapse blockchain/AI pillars overnight. Cultural friction's stubbornness defies even VAR's charm—Hofstede's dimensions root in centuries, resistant to drills or holograms. The 2035 vision demands perpetual adaptation, human-AI symbiosis where commanders wield Fingerspitzengefühl to override, not defer to, machines. NATO's triumph lies not in conquering fog but in mastering its navigation—agile enough to exploit friction gaps, resilient enough to endure its persistence, wise enough to keep humans commanding the storm.

Clausewitz closes *On War* with a truth this analysis vindicates: "In war more than anywhere else, things do not turn out as we expect" (1832, p. 85). Multi-domain friction ensures his words resonate louder in 2026 than in 1832. But dispersed C2, fortified by AI and trust, offers NATO what Prussia found at Königgrätz⁹—not certainty, but calibrated courage to act amid irreducible uncertainty. The fog will never lift. The question is whether we design C2 to freeze within it or dance.

9 At Königgrätz (1866), Prussia decisively defeated Austria using the decentralized command culture and independent-corps maneuver doctrine developed in direct response to its catastrophic 1806 defeat at Jena-Auerstedt. Scharnhorst and Gneisenau's post-Jena reforms—institutionalizing Auftragstaktik and training officers for independent judgment—had matured over six decades into a demonstrable battlefield advantage, proving that structural adaptation to friction, not its elimination, secures victory (Shamir 2011, pp. 33-40).

References

- Arquilla, J. (2021). *Bitskrieg: The new challenge of cyberwarfare*. Polity Press.
- Auerswald, David P., and Stephen M. Saideman. *NATO in Afghanistan: Fighting Together, Fighting Alone*. Princeton University Press, 2014. <https://doi.org/10.2307/j.ctt4cgbc2>.
- Bose, R. (2020). *The kill chain: Defending America in the future of high-tech warfare*. Hachette Books.
- Borchert, H., Schütz, T., & Verbovsky, J. (2021). *Beware the hype: What military conflicts in Ukraine, Syria, Libya, and Nagorno-Karabakh don't tell us about the future of war* (DAIO Study 21/01). Defense AI Observatory, Helmut Schmidt University.
- Boyd, J. R. (2018). *A discourse on winning and losing* (G. T. Hammond, Ed.). Air University Press.
- Britt, A. B. (2003). *The wars of Napoleon*. Square One Publishers.
- Cebrowski, A. K., & Garstka, J. H. (1998). *Network-centric warfare: Its origin and future*. Naval Institute Proceedings, 124(1), 28–35.
- Clausewitz, C. von (1989). *On war* (M. Howard & P. Paret, Eds. & Trans.). Princeton University Press. (Original work published 1832)
- Elron, E., Shamir, B., & Ben-Ari, E. (1999). *Why don't they fight each other? Cultural diversity and operational unity in multinational forces*. Armed Forces & Society, 26(1), 73–97.
- Endsley, M. R. (1987). *The application of human factors to the development of expert systems for advanced cockpits*. Proceedings of the Human Factors Society 31st Annual Meeting, 1388–1392.
- Eric V. Larson et al., *Understanding and Influencing Publics: Psychological Operations and Communication Strategies in Information Operations* (Santa Monica, CA: RAND Corporation, 2012).
- Gardner, N. (2024). *Clausewitzian friction and twenty-first-century war: The paradox of technology*. Naval War College Review, 77(1), Article 5.
- Gerasimov, V. (2013, February 27). *The value of science in prediction*. Voyenno-Promyshlennyy Kurier.
- Google Threat Analysis Group. (2023, February 15). *Fog of war: How the Ukraine conflict transformed the cyber threat landscape*.
- Hofstede, G., Hofstede, G. J., & Minkov, M. (2010). *Cultures and organizations: Software of the mind* (3rd ed.). McGraw-Hill.
- Hughes, D. J. (Ed.). (1993). *Moltke on the art of war: Selected writings*. Presidio Press.
- Kania, E. (2019). *Artificial intelligence in future Chinese command decision making*. In N. D. Wright (Ed.), *Artificial intelligence, China, Russia, and the global order* (pp. 155–168). Air University Press.
- Lindsay, J. R. (2020). *Information technology and military power*. Cornell University Press.
- NATO. (2015). *AJP-3.14: Allied joint doctrine for force protection (Edition A, Version 1)*. NATO Standardisation Office.
- NATO. (2017). *AJP-01: Allied joint doctrine (Edition E, Version 1)*. NATO Standardisation Office.
- NATO. (2022). *Allied Joint Doctrine (AJP-01, Edition F, Version 1)*. NATO Standardization Office.
- NATO C2 Centre of Excellence. (2019a). *The future of the command post: Part 1*.
- NATO C2 Centre of Excellence. (2019b). *The future of the command post: Part 2*.
- NATO C2 Centre of Excellence (2020), *Whitepaper NATO C2COE C2 Demonstrator 1.0*.
- NATO C2 Centre of Excellence. (2025). *From concept to capability: NATO's C2 of multi-domain operations—History, evolution, and challenges*.

NATO C2 Centre of Excellence. (2026). *From concept to capability: NATO's C2 of multi-domain operations—History, evolution, and challenges*.

Ochs, R. (2020). *Mission command: An exploration of its issues in multinational units*. *Journal of Military History and Defence Studies*, 1(2), 75–94.

Rid, T. (2013). *Cyber war will not take place*. Hurst.

Rouse, M. (2018). *Blockchain*. In *Essential guide: Special report—AI apps come of age*. TechTarget.

Scharre, P. (2018). *Army of none: Autonomous weapons and the future of war*. W. W. Norton.

Shamir, E. (2011). *Transforming command: The pursuit of mission command in the U.S., British, and Israeli armies*. Stanford University Press.

Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenian, A., Shah, J., Tambe, M., & Teller, A. (2016). *Artificial intelligence and life in 2030: One hundred year study on artificial intelligence* (Report of the 2015 Study Panel). Stanford University.

U.S. Army. (2022). *ADP 6-0: Mission command: Command and control of Army forces*.

Watts, B. D. (1996). *Clausewitzian friction and future war (McNair Paper 52, Revised Edition)*. Institute for National Strategic Studies, National Defense University.

Author biography

Antonio Legaz Sanz

Antonio Legaz Sanz is a defence and security researcher currently working at Indra Sistemas, where he focuses on command and control, NATO structures and European defence policy. In his role he contributes to defence-technology and C2-related projects, combining operational insights with strategic analysis. His work has led to research and policy writing on defence organisation, military planning, and command-and-control systems in multinational environments. He has published in academic and policy-oriented outlets on defence and security topics, including contributions to journals and platforms specialised in international relations and defence studies.

Authors Note

The authors have not received any funding for writing this paper. There are no conflicts of interest.

Acknowledgements

Funded by the European Union under the European Defence Fund (GA no. 101121403 - NEWSROOM). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the granting authority can be held responsible for them.

The author declares no financial, professional or personal conflicts of interest that could influence the objectivity of this work.



RUSSIAN COMMAND AND CONTROL IN INFORMATION WARFARE AND REFLEXIVE INFLUENCE

Implications for NATO and Counter-Reflexive
Control

Dr. Ivana Stradner, Rear Admiral (Ret.) Mark Montgomery

The authors would like to thank Stephen Zhukov for providing valuable research assistance for this article.



This paper reviews Russia's approach to C2 in the information environment. This approach relies on the Soviet technique of "reflexive control," in which Russian operatives model an adversary's decision-making structure and plant subversive information that leads the adversary to make seemingly logical decisions that benefit Moscow. Using reflexive control, the Kremlin aims to subvert NATO's decision-making infrastructure by inciting disunity and disfunction in the alliance. Moscow uses nuclear signaling in the framework of reflexive control to try to fracture the alliance's resolve against Russia. The paper concludes with policy recommendations for NATO's consideration: the introduction of a Perception Management Warfare Assessment Unit and the implementation of a tactic of Counter-Reflexive Control.

Keywords:

Cognitive warfare, nuclear signaling, perception management, reflexive control, Russia

Introduction

In the modern-day digital environment, those who can effectively leverage the digital space wield the most power in the global system. Learning how to conduct not only cyber-attacks, but information operations (IO), psychological operations, and social media campaigns afford beneficiaries significant power over the cognitive functions of their adversaries. Russia has not only learned these tactics but has institutionalized them into its Command and Control (C2) infrastructure (NATO C2OE, 2026). Within this framework, C2 assumes a fundamentally different meaning from how it is understood in the West. While NATO typically views C2 as a purely technical means of executing strategic and operational goals, the Kremlin uses C2 to target not only NATO's technical architecture, but also its cognitive resiliency, seeking to degrade NATO's ability to plan and execute missions from its very foundations.

This situation calls for an immediate reevaluation of NATO IOs, and the creation of new measures that can mitigate the effects of a vulnerable information environment. This paper argues for two key recommendations: The creation of the "Perception Management Warfare Assessment Unit," (PMWAU) and the implementation of a tactic of "counter-reflexive control" (CRC). These initiatives will address the exploitation of NATO's operations while helping the alliance develop robust protocols to avoid exposing its decision-making systems to Russia and thus bypass the carefully laid traps the Kremlin has set for it. In reducing these vulnerabilities, NATO will build a more transparent

and effective information system and level the playing field in the information confrontation with Russia.

I Russia's Information Warfare C2 Model

Russia's approach to cyber and information warfare (IW) differs fundamentally from the Western model. The United States, and much of NATO, conceptualizes cyber space as a military domain analogous to air, sea, land, or space, focused on network defense, operational disruption, and domain superiority (NATO, 2024). In Russia, by contrast, cyber is not primarily about defending networks or even technical exploitation alone. It is part of a broader political warfare infrastructure designed to shape perception, influence decision-making, and manage cognitive environments (Creery, 2018). Instead of the term "cyber security," (кибербезопасность) Russia's 2021 National Security Strategy uses the term "information security" (информационная безопасность) (President of the Russian Federation, 2021, p. 19). While this may appear as a semantic difference, it is actually intentional and consequential in the language of the Kremlin. Hence, in this paper we use Russia's term "information security." The Kremlin's prioritization of IW is exemplified in the 2021 National Security Strategy, which states "...information security is achieved through... the development of the forces and means of information warfare." (President of the Russian Federation, 2021, p. 21)

Russia has holistically integrated IW into its C2 system. Per the Ministry of Defense's 2011 Conceptual Views Regarding the Activities of

the Armed Forces of the Russian Federation in the Information Space, “The leadership and the command staff of all levels directly participate in the organization of the activity in the information space during the peacetime and the wartime, in the preparation and in the course of the operations (military engagement).” (Ministry of Defense of the Russian Federation, 2011, p. 8) Influencing the information space is not isolated to wartime activities for the Kremlin; its intelligence agencies work tirelessly to dominate the cyber sphere, using C2 as its framework.

Moscow’s IW C2 effort centers around targeting three areas of cognitive influence. First, Moscow has built an offensive response to adversarial information and psychological operations. Its intelligence agencies operate as if an attack is imminent, developing, planning, and implementing preventative offensive operations and striking preemptively (Lozovsky et al., 2019). Second, Moscow’s C2 works to constantly weaken its adversaries’ information capabilities. This includes anything from outright sabotage, targeting power plants and electrical infrastructure abroad to sowing doubts in existing information capabilities. Finally, Moscow conducts offensive IOs against adversary troops and populations to create an information environment that is beneficial to Kremlin policies.

Unlike NATO where a centralized cyber or IW command may oversee operations, Russian C2 is distributed and reflexive (Vasara, 2020, p. 71). The system is organized as a hierarchy: priorities, goals, and targets are defined at the highest level of state authority, and execution is segregated by responsibility and dispatched to the corresponding agency. By consolidating priorities at the top, the regime can ensure that all actions are aligned with the leadership’s strategic objectives, regardless of which department executes them (Vasara, 2020, p. 71). Sitting directly underneath the executive leadership, Russia’s Security Council combines intelligence assessments across agencies,

integrating leadership from the Federal Security Service (FSB), the Main Directorate of the General Staff (GU) and the Foreign Intelligence Service (SVR), alongside the Kremlin-sponsored media ecosystem (President of the Russian Federation, n.d.).

Federal Security Service (FSB)

The FSB primarily focuses on shaping and protecting the domestic information environment while also contributing to offensive IOs. To shield Russian decision-making from external manipulation, the FSB secures the domestic cognitive environment combining internal control with counter-reflexive shielding. Currently, the FSB is responsible for the implementation of large-scale internet suppression and censorship in Russia (Kolmychenko, 2026)¹. By monitoring public sentiment, and limiting foreign informational influence, the FSB preserves Russia’s cognitive resilience (Government of the Russian Federation, 1995).

The Main Directorate of the General Staff (GU)

The GU operationalizes Russia’s C2 doctrine most directly, bearing the primary responsibility for modelling adversary reactions, testing escalation thresholds, and adapting tactics based on the observed responses from the target. Operationally, the GU operates the most direct equivalent of an embedded cognitive assessment capability within military planning, with specialized Units such as 29155, 26165 or 74455 working to integrate cyber-attacks and psychological operations into cohesive campaigns (UK Foreign, Commonwealth, and Development Office, 2025).

Foreign Intelligence Service (SVR)

The SVR primarily conducts cyber espionage operations to collect intelligence in support of Russian strategic objectives.² However, the infrastructure it develops and the data it steals

¹ <http://www.kremlin.ru/acts/assignments/orders/79927> and <https://home.treasury.gov/news/press-releases/jy1572>

² More about Russia’s SVR operations <https://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-View/Article/3931959/nsa-issues-updated-guidance-on-russian-svr-cyber-operations/>. Also, see Russia’s APT 29 <https://www.cyber.nj.gov/threat-landscape/nation-state-threat-analysis-reports/russia-cyber-threat-operations/russia-apt29>

are also used to enable cyber-enabled influence and IOs (Bowen, 2022). The U.S. government identified the SVR as one of the two Russian cyber units responsible for cyber campaign-related activities during the 2016 U.S. presidential election. Also, the SVR targets adversary political elites, think tanks, academic institutions, and business networks (Panyi, 2026). Rather than immediate destabilization, the SVR seeks gradual penetration of decision-making ecosystems, seeking to alter critical decision-making infrastructure. This approach reflects Russian doctrine's view of IW as a sustained competition over influence rather than a series of isolated events.

Kremlin-Sponsored Media Ecosystem

Russia's media ecosystem plays a vital role in its IW environment, functioning as both an amplifier and a sensor. State-owned media outlets such as Sputnik and RT (formerly Russia Today) broadcast Russian narratives across the world in a variety of languages and media formats (U.S. Department of State, 2022). They aim not only to propagate the Russian government's version of reality, but more fundamentally to foment social division and "to undermine the official version of events, even the very idea that there is a true version of events and foster a kind of policy paralysis." (MacFarquhar, 2016). Kremlin-sponsored bloggers on social media sites like Telegram allow further dissemination of narratives, often disguising their true origins (Global Disinformation Unit, BBC World Service, 2025). This architecture constitutes a distributed but coherent command-and-control system for IW. It is centralized in intent, decentralized in execution, and continuous in operations, whether peacetime or conflict (Bondar, 2026). This structure creates the ideal circumstance for IOs, in which leadership can quickly distribute instructions to initiate new campaigns or adjust those in progress. At the same time, decentralized execution makes it difficult to attribute operations back to the Kremlin, while acting as a force multiplier as IOs are carried out from multiple outlets and directions.

II Reflexive Control as the Core Logic

Russian information-warfare command and control can best be understood as the institutionalization of reflexive control (Snegovaya, 2015). Reflexive control (рефлексивное управление) has been central to Russian strategy since the Soviet era. Developed by Vladimir Lefebvre in the 1960s, reflexive control seeks to transmit carefully prepared information to an adversary so that the adversary voluntarily makes decisions favorable to the initiator wrongly believing those decisions are autonomous. It aims to structure the opponent's perception so that certain choices appear logical and necessary, while actually being predetermined and subversive (Lefebvre & Smolyan, 1968).

Lefebvre's framework seeks to mathematically model decision-making by mapping a target's layered perceptions: not only how the target perceives himself and relevant others, but how he believes others perceive him and each other in turn (Lefebvre & Smolyan, 1968, p. 20). Understanding the system, cultural nuances, and global environment that decisions are embedded in is key to shaping an adversary's decision-making process. It models how injecting specific information into the system at select times would shape decision making not just at the crux, but at every layer of the decision-making process (Snegovaya, 2015, p. 10). Beyond influencing a single decision, the reflexive game theory model alters how the adversary analyzes every input from the system through the desired lens, unknowingly forfeiting control to Russia.

Russian IW therefore does not primarily aim to "win arguments." It seeks to overload cognition, exploit bias and trigger desired adversary's responses. The adversary is treated as a system with different inputs so Russia can design stimuli that generate desired outputs (Merriam, 2023).

In practice, this operational loop of reflexive control begins with modeling the adversary. With this model, Russia can construct detailed assessments of political culture, ideological fault lines, media behavior or legal constraints (Thomas, 2004). The adversary modeling phase is continuous and feeds strategic planning (Merriam, 2023). Despite their varied responsibilities, each agency's approach

remains aligned with the goals of the Presidential administration, a key feature of the system's organization (Galeotti, 2020).

A core goal of Russia's IW C2 model is to disrupt or disorganize the functionality of its adversary's command-and-control structure. Timothy Thomas wrote in detail how Russia has been using reflexive control to disrupt adversary's C2 (Thomas, 2019). Major General Yuriy Lastochkin, head of the Defense Ministry's electronic warfare force (REB), established "disorganizing adversary C2" as the guiding principle of the REB's combat employment strategy.

III NATO as a Target of Reflexive Influence

Russia's strategic approach to NATO is grounded in the logic of reflexive control. Rather than prioritizing direct kinetic confrontation with NATO, Moscow seeks to influence how the Alliance interprets signals and how it deliberates under pressure. The central target is not territory or force posture alone, but the Alliance's decision-making process itself (Thomas, 2016).

Russia models NATO as a multi-member, consensus-driven alliance characterized by diverse political cultures and varying threat perceptions (NATO, n.d.). For example, Slovakia and Hungary have very different interests from France or the United Kingdom. Countries on the alliance's Eastern flank are frequently more "hawkish" on Russia than other member states (Kamiński & Śliwa, 2023). Decision-making requires consultation and unanimity, and Moscow has sought to exploit these differences. Russia's objective is not necessarily to fracture NATO outright, but to maximize indecision or produce asymmetric responses (Galeotti, 2025). One state may call for escalation, another for caution that would slow action and project fragmentation. Even temporary misalignment can produce strategic leverage. Russia also conducts such attacks to test NATO's reaction and build a strategy for future endeavors based on any discovered weaknesses.

For instance, Russia might conduct a drone incursion in Poland to cause chaos in NATO decision-making and strain the viability of their key

Article 5 mutual defense clause (NATO, n.d.). A drone incursion would likely warrant an immediate air defense alert and military aircraft response, and likely an emergency session to consult NATO and determine an appropriate response in the following days. If this process leads to indecision and Article 5 is not invoked, Russia could point to this as another example of NATO's weakness and inability to protect its own, sowing potential mistrust amongst its more vulnerable members.

This exact scenario played out in September 2025, when a Russian drone invaded Polish airspace. Polish Foreign Minister Radosław Sikorski noted that the Kremlin "wanted to test the readiness of NATO allies" (O'Connor, 2025). Russia likely modeled the severity of the immediate military response, determining how quickly members could convene and if the incursion could be considered a violation of Article 5. In the ensuing months, Russia's operational loop of reflexive control has likely taken in the inputs of how the incursion impacted NATO's decision-making and fed these into models to understand and evaluate which desired outcomes were achieved and what operational changes might need to be made for future attacks.

IV Nuclear Information Operations as Russia's Cognitive Weapon

Nuclear weapons are a mainstay of Russian foreign policy, and the Kremlin frequently defers to nuclear signaling when it faces international opposition. Putin has repeatedly threatened to use "all available means" to defend Russia and has leveraged its arsenal of nuclear capable missiles to discourage NATO support for Ukraine (The Associated Press, 2022). Indeed, when losing on the battlefield in Ukraine in September 2024, Putin resorted to saber rattling, conducting an exercise with possible nuclear weapons on the border of Ukraine (Schnieder, 2024). Just as recently as May 2026, Russia organized a three-day nuclear exercise in Belarus, where they first began repositioned tactical nuclear forces in 2022 (Arms Control Association, 2024; The Associated Press, 2026). Nuclear reflexive control is not primarily about actual nuclear weapons use. It is about shaping NATO behavior through manipulated perception.

Nuclear signaling serves a broader strategic purpose than purely deterring NATO: following Russia's principle of "escalate to de-escalate," (Haney, 2016) the Kremlin leverages nuclear signaling to both shape and fracture NATO responses, forcing NATO members to rapidly evaluate the credibility of the threat and unanimously decide how to "escalate or de-escalate". Unanimity is considerably more difficult when faced with a direct threat, often causing fragmentation among the more risk-averse members and the more "volatile" members over an appropriate response. To avoid direct confrontation or unnecessary loss of life, Russia projects that NATO will choose de-escalation or diplomacy, making nuclear signaling a very effective weapon (Horowitz & Arndt, 2022).

In the context of nuclear IOs, Russia's IW begins with perception targeting. Russia may announce nuclear exercises, highlight nuclear force deployments, adjust readiness postures (even just rhetorically), or allow selective leaks suggesting escalatory intent (Matle & Rácz, 2026). These actions do not necessarily indicate imminent weapons use. Instead, they are designed to elevate perceived risk inside NATO capitals. By increasing uncertainty around intent, Moscow seeks to widen the gap between objective capability and subjective threat perception.

Reflexive nuclear IOs unfold in stages. First, Russia targets NATO's perception of threat by elevating uncertainty around intent. Second, it exploits alliance decision patterns, such as consensus requirements, risk aversion, and domestic constraints. Third, it operates at the level of "perception of perceptions": shaping what NATO believes other NATO members believe (Selhorst, 2016). Finally, it leverages multi-member variability. A frontline state may perceive imminent danger; another may interpret posturing.

One consistent success the Kremlin has achieved through its reflexive information campaigning is its ability to muddy the decision-making process surrounding NATO's relationship to Ukraine. NATO members have long disagreed on the feasibility and benefits of admitting Ukraine into NATO. Because a consensus is difficult to reach, the Kremlin has found multiple avenues of systematic exploitation of

NATO's indecision. By framing NATO as an internally fractured and indecisive body unwilling to take risks to fully support Ukraine, and amplifying narratives of the futility of NATO accession to Ukrainian audiences, the Kremlin is encouraging further decision-making paralysis and infighting (Langelonatišvili, 2016, p. 17-19). A prime example of NATO's fractured response to the conflict in Ukraine was in 2022 and 2023, when Ukraine requested long-range missiles from the United States (Trevelyan & Osborn, 2022). These missiles would afford Ukraine the ability to reach targets deeper inside Russian occupied territory, including supply lines, airbases and rail networks. Despite their clear tactical advantage and Washington's vocal support for the Ukrainian war effort, Kremlin saber rattling proved effective in the short-term; the US repeatedly denied Ukrainian requests in the first years of the war as Putin threatened that NATO involvement would be met with swift escalation. This position was notably backed by Germany and other NATO members, who worried that Russia's "red lines" were a credible threat. Through nuclear signaling, Russia leveraged NATO's varied risk-appetites to achieve its desired end goal of limiting military support for Ukraine.

Reflexive nuclear IOs also operates at a meta-perception level, what might be called "perception of perceptions". Russian planners consider not only what NATO believes Russia is doing, but what NATO believes other NATO members believe. By shaping the information environment through media amplification, Russia can cloud diplomatic engagements with outside noise (false narratives), hampering efforts to clarify state intentions and perceptions (Stradinger, 2024).

"Perception of perceptions" is an important example of Russia's emphasis on C2 disorganization. By targeting NATO's ability to organize and understand its own members positions, while facing the threat of nuclear escalation, Russia stuns the alliance's ability to act quickly and in a unitary fashion. Disorganization is especially pronounced and effective in a multi-member alliance. Different states may interpret identical signals differently.

V “Perception Management Warfare Assessment Unit,” (PMWAU)

Neutralizing reflexive control in the nuclear domain requires protecting decision-making from manipulated perception. To this end, NATO should develop a Perception Management Warfare Assessment Unit (PMWAU). The establishment of a PMWAU within NATO would offer a structured countermeasure to mitigate this vulnerability. The purpose of a PMWAU would not be to dominate narratives or outcompete disinformation in the public sphere. Rather, it would protect Alliance decision-making from adversary-induced distortion.

A PMWAU would perform three core functions. First, it would provide real-time threat assessment, distinguishing actual military or operational risk from perception-driven escalation. Second, it would identify cognitive biases or distortions introduced by Russia’s signaling and manipulating emotional response. Third, it would deliver corrective feedback to decision-makers, ensuring that policy choices are grounded in calibrated assessments rather than manipulated decisions.

Operationally, the integration of a PMWAU into NATO would reduce both overreaction and undue hesitation. By providing perception-corrected analysis, it would narrow inter-member divergence and strengthen synchronized responses. In effect, it would transform NATO from a reflexive target reacting predictably to information shocks into a more resilient decision network capable of absorbing and contextualizing Russia’s cognitive stimuli.

VI Modeling a Perception Management Warfare Assessment Unit response to Kremlin Reflexive Nuclear Information Operations

The first function of a PMWAU in a nuclear signaling scenario would be signal verification. The PMWAU would evaluate reported adversarial deployments, exercises, or readiness changes to determine whether they reflect substantive shifts in capability or posture, or whether they constitute demonstrative signaling. This assessment would integrate satellite imagery, cyber intelligence, national contributions, and doctrinal analysis. By grounding deliberation

in verified indicators, the PMWAU would reduce the likelihood of perception-driven escalation, protecting NATO decision-making from attempted reflexive influence.

The PMWAU’s second function would be bias correction. Nuclear signals are designed to amplify fear. A PMWAU would provide perception-adjusted assessments that distinguish between credible threat indicators and political messaging. By explicitly separating emotional salience from operational likelihood, it would build NATO’s resiliency against cognitive manipulation.

Third, a PMWAU would facilitate coordinated decision-making. Asymmetries in information distribution and skepticism regarding foreign intelligence assessments frequently impede the NATO decision-making process and open gaps our adversaries can exploit.

Finally, a PMWAU would provide feedback and adaptation. For instance, in late 2025, Russia announced it was moving nuclear weapons close to the EU border (Khomenko, 2025). NATO could use the PMWAU to analyze Russian signaling patterns such as recurring drills near NATO borders or synchronized media amplification over time to identify reflexive tactics. Continuous monitoring would enable the Alliance to anticipate and contextualize future signals, reducing their disruptive potential. The aim is to prevent each episode from being treated as unprecedented when it follows an identifiable pattern.

VII Counter-Reflexive Control (CRC): How to put Russia on the defensive

As Russia increasingly employs reflexive control tactics against NATO, alliance disputes and hesitation will become more commonplace. To rectify the situation before a more serious disagreement arises, it is imperative for NATO to develop measures to maintain a united front. To do this, NATO should begin implementing the tactic of Counter-Reflexive Control, aimed at facilitating inter-alliance communication and mitigating the effects of Russian IW tactics. CRC should deny Russia the ability to predict or manipulate Alliance decision-making through cognitive exploitation and

reaction-forcing tactics. Beyond the PMWAU, the use of CRC provides a comprehensive framework for protecting Alliance decision autonomy, one which recognizes that adversaries exploit NATO's transparency, consensus procedures, legal thresholds, and pressure for speed (Merriam, 2023).

CRC tactics also include measures to cut the enemy off from information about one's own decision-making, which would vastly degrade adversarial reflexive cognitive capabilities. For Western countries which place value on transparency, this can be a challenge. However, measures can be taken which limit Russian access to feedback information, within the limits set by the legal requirements for transparency (Vasara, 2023).

Finally, CRC techniques offer a means to better understand and categorize threats, which can reduce the time pressure that frequently pushes actors towards making decisions that are either more predictable or not fully vetted for potential adversarial reflexive influence. As discussed above, nuclear threats particularly take advantage of time stress to put pressure on decision-making.

NATO should measure the success of CRC tactics not by narrative dominance but by preserved cohesion, avoided premature escalation, maintained autonomy, and increased adversary uncertainty. Experts point to the publication of Russian troop movements and attack plans in early 2022 as a particularly effective example of countering reflexive control techniques, one which disrupted Kremlin propaganda and paved the way to strengthening and expanding the NATO alliance (Vasara, 2023).

VIII Strategic Implications

The central insight from analyzing Russia's C2 modeling and reflexive influence campaigning against NATO is that decision-making itself is a contested capability. Russia's command and control system institutionalizes reflexive control, integrating intelligence, media, cyber operations, and political direction to shape adversary cognition. Russia's conception of C2, as well as its understanding of hybrid war, persistently place it on the offensive against NATO's cognitive space. NATO's structural

strengths of transparency, consultation, and democratic legitimacy can become vulnerabilities if rendered predictable. Putin's saber rattling has already made it more difficult for NATO countries to maintain a unified response, leveraging Russia's nuclear arsenal as both a kinetic and psychological weapon. Using response reflexive control modelling, the Kremlin clouds threat credibility, exploits vulnerabilities in NATO's consensus-driven system, and alters perceptions.

Ultimately, Russia's IW model seeks to govern adversary reactions rather than defeat adversary forces directly. The information space has risen to prominence as the primary battleground, and hybrid warfare is the new norm for NATO adversaries. Putin's regime is actively using Ukraine as a test ground for its reflexive control tactics, learning how to leverage perceptions and strategic information input to shape real outcomes. In this space, where ideas are weapons, NATO should build cognitive resiliency.

To this end, establishing a PMWAU will be paramount in regulating NATO's information space, assessing threat credibility, evaluating Russian nuclear signaling based on historical precedents, and coordinating alliance decisions to render Russian psychological tactics ineffective. Although this article focuses on reflexive control in the context of nuclear weapons, the concept can be broadly applied to other scenarios in which Russia seeks to influence, disrupt, or manipulate NATO's C2 decision-making processes.

NATO should also invest in Counter-Reflexive Control tactics and techniques to ensure that it retains freedom of action by defending its cognitive domain. The PMWAU and CRC will work together to empower NATO to make decisions without any Kremlin poison seeping in. In a monitored and coordinated information space, Russian strategies that rely on altered perceptions will no longer have the same effect on NATO decision-making, affording the alliance to make unified, informed decisions to achieve strategic goals. If NATO does not invest in these critical resources, it risks losing before the fight even begins.

References

- Associated Press. (2021). *Cornered by war, Putin makes another nuclear threat*. The Hill. <https://thehill.com/homenews/ap/ap-international/ap-cornered-by-war-putin-makes-another-nuclear-threat/>
- Copp, T., & Weisgerber, M. (2022). Poland offers fighter jets for Ukraine, but US rejects swap-and-send idea. *Defense One*. <https://www.defenseone.com/threats/2022/03/poland-offers-fighter-jets-ukraine-will-us-agree-twist/362928/>
- Creery, M. (2018). *Russia strategic understanding of cyber: Not an information war – a war on information*. Georgetown Security Studies Review. <https://georgetownsecuritystudiesreview.org/2018/12/05/russia-strategic-understanding-of-cyber-not-an-information-war-a-war-on-information/>
- Bondar, K. (2026). *How Russia is reshaping command and control for AI-enabled warfare*. Center for Strategic and International Studies. <https://www.csis.org/analysis/how-russia-reshaping-command-and-control-ai-enabled-warfare>
- Bowen, A. (2022). *Russian cyber units*. Congressional Research Service. <https://www.congress.gov/crs-product/IF11718>
- Bowen, A. (2025). *Russia's Foreign Intelligence Services*. Library of Congress. <https://www.congress.gov/crsproduct/IF12865#:~:text=As%20Russia's%20primary%20civilian%20foreign,Russia%20or%20the%20Russian%20government>
- Galeotti, M. (2020). *The Presidential Administration: The Command and Control Nexus of Putin's Russia*. George C. Marshall European Center for Security Studies. <https://www.marshallcenter.org/en/publications/security-insights/presidential-administration-command-and-control-nexus-putins-russia-0>
- Galeotti, M. (2025). *When disinformation meets disruption: Russia's strategy of paralysis*. Фондация за хуманитарни и социални изследвания -- София. <https://www.ceeol.com/search/article-detail?id=1373757>
- Global Disinformation Unit, BBC World Service. (2025). *The Telegram channels spreading pro-Russian propaganda in Poland*. BBC News. <https://www.bbc.com/news/articles/c0eqdd53dd70>
- Government of the Russian Federation. (1995). *Федеральный закон "О федеральной службе безопасности"*. https://www.consultant.ru/document/cons_doc_LAW_6300/
- Haney, C. (2016). *Project on nuclear issues capstone conference*. U.S. Strategic Command. <https://www.stratcom.mil/Media/Speeches/Article/986478/project-on-nuclear-issues-capstone-conference/>
- Horovitz, L. & Arndt, A. (2022). *Russia's catch-all nuclear rhetoric in its war against Ukraine*. Stiftung Wissenschaft und Politik. <https://www.swp-berlin.org/en/publication/russias-catch-all-nuclear-rhetoric-in-its-war-against-ukraine>
- Kamiński, M., & Śliwa, Z. (2023). *Poland's threat assessment: Deepened, not changed*. PRISM. https://www.researchgate.net/publication/371874715_Poland's_Threat_Assessment_Deepened_Not_Changed
- Khomenko, I. (2025). *Russia parks Oreshnik nuclear missiles in Belarus—but aims at the West*. United24 Media. <https://united24media.com/latest-news/russia-parks-oreshnik-nuclear-missiles-in-belarus-but-aims-at-the-west-14603>
- Kolomychenko, M. (2026). *Who is responsible for the demise of the Russian internet?* Carnegie Endowment for International Peace. <https://carnegieendowment.org/russia-eurasia/politika/2026/04/russia-internet-crackdown>
- Lange-Ionatamishvili, E. (2016). *Handbook of Russian information warfare*. NATO Strategic Communications Centre of Excellence. https://stratcomcoe.org/cuploads/pfiles/russian_information_campaign_public_12012016fin.pdf
- Lefebvre V., & Smolyan G. (1968). *Conflict Algebra*. Moscow: Znanie Press.

- Lozovsky, V., Kuzmenko, A., & Kuznetsov, A. (2019). *Информационная безопасность: защита от внутренних и внешних воздействий в киберпространстве*. Voyennaya Mysl. <https://cyberleninka.ru/article/n/informatsionnaya-bezopasnost-zaschita-ot-vnutrennih-i-vneshnih-vozdeystviy-v-kiberprostranstve/viewer>
- NATO Command and Control Centre of Excellence. (2026). *Information operations in Russian command and control*. <https://c2coe.org/wp-content/uploads/Library%20Documents/Articles/20260520%20Information%20Operations%20in%20Russian%20Command%20and%20Control%20Shortread.pdf>
- North Atlantic Treaty Organization. (n.d.). *Collective defence and Article 5*. <https://www.nato.int/en/what-we-do/introduction-to-nato/collective-defence-and-article-5>
- North Atlantic Treaty Organization. (2024). *Cyber defence*. <https://www.nato.int/en/what-we-do/deterrence-and-defence/cyber-defence>
- North Atlantic Treaty Organization. (n.d.). *What is NATO?* <https://www.nato.int/en/what-is-nato>
- MacFarquhar, N. (2016). *A powerful Russian weapon: The spread of false stories*. The New York Times. <https://www.nytimes.com/2016/08/29/world/europe/russia-sweden-disinformation.html>
- Matle, A. & Rácz, A. (2026). *Russia's nuclear signaling in 2026 and the implications for European security*. German Council on Foreign Relations. <https://dgap.org/en/research/publications/russias-nuclear-signaling-2026-and-implications-european-security>
- Merriam, J. (2023). *One move ahead – diagnosing and countering Russian reflexive control*. The Journal of Slavic Military Studies. <https://www.tandfonline.com/doi/full/10.1080/13518046.2023.2201113>
- Ministry of Defense of the Russian Federation (2011). *Концептуальные взгляды на деятельность вооруженных сил Российской Федерации в информационном пространстве*. <https://info.publicintelligence.net/RU-CyberStrategy.pdf>
- O'Connor, T. (2025). *Poland deputy PM seeks 'toughest' pressure on Russia after drone incursion*. Newsweek. <https://www.newsweek.com/poland-will-take-toughest-measures-over-russia-drone-incursion-deputy-pm-2131497>
- Panyi, S. (2026) *The think-tank spy: Russian diplomat infiltrating Orbán's right-wing elite quietly expelled*. VSquare. <https://vsquare.org/the-think-tank-spy-russian-diplomat-infiltrating-orbans-right-wing-elite-quietly-expelled/>
- President of the Russian Federation. (2021). *Стратегия национальной безопасности Российской Федерации*. <http://www.scrf.gov.ru/media/files/file/9htlEYM6sh7mr0riotJdAd7atKKM7w9J.pdf>
- President of the Russian Federation. (n.d.). *Members – Security Council structure*. <http://en.kremlin.ru/structure/security-council/members>
- Reynolds, J. (2026). *"Russia blames UK for deadly Storm Shadow missile strike on munitions factory,"* The Independent. <https://www.independent.co.uk/news/world/europe/russia-britain-storm-shadow-missile-factory-ukraine-b2937047.html>
- Schnieder, M. (2024). *The Russian non-strategic (tactical) nuclear exercise*. RealClear Defense. https://www.realcleardefense.com/articles/2024/09/17/the_russian_non-strategic_tactical_nuclear_exercise_1058803.html
- Selhorst, A. (2016). *Russia's perception warfare: The development of Gerasimov's doctrine in Estonia and Georgia and it's application in Ukraine*. Militaire Spectator. <https://www.militairespectator.nl/sites/default/files/uitgaven/inhoudsopgave/Militaire%20Spectator%204-2016%20Selhorst.pdf>
- Snegovaya, M. (2015). *Putin's information warfare in Ukraine: Soviet origins of Russia's hybrid warfare*. Institute for the Study of War. <https://understandingwar.org/wp-content/uploads/2024/05/Russian20Report20120Putin27s20Information20Warfare20in20Ukraine-20Soviet20Origins20of20Russias20Hybrid20Warfare.pdf>

Stradinger, J. (2024). *Narrative intelligence: Detecting Chinese and Russian information operations to disrupt NATO unity*. Foreign Policy Research Institute. <https://www.fpri.org/article/2024/11/intelligence-china-russia-information-operations-against-nato/>

Thomas, T. (2019). *Russian military thought: Concepts and elements*. The MITRE Corporation. https://oe.t2com.army.mil/site-content/OEE/FMSO%20Books/2019_08_01-Russian-Military-Thought_Concepts-and-Elements-_Timothy-Thomas_.pdf

Thomas, T. (2016). *Russia's 21st century information war: Working to undermine and destabilize populations*. NATO Strategic Communications Centre of Excellence. https://stratcomcoe.org/cuploads/pfiles/timothy_thomas.pdf

Thomas, T. (2004). *Russia's reflexive control theory and the military*. The Journal of Slavic Military Studies. <https://www.tandfonline.com/doi/abs/10.1080/13518040490450529>

Trevelyan, M., & Osborn, A. (2024). *Putin draws a nuclear red line for the West*. Reuters. <https://www.reuters.com/world/europe/putin-draws-nuclear-red-line-west-2024-09-27/>

Vasara, A. (2023). *How to respond to the challenge of reflexive control?* Russian Military and Security Research Group. <https://rusmilsec.blog/2023/12/12/how-to-respond-to-the-challenge-of-reflexive-control/>

Vasara, A. (2020). *Theory of reflexive control: Origins, evolution and application in the framework of contemporary Russian military strategy*. National Defence University. https://www.doria.fi/bitstream/handle/10024/176978/Vasara_FDS22_Theory%20of%20Reflexive%20Control%20%28web1%29-1.pdf

UK Foreign, Commonwealth, and Development Office. (2025). *Profile: GRU cyber and hybrid threat operations*. <https://www.gov.uk/government/publications/profile-gru-cyber-and-hybrid-threat-operations/profile-gru-cyber-and-hybrid-threat-operations>

U.S. Department of State, Global Engagement Center. (2022). *Kremlin-funded media: RT and Sputnik's role in Russia's disinformation and propaganda ecosystem*. https://2021-2025.state.gov/wp-content/uploads/2022/01/Kremlin-Funded-Media_January_update-19.pdf

Author biography

Dr. Ivana Stradner

Rear Admiral (Ret.) Mark Montgomery

Dr. Ivana Stradner serves as a research fellow with FDD's Barish Center for Media Integrity. She studies Russia's security strategies and military doctrines to understand how Russia uses information operations for strategic communication. Her work examines both the psychological and technical aspects of Russian information security.

Rear Admiral (Ret.) Mark Montgomery serves as senior director of the Foundation for Defense of Democracies' (FDD) Center on Cyber and Technology Innovation (CCTI). Mark served for more than three decades in the U.S. Navy, held senior leadership roles in Congress, and is a recognized expert on cyber and technology policy.



NATO
C2COE

ACCOUNTABILITY UNDER AUTONOMY

In Human Oversight in AI-Enabled Command
and Control

The integration of artificial intelligence into military Command and Control (C2) is transforming how decisions are generated, executed, and supervised, particularly in the context of Lethal Autonomous Weapon Systems (LAWS). While these systems enhance operational effectiveness, they also challenge established frameworks of accountability and ethical oversight. When elements of target selection and engagement are delegated to autonomous systems, responsibility becomes embedded within complex socio-technical processes, raising questions about how accountability can be maintained.

This article examines how accountability and ethical oversight can be preserved under these conditions. It adopts the Comprehensive Human Oversight Framework (CHOF), expanding the notion of 'Meaningful Human Control' (MHC). This framework conceptualizes oversight as a continuous and distributed process across governance, socio-technical, and technical layers, and across the phases before, during, and after deployment. It enables analysis of how accountability is institutionally organized within AI-enabled C2 systems.

Empirically, the study draws on a qualitative case study of the Royal Netherlands Army (RNLA), based on semi-structured interviews and document analysis. Three key findings emerge. First, the delegation of decision-making intensifies the demand for demonstrable system reliability as a condition for lawful and accountable use. Second, the technical characteristics of AI systems, including data-dependency, model updates, and opacity, render predictability conditional and fragile. Third, organizational fragmentation across the system lifecycle creates gaps between formal accountability and practical oversight capacity, undermining traceability and complicating legal review mechanisms.

The findings suggest that accountability in AI-enabled C2 depends less on direct human control at the moment of engagement, and more on the institutional production of reliability through continuous validation, technical transparency, and lifecycle-wide integration. CHOF provides a structured lens for identifying these requirements and their associated tensions in practice.

Keywords:

Lethal Autonomous Weapon Systems (LAWS), Accountability, Human Oversight, AI Governance

Introduction

The governance, Command and Control (C2), of Lethal Autonomous Weapon Systems (LAWS) poses an acute challenge, as these intelligent devices can select and engage targets through kinetic force without direct human intervention once deployed (Taddeo & Blanchard, 2022). The integration of artificial intelligence into military C2 is fundamentally transforming how decisions are generated, executed, and supervised. AI-enabled systems enhance situational awareness, accelerate information processing, and improve operational efficiency in increasingly complex and data-saturated environments (van Rijn et al., 2025). All this makes the governance challenge multifaceted and complex. The exercise of force must remain

subject to final human authority and embedded within established ethical, legal, and doctrinal frameworks (van Rijn et al., 2025). When algorithmic systems enter the decision cycle, C2 is reshaped not only as a decision-support architecture but as an accountability structure through which authority, responsibility, and justification are organized.

How to meet this challenge, is subject of this paper. Our central question is: how can accountability and ethical oversight be preserved when lethal decision-making is partially delegated to autonomous systems within military C2? To address this question, the article proceeds in five steps. First, we discuss the scope of the governance challenge. Second, we examine the ethical tensions introduced by autonomous weapon systems, focusing on the

pressure these systems place on core principles of International Humanitarian Law. Third, we analyze how autonomy disrupts established accountability structures within military command chains. Fourth, we examine how concepts of human oversight, particularly Meaningful Human Control (MHC) and the Comprehensive Human Oversight Framework (CHOF), helps to preserve human authority over autonomous systems. Finally, these theoretical perspectives subsequently inform the empirical analysis of oversight practices within the Royal Netherlands Army.

LAWS' rise and the scope of the governance challenge

Recently, major powers have invested heavily in LAWS to maintain military competitiveness (AIV & CAVV, 2021; Kreps, 2021). The Dutch defense strategy similarly emphasizes AI-enabled modernization to ensure operational adaptability (Ministerie van Defensie, 2020, 2023). Recent conflicts demonstrate that autonomous and semi-autonomous systems significantly alter the tempo and character of operations, particularly in high-density drone warfare environments (Bondar, 2025). Strategically, such systems are highly attractive to contemporary armed forces. They reduce decision latency, enable sustained operational tempo, process vast streams of sensor data, and decrease personnel exposure in contested environments (Altmann & Sauer, 2017; Verbruggen & Boulanin, 2017a). Crucially, they are also highly lethal. In Ukraine, autonomous last-mile navigation has increased FPV drone strike success rates from below 20% under manual control to 70–80%, with domestic production reaching approximately two million units in 2024 alone (Bondar, 2025).

Operationally, autonomy resides in core military functions such as targeting, navigation and intelligence (Verbruggen & Boulanin, 2017). In targeting, Automatic Target Recognition systems detect, classify, and prioritize objects at machine speed. In navigation, platforms autonomously traverse contested terrain using sensor fusion and real-time environmental assessment. Intelligence functions allow systems to process and interpret large volumes of data, supporting situational

awareness and threat analysis. Beyond these core functions, autonomy also enables new forms of multi-platform coordination including machine-to-machine communication and swarming behaviors, as part of joint operations. However, such coordination constitutes an operational requirement rather than a core military function in itself.

The disruptive implications of LAWS for C2 can be understood through four defining characteristics (Taddeo et al., 2022). First, autonomy enables execution of engagement-related tasks without continuous human intervention (ICRC, 2019). Second, adaptive capability, often driven by machine learning, allows systems to modify internal decision parameters based on data exposure, potentially altering behavior across contexts (Bode & Watts, 2023). Third, human control may be positioned at varying levels, in, on, or out of the loop, reshaping how authority is retained or delegated (Taddeo et al., 2022). Fourth, the purpose of use anchors these systems in kinetic military application, directly implicating legal and ethical norms governing the use of force. Together, these characteristics relocate elements of target selection, prioritization, and engagement into socio-technical systems whose behavior may be adaptive, opaque, and difficult to reconstruct.

Within this functional and substantive scope, scholars have proposed conceptual frameworks to address these governance challenges, notably Meaningful Human Control (MHC; Santoni de Sio & Van den Hoven, 2018) and the Comprehensive Human Oversight Framework (CHOF; Verdiesen et al., 2021). Both seek to preserve human authority and accountability in the use of autonomous systems; we elaborate this in detail below. However, these frameworks remain largely conceptual, and their practical implementation within military C2 structures has received limited empirical attention. In particular, little is known about how lifecycle-wide human oversight can be operationalized once autonomous capabilities are embedded within operational decision cycles. Before discussing aspects of oversight itself, we first need to address pertinent ethical demands.

Ethical demands: proportionality, distinction, necessity and humanity

The deployment of LAWS places structural pressure on the core principles of International Humanitarian Law (IHL) and, by extension, on military doctrines that embed them. IHL regulates the conduct of hostilities through principles including proportionality, distinction and military necessity and humanity (ICRC, 2011). These principles presuppose evaluative judgment: combatants must distinguish lawful targets from civilians, assess proportionality between anticipated military advantage and incidental harm, and ensure that force is necessary to achieve legitimate objectives, upholding moral responsibility (Boutin & Woodcock, 2024; Hughes, 2020; Kwik, 2024). We will now discuss these principles in turn.

Proportionality requires evaluative balancing between anticipated military advantage and incidental civilian harm. Machine learning systems, however, operate through probabilistic pattern recognition rather than normative reasoning. While operators may optimize accuracy within defined parameters, they lack contextual awareness and ethical deliberation, rendering proportionality assessments dependent on algorithmic thresholds and training data distributions (Bode & Watts, 2023; Devitt, 2024). As Schwarz (2021) notes, determining what counts as “excessive” harm is inherently normative and cannot be reduced to purely computational inputs.

Similarly, the principle of distinction relies on context-sensitive identification of lawful targets. Although technologies such as Automatic Target Recognition improve classification under controlled conditions, their performance degrades in complex environments, and they lack the interpretive capacity to assess intent or ambiguous behavior (Booker, 2024; Christie et al., 2024; Kwik, 2024). Errors under degraded operational conditions, including sensor interference or incomplete data, further amplify risk (Boulanin & Lewis, 2023).

Necessity and humanity raise deeper concerns. Encoding “imminent threat” or “military objective” into computational models risks transforming ethical judgment into predefined technical criteria

(Amoroso & Tamburrini, 2021; Devitt, 2024). As Santoni de Sio and van den Hoven (2018) argue, moral responsibility requires agency, understood as “the capacity to act intentionally, make normative judgments, and be held accountable for one’s decisions”, alongside the capacity for justification, attributes absent in machine systems.

Accountability ‘vacuum’

These IHL principles raise subsequential questions concerning accountability: how are these principles upheld? Who is responsible when unlawful harm occurs? Conventional military accountability frameworks rely on identifiable human agency and traceable chains of command, allowing decisions to be legally and institutionally scrutinized (Smith, 2022; Wood, 2023).

In AWS-enabled operations, however, decision-making becomes distributed across developers, procurement bodies, commanders, operators, and algorithmic systems, without any single actor necessarily retaining full situational control at the decisive moment (Smith, 2022; Wood, 2023). When unlawful harm occurs, this dispersion creates what has been described as an “accountability vacuum,” in which responsibility cannot be clearly located within established command hierarchies (Cavalcante Siebert et al., 2023; Ebrahimi, 2024; Umbrello, 2021).

The “accountability vacuum” is intensified by the opacity of machine-learning systems. Operating as technical “black boxes,” these systems often lack accessible internal reasoning pathways, complicating both traceability and explainability in post-strike investigations (Bode & Watts, 2023; Christie et al., 2024). Attribution dilemmas emerge when misclassification or erroneous engagement occurs, as it becomes difficult to determine whether failure stemmed from flawed training data, algorithmic limitations, operational misuse, or unpredictable battlefield conditions (Devitt, 2024). Without reliable auditing mechanisms, legal review and ethical evaluation are weakened, undermining the integrity of oversight processes within democratic military institutions (Kwik, 2024).

Human Oversight: tracing, tracking and layered institutionalization

In response to accountability issues, the concept of MHC has emerged as a normative framework for preserving human authority over lethal decision-making. Among the most influential operationalizations of MHC are the conditions of tracking and tracing proposed by Santoni de Sio and van den Hoven (2018). Together, these conditions seek to ensure that autonomous systems remain embedded within human moral and institutional responsibility structures.

Tracking requires that an autonomous system be responsive to morally and legally relevant features of its operational environment. In practice, this means that system behavior must align with the reasons and intentions of human decision-makers, including compliance with IHL principles such as distinction and proportionality. Tracking requires LAWS to prospectively register civilian presence, contextual constraints, and rules of engagement, in line with ruling normative frameworks. Tracking thus operationalizes the normative conditions for MHC. However, this presupposes that system behavior remains sufficiently predictable and aligned with human intent across varying operational contexts. As highlighted in the literature on AI-enabled systems, such alignment cannot be assumed (Christie et al., 2024; Schwarz, 2021). Machine-learning systems operate through probabilistic pattern recognition and may exhibit context-sensitive variability, particularly when exposed to environments that diverge from their training conditions. This introduces uncertainty to the extent to which systems can reliably track morally relevant features, thereby challenging the practical realization of this condition.

Tracing, by contrast, concerns the retrospective dimension of accountability. It demands that system actions be traceable to identifiable human agents who can explain and justify deployment decisions. This requires technical mechanisms such as logging, data provenance documentation, and explainable system outputs, enabling post-strike reconstruction and institutional review (Christie et al., 2024). Without such mechanisms,

accountability risks becoming symbolic rather than enforceable. Yet, the feasibility of tracing is complicated by the opacity of many AI systems. The limited interpretability of algorithmic decision-making processes, combined with potential gaps in data availability or accessibility, raises questions about whether traceability can be meaningfully achieved in practice.

Taken together, tracking and tracing provide a strong normative articulation of MHC (Santoni de Sio & van den Hoven, 2018). However, MHC remains conceptually limited in accounting for how these conditions can be realized in complex socio-technical systems. The concept is applied inconsistently across domains and often presupposes direct human intervention, which is no longer feasible once autonomous systems are deployed (Ekelhof, 2019; Verdiesen et al., 2021). Moreover, MHC primarily focuses on human-machine interaction, offering only a partial account of how responsibility and accountability are structured beyond the level of individual operators (Verdiesen et al., 2021).

This raises the question of how tracking and tracing can be institutionally organized within military command structures. Verdiesen et al. (2021) address this by proposing the CHOF, a distributed architecture of oversight spanning multiple layers of military organization across the lifecycle of autonomous weapon systems. By shifting the analytical focus from control to oversight, CHOF conceptualizes accountability as a continuous and distributed process across the phases before, during, and after deployment (Verdiesen et al., 2021).

CHOF distinguishes between three architectural layers: governance, socio-technical, and technical. At the governance layer, oversight is embedded in doctrinal guidance, legal review procedures, and the institutional allocation of command responsibility, ensuring that the development and deployment of autonomous capabilities remain aligned with international humanitarian law and national military doctrine. The socio-technical layer concerns the operational practices through which humans interact with autonomous systems in practice,

including training procedures, rules of engagement, and the integration of human-machine interaction within command structures. Finally, the technical layer addresses the design and verification of the systems themselves, including data governance, algorithmic validation, testing procedures, and traceability mechanisms that enable post-use reconstruction of system behavior.

LAWS' lifecycle process spans three phases: pre-deployment, deployment, and post-deployment. Oversight begins in the pre-deployment phase, where governance decisions, procurement processes, system design choices, and testing procedures establish the conditions under which autonomy may be responsibly employed. During deployment, oversight shifts toward operational control, including the interaction between human operators and autonomous systems, the enforcement of rules of engagement, and the monitoring of system behavior under real-world conditions. Finally, the post-deployment phase focuses on retrospective accountability and institutional learning through

mechanisms such as logging, after-action reviews, and battle damage assessments. By structuring oversight across these temporal phases, CHOF emphasizes that MHC cannot be confined to the moment of system activation but must be sustained throughout the entire lifecycle of autonomous capabilities.

Taken together, the layered and lifecycle dimensions of the CHOF provide a structured lens for examining how ethical oversight and accountability are operationalized in military organizations (Figure 1). Rather than treating oversight as a single decision point, CHOF conceptualizes oversight as a distributed process that spans institutional governance, operational practice, and system design across the full lifecycle of autonomous capabilities. In this study, the framework is used to examine, through empirical investigation, how oversight mechanisms are organized within the RNLA. We focus on tensions emerging when autonomous functionalities are integrated into so-called operational C2 structures.




Figure 1: The Comprehensive Human Oversight Framework (CHOF) as developed by Verdiesen et al. (2021), organizing oversight across time (before, during, after) and system environment (technical, socio-technical, governance)

Method

This study adopts a qualitative case study design to examine how lifecycle-wide human oversight is operationalized within a national military context. The RNLA was selected as a relevant case due to its explicit ambitions to integrate AI-enabled systems into future force structures and its established commitment to compliance with International Humanitarian Law.

Empirical data were collected through eleven semi-structured interviews with military stakeholders across operational, legal, technical, and procurement domains, assembled through purposive sampling strategy in close dialogue with the RNLA. Interviews were held with two legal advisors, four operational commanders, two military AI specialists, all of the Dutch Army, two policymakers from the Ministry of Defense and one external drone expert, during April and May 2025. The interviews used a semi-structured format, covering the themes discussed above (see Appendix). Empirical data also consisted of policy documents from the Dutch parliament, NATO, and an advisory report from the AIV & CAVV, mostly on advice of the interviewees, complemented by literature research on core topics. The data were thematically analyzed using the layers of the CHOF as an analytical lens through Atlas-TI software. More specifically, the data reveals concrete organizational structures, validation processes, accountability procedures, and technical reliability conditions as they currently exist or are being developed within the RNLA. This approach allows assessment of how tracking and tracing requirements are institutionally embedded and where oversight architectures encounter practical strain once autonomy enters operational C2 structures. Thematic analysis pointed out three core issues: the commanders' need for reliability, technical demands, and governance fragmentation.

Finding One. Commanders' need for Reliability

The first finding is that the delegation of decision-making intensifies the institutional demand for reliability. Delegation of execution authority is not itself new; it is foundational to Western military doctrine (Kalloniatis, 2025). Mission command (Auftragstaktik) has long premised

C2 on commanders specifying intent and enabling subordinates to exercise initiative within operational constraints (Kalloniatis, 2025). The integration of LAWS into the engagement cycle follows the same structural logic. What changes is not the fact of delegation but the nature of the delegate. Human subordinates exercise contextual judgment. Humans can deviate from orders when circumstances demand it, and bear personal moral and legal responsibility for their actions. Autonomous systems can do none of these things. Delegation to LAWS therefore generates a structural imperative that mission command does not. Because the system cannot adapt wisely or be held accountable for its own decisions, the full weight of responsibility concentrates in the quality of the prior authorization and the demonstrable reliability of system behavior within predefined boundaries. One commander sketches the process of working with autonomous systems: *"Do I have sufficient information about my surroundings? About the enemy, about the terrain. About the performance factors of an autonomous system. About the presence or absence of civilians. How certain am I of that? What is my most recent information? All questions are directed to the commander. And then he assesses the situation and decides whether the use of autonomy is necessary and applicable."*

This distinction reshapes the logic of predictability within military C2. In conventional C2, predictability is anchored in trained human judgment and hierarchical oversight, sustaining real-time intervention. With autonomous systems, however, delegation becomes conditional, and predictability must be engineered and validated. Command responsibility remains formally unchanged; commanders retain legal and doctrinal accountability for the use of force. Yet their ability to justify delegation now depends on whether system behavior can be reasonably anticipated across variable operational contexts. An AI researcher within the RNLA connects predictability to accountability: *"If you're going to hold somebody accountable, there has to be a level of predictability in that it was going to go wrong. You can't hold somebody accountable for something that they could have in no way known was going to happen."*

Our respondents indicate that autonomy amplifies sensitivity to system performance characteristics. One comment was: *“the more usefully autonomous, the more advantages you can draw from the autonomy, the more black box it generally becomes. So, we will have to find a path on what we still find acceptable and what not”* Accuracy alone is insufficient. What becomes critical is behavioral stability across updates, robustness under degraded conditions, and alignment between training data and deployment environment. The more autonomy is introduced into detection, classification, or prioritization functions, the stronger the institutional demand for demonstrable predictability. Without it, delegation risks exceeding what can be lawfully or ethically justified.

Autonomy thus does not reduce the burden of oversight but reconfigures it. It shifts oversight upstream, toward validation, testing, and system understanding. Reliability becomes the operational bridge that sustains lawful delegation within C2. Where predictability cannot be demonstrated, accountability structures are not formally removed, but they become increasingly difficult to operationalize in practice.

Finding Two. Reliability as a technical condition

The second finding is that about reliability's technical supply. While autonomy generates an institutional demand for predictability, the technical characteristics of AI-enabled systems are becoming more demanding. Within the RNLA context, several structural tensions emerged between doctrinal responsibility and the technical behavior of adaptive systems. This poses three challenges: data-domain alignment, system updates, and opacity.

First of all, respondents point out the problem of *data-domain alignment*: *“Quality of AI increases with the quality and quantity of the data you use”*. But the alignment between the training data and the actual operational theater is not always guaranteed. Autonomous systems trained on specific datasets may encounter operational environments that diverge significantly from training conditions. Even minor shifts in terrain, sensor input quality, or adversarial tactics can affect classification performance. As autonomy increases, such shifts

no longer merely influence advisory outputs; they shape decision-relevant system behavior. Predictability therefore becomes contingent upon the continued, always evolving, congruence between training data and deployment contexts. One officer questions whether it is even possible to create complete congruence between data and operational theater: *“We try to simulate operational situations during exercises, but I don't think the quality is realistic enough to train an AI-model.”*

Dependency on data and algorithms also creates a second challenge; system updates and retraining cycles. Machine-learning models may evolve through software updates, parameter adjustments, or integration with new data streams. While such updates aim to improve performance, they simultaneously destabilize prior validation results. Reliability becomes a moving target: a system verified at one moment may exhibit altered behavior after modification. This dynamic complicates the assurance that previously established operational boundaries remain valid. One legal adviser tells how a retraining with different data can alter performance, and the subsequent legal review status: *“It's certainly possible that it will be further developed, that at some point there will be a kind of data update to the system, which will suddenly allow you to do certain things better, but also differently. The moment its operation changes significantly, or its effect changes significantly, then you'll really have to conduct another review.”*

Third, opacity undermines traceability. Even when performance metrics are available, internal decision pathways based on algorithms remain inaccessible or difficult to reconstruct. In the event of erroneous classification or unintended engagement, distinguishing between data bias, algorithmic limitation, or contextual misapplication becomes technically complex. Without granular logging and explainability mechanisms, post-use accountability relies on probabilistic inference rather than demonstrable reconstruction. One AI researcher however is skeptical about the ability to understand algorithms: *“Humans can't see inside that algorithm. And we won't solve that by using explainable AI either. Because for instance, an image recognition algorithm, you can never make that fully explainable.”* One operational officer refuse to use

autonomous systems he cannot explain: *"You end up with something like a black box and as a user, I don't want that. Because I don't trust it"*

Together, these challenges reveal a central tension: autonomy intensifies the requirement for predictability, yet the technical architecture of AI-enabled systems renders predictability conditional and fragile. Reliability, rather than being a fixed property, becomes dependent on continuous validation, monitoring, and institutional integration. Where such mechanisms are insufficiently developed, the foundation of lawful delegation within C2 is placed under structural strain.

Finding Three. Governance Fragmentation and Distributed Responsibility

Third, these (socio)technical tensions are amplified by organizational fragmentation across the lifecycle of autonomous systems. While CHOF presupposes integrated oversight spanning governance, socio-technical, and technical layers, the empirical findings indicate that responsibilities are often distributed across institutional boundaries without full systemic coordination.

Procurement, development, validation, and operational deployment are typically handled by different actors, including acquisition bodies, external industry partners, technical specialists, and operational commanders. This distribution is not inherently problematic as modern defense organizations rely on specialization. However, when autonomy is introduced into core decision functions, fragmented responsibility complicates the maintenance of reliability across the LAWS' lifecycle. One AI expert explains: *"There's the classic discussion, is it the programmers? Is it the company as a whole? Is it the commander who deployed it that is responsible? And also, which company, because these things are normally, fully autonomous systems, may have multiple software's on it from different companies. So, is it one specific piece of that autonomous system that went wrong? Is it the system as a whole that went wrong?"*

In particular, technical expertise necessary to evaluate model behavior, data provenance, and retraining implications may reside outside operational command structures. Commanders remain accountable for system use, yet may lack direct access to the technical transparency required to independently assess behavioral stability. A representative of an AI software company even says: *"I do not expect a corporal, sergeant or officer to know exactly how a neural network works. That is an organizational responsibility that must be placed elsewhere in the organization. With the weapons system manager and the norm setter. The user knows the general limitation."*

Similarly, systems introduced through procurement channels or industry collaboration may not be fully integrated into operational validation cycles. This may cause problems, like this officer notes: *"But I don't believe in collaborating with industry on AI. I don't think we should be buying our tactics and decision-making from companies. If you start buying that decision-making from a company, you're essentially buying their vision of warfare. I think we should always retain control over those decision-making algorithms. And even if we don't develop it ourselves, we should at least be able to assess for ourselves whether they're accurate."* Another officer is also opposed to externally developed AI: *"So from the stance of the Dutch Army, you can't really control how those algorithms work or what kind of data they use. To some extent, yes, but you do need to understand it. But how would you do that? You need an army of tech experts who understand this."*

This institutional dispersion not only complicates operational oversight, but also challenges formal legal review mechanisms, such as Article 36 weapons reviews (the obligation under Additional Protocol I (1977) for states to assess whether new weapons comply with international law before deployment), which depend on sufficient transparency, explainability, and traceability to assess compliance with international humanitarian law (Christie et al., 2024; Verbruggen & Boulanin, 2017b). One legal officer explains how the RNLA does not have answers for this problem: *"If you start building autonomous weapon systems that can independently select and engage targets, it is still an*

open question as to which parts of the software need to be transparent to your organization in order to conduct a legal review. That obligation comes from Article 36. These are all issues for which, honestly, much still needs to be resolved.'

This institutional dispersion creates gaps between formal accountability and practical oversight capacity. Reliability becomes dependent not only on technical design, but on sustained coordination between procurement authorities, developers, legal advisors, and operational users. Where lifecycle integration is incomplete, the tracing condition, understood as the ability to reconstruct and attribute system behavior to responsible actors, is significantly undermined. Autonomy therefore exposes an underlying structural misalignment: accountability remains hierarchically defined, while control over technical behavior is distributed across a network of actors. Without robust governance integration, the predictability imperative generated by autonomy cannot be fully satisfied within existing C2 arrangements.

Conclusion

The RNLA case demonstrates that the CHOF provides a robust diagnostic lens for assessing human control across the lifecycle of autonomous systems. It helps to reveal various tensions between different levels and phases.

One finding is how structural strain notably occurs at the technical level. CHOF presumes that governance, operational practice, and system design can be coherently aligned through lifecycle-wide oversight. In practice, this alignment becomes contingent upon demonstrable reliability. At the governance level, oversight structures remain formally intact. Legal review processes, doctrinal accountability principles, and command responsibility frameworks continue to structure decision-making. At the socio-technical level, training procedures, operational boundaries, and rules of engagement reflect awareness of autonomy-related risks. However, autonomy shifts critical decision-relevant functions into algorithmic processes whose predictability depends on complex conditions of data integrity, model stability, retraining control, and traceability

mechanisms. Where these conditions are insufficiently institutionalized, the linkage between legal norms and technical behavior weakens.

Moreover, the translation of abstract ethical principles into enforceable technical requirements proves challenging. While governance structures articulate compliance with IHL, ensuring that system architectures consistently reflect those norms requires granular insight into training datasets, model behavior under environmental variation, and update procedures. CHOF identifies the need for oversight across these layers, but the RNLA case shows that operationalizing this integration demands technical transparency and cross-domain expertise that are not automatically present.

The central tension therefore lies not in the absence of normative commitment, but in the conditional nature of technical predictability. Autonomy does not eliminate accountability structures; rather, it renders them dependent on continuous validation and monitoring. Where reliability cannot be demonstrated across the system lifecycle, lawful delegation becomes increasingly difficult to justify.

CHOF remains analytically valuable because it exposes precisely where oversight depends on technical robustness. Yet the findings indicate that lifecycle-wide human control under AI-enabled autonomy requires deeper institutional integration between governance, procurement, and technical validation processes than current structures fully provide.

The empirical findings also highlight the disruptive implications of several defining characteristics of autonomous weapon systems. Adaptive capability challenges traditional validation models, as system behavior may evolve through retraining or environmental variation, rendering static certification procedures insufficient. Similarly, the positioning of human control in, on, or out of the loop alters how authority is exercised within C2 structures, shifting oversight from real-time supervision toward upstream processes such as system validation and operational boundary setting.

Finally, the purpose of use of these systems, namely

the application of lethal force, intensifies the ethical and legal stakes associated with autonomy, placing greater pressure on accountability and oversight mechanisms. Taken together, these characteristics do not eliminate human responsibility, but they fundamentally reconfigure how control and accountability must be organized within military command structures. On this basis, we formulate five recommendations.

Practical recommendation

The first recommendation is to operationalize the lifecycle-wide oversight envisioned by CHOF. The empirical findings indicate that linear procurement models and one-off legal reviews are insufficient where system behavior may evolve through retraining, data updates, or contextual drift. Reliability and predictability cannot be presumed at the point of certification; they must be institutionally produced and continuously revalidated. Accordingly, the RNLA must adapt acquisition, validation, and accountability practices to the realities of adaptive AI-enabled systems.

Second, a central implication for military C2 is the need for continuous, scenario-realistic Testing, Evaluation, Verification, and Validation (TEVV). Rather than treating legal and technical review as pre-deployment thresholds, oversight must incorporate structured triggers for reassessment when operational design domains, training data, or model architectures change. Predictability becomes a dynamic property requiring sustained monitoring.

Third, this transformation also demands greater algorithmic transparency. Proportionality and necessity cannot be meaningfully operationalized through static performance metrics alone. Translating these principles into enforceable technical constraints requires access to model calibration procedures, traceable data lineage, and inspectable logging mechanisms. Opening the “black box” is not solely a technical aspiration but a condition for credible accountability within democratic military institutions.

Fourth, organizational adaptation is equally critical. Hybrid expertise must bridge operational command and technical development, ensuring that commanders retain substantive control rather than functioning as passive authorizers of opaque systems. Data governance structures should secure institutional ownership over datasets, model states, and telemetry, safeguarding operational sovereignty and enabling forensic reconstruction. Procurement relationships must evolve from transactional acquisition toward collaborative co-development models that guarantee transparency over system updates and algorithmic behavior.

Finally, accountability mechanisms must integrate standardized system logs into ‘After Action Reviews’ and ‘Battle Damage Assessments’, strengthening the evidentiary foundation for legal and ethical scrutiny. Only by embedding continuous validation, technical transparency, hybrid expertise, and robust data governance can lifecycle oversight move from formal commitment to operational credibility. Under AI-enabled autonomy, ethical C2 depends less on declarative compliance and more on the ‘live’ institutional production of reliability.

References

- AIV, & CAVV. (2021). Autonome wapensystemen: Het belang van reguleren en investeren.
- Altmann, J., & Sauer, F. (2017). Autonomous weapon systems and strategic stability. *Survival*, 59(5), 117-142.
- Amoroso, D., & Tamburrini, G. (2021). Toward a normative model of meaningful human control over weapons systems. *Ethics & International Affairs*, 35(2), 245-272.
- Bode, I., & Watts, T. F. A. (2023). Loitering munitions and unpredictability: Autonomy in weapon systems and challenges to human control.
- Bondar, K. (2025). Ukraine's Future Vision and Current Capabilities for Waging AI-Enabled Autonomous Warfare.
- Booker, J. O. (2024). War Machines of Tomorrow: Accountability and Oversight in the Age of Lethal Autonomous Weapon Systems.
- Boulanin, V., & Lewis, D. A. (2023). Responsible reliance concerning development and use of AI in the military domain. *Ethics and Information Technology*, 25(1), 8.
- Boutin, B., & Woodcock, T. (2024). Aspects of realizing (meaningful) human control: a legal perspective. In *Research Handbook on Warfare and Artificial Intelligence* (pp. 179-196). Edward Elgar Publishing.
- Cavalcante Siebert, L., Lupetti, M. L., Aizenberg, E., Beckers, N., Zgonnikov, A., Veluwenkamp, H., Abbink, D., Giaccardi, E., Houben, G.-J., & Jonker, C. M. (2023). Meaningful human control: actionable properties for AI system development. *AI and Ethics*, 3(1), 241-255.
- Christie, E. H., Ertan, A., Adomaitis, L., & Klaus, M. (2024). Regulating lethal autonomous weapon systems: exploring the challenges of explainability and traceability. *AI and Ethics*, 4(2), 229-245.
- Devitt, S. K. (2024). Meaningful human command: Advance control directives as a method to enable moral and legal responsibility for autonomous weapons systems. In *Research Handbook on Meaningful Human Control of Artificial Intelligence Systems* (pp. 53-80). Edward Elgar Publishing.
- Ebrahimi, V. (2024). The Role of Law in Governing Artificial Intelligence in the Context of Global Warfare. *Legal Studies in Digital Age*, 3(3), 23-30.
- Ekelhof, M. (2019). Moving beyond semantics on autonomous weapons: Meaningful human control in operation. *Global Policy*, 10(3), 343-348.
- Hughes, J. G. (2020). The law of armed conflict issues created by programming automatic target recognition systems using deep learning methods. *Yearbook of International Humanitarian Law*, Volume 21 (2018), 99-135.
- ICRC. (2019). Autonomy, artificial intelligence and robotics: Technical aspects of human control
- Kalloniatis, A. (2025). A 'Brief' History of Command and Control. *Annals of Command and Control*, 2(1), 6-43.
- Kreps, S. (2021). Democratizing harm: Artificial intelligence in the hands of nonstate actors. *Foreign Policy*.
- Kwik, J. (2024). *Lawfully Using Autonomous Weapon Technologies*. Springer.
- Ministerie van Defensie. (2020). Defensievisie 2035.
- Ministerie van Defensie. (2023). Defensie Strategie Data Science en AI 2023-2027.
- Schwarz, E. (2021). Autonomous weapons systems, artificial intelligence, and the problem of meaningful human control. *Philosophical Journal of Conflict and Violence*.
- Smith, P. T. (2022). Resolving responsibility gaps for lethal autonomous weapon systems. *Frontiers in big data*, 5, 1038507.
- Taddeo, M., & Blanchard, A. (2022). A Comparative Analysis of the Definitions of Autonomous Weapons Systems. *Sci Eng Ethics*, 28(5), 37. <https://doi.org/10.1007/s11948-022-00392-3>

Taddeo, M., McNeish, D., Blanchard, A., & Edgar, E. (2022). Ethical principles for artificial intelligence in national defence. In *The 2021 Yearbook of the Digital Ethics Lab* (pp. 261-283). Springer.

Umbrello, S. (2021). Coupling levels of abstraction in understanding meaningful human control of autonomous weapons: A two-tiered approach. *Ethics and Information Technology*, 23(3), 455-464.

van Rijn, M., Saylam, R., Scherrenburg, M., & Dönmez, M. (2025). AI in Military C2 Systems: An Introduction and Recent Advances. *Annals of Command and Control*, 2(1), 86-100.

Verbruggen, M., & Boulanin, V. (2017a). Mapping the development of autonomy in weapon systems.

Verbruggen, M., & Boulanin, V. (2017b). SIPRI compendium on article 36 reviews.

Verdiesen, I., Santoni de Sio, F., & Dignum, V. (2021). Accountability and control over autonomous weapon systems: a framework for comprehensive human oversight. *Minds and Machines*, 31(1), 137-163.

Wood, N. G. (2023). Autonomous weapon systems and responsibility gaps: a taxonomy. *Ethics and Information Technology*, 25(1), 16.

Author biography

Kjell Schipper, MSc

Prof. Dr. Arnoud Legendijk

Kjell Schipper, MSc recently graduated cum laude in Human Geography (Conflicts, Territories and Identities) with a thesis on the Ethics of Lethal Autonomous Weapon Systems. Before completing this master's degree, he obtained degrees in both Business Administration and Psychology.

Arnoud Legendijk is Professor of Economic Geography at Radboud University Nijmegen and a board member of the Dutch Education Union (AOB). He has a long track record in the field of regional development and policy, and has published extensively on innovation, clustering, strategy formulation, regional identities, and mobility. He is currently working on regional coalition-building aimed at the transition to a circular economy

Authors Note

The authors have not received any funding for writing this paper. There are no conflicts of interest.

Acknowledgements

I would like to thank Lieutenant-Colonel Joost Tuinman for his support throughout this research and for his thoughtful advice, generously given in his own time. I am also grateful to Captain Aya Dürst-Britt of the NATO Command and Control Centre of Excellence for her belief in this study and for encouraging me to participate in the Student Challenge; her energy and dedication are admirable. Finally, I thank all experts who shared their time and expertise.

The interviews covered the followings topics (used variably depending on interviewees)

Introduction: general function and tasks; timeframe, role and collaboration in relation to AWS; the AWS development process within the Army: role, significance, drivers, stakeholders (collaboration), existing requirements (e.g. regarding algorithms),

Training data: production, nature, use, learning process, validation, sharing of results.

Autonomy: interpretation (differences), envisaged forms of operational implementation and decision-making, use in the targeting cycle (deliberate and dynamic targeting), adjustments to Rules of Engagement or Doctrine, ethical design and accountability (necessity, reliability, proportionality, role of context,

Strategy: envisaged use of autonomous weapon systems by the Army, with what scope of ethical and concrete standards (how translated into technical requirements)?

Meaningful Human Control: interpretation (differences), form of 'human-in-the-loop control', accountability mechanisms and channels (specifically in the event of errors), limitations of AWS logs. ICH has

Development: learning capability (scope and form of feedback loop regarding the use of AWS), improving ethical oversight of AWS within the Ministry of Defense, desired education and training in safeguarding ethics



NATO
C2COE

A COGNITIVE FRAME- WORK OF ASSESSING LEADERSHIP

In the Context of Multi-Domain Operations

Mehmet Akif Morkaya

DIR of International Business Development

CNK GLOBAL, Türkiye



This paper proposes a new systematic framework for assessing leadership as a fundamental component of Command and Control functions in the context of NATO Multi-Domain Operations. The study specifically explores the cognitive dimension of MDO beyond its physical and virtual dimensions. The central research question is structured as follows: How should a Cognitive Framework of Assessing Leadership be configured to enhance Command and Control (C2) effectiveness in the context of MDO? This study addresses the central research question based on two perspectives: On the one hand, there is sparse NATO doctrine or procedures for assessing the success of leadership within C2 structures, which may be because formal leader development is a national responsibility. On the other hand, Allies must reconsider building their recruitment, selection, training, promotion, and retention systems to deliver the right people with the right skills, as accentuated in the concept of MDO.

This study has identified a lack of leadership assessment in NATO and argues that the success of NATO civil & military leadership needs to be predictably assessed through a cognitive approach to ensure resilient and adaptive C2. Therefore, this study adopts a conceptual approach by analyzing relevant literature and insights from leadership theories. Besides, the cognitive framework focuses on knowledge dimensions and cognitive processes. In conclusion, the Cognitive Framework of Assessing Leadership contributes to the C2 literature by providing an institutional mechanism to assist the Alliance in identifying leaders and staff who can cope with the future environment's speed, complexity, and uncertainty.

Keywords:

Assessing Leadership, Command and Control (C2), Knowledge Dimensions, Cognitive Processes.

1.0 INTRODUCTION

The paper proposes a definitive synthesis based on the assumption that culture clusters and desired/preferred leadership behaviours or styles are categorically listed in NATO. Thereby, an effective cross-cultural leadership demands propositions explicitly tailored to the operational environment, anchored by a foundational understanding of the prevailing culture cluster. This cultural mapping mandates the specific values and expectations that shape organizational dynamics, thereby dictating the requirement for a culturally congruent preferred leadership style – one recognized as legitimate and effective by the cluster's members. It is asserted that a gap exists in the current academic literature regarding sufficient clarity on the precise point at which 'leadership' commences within a command structure, or how it should be corporately defined specifically within the NATO civil and military leadership context.

This critical ambiguity necessitates the strategic synthesis of the key planning elements of the Cognitive Framework of Assessing Leadership, which shifts its primary focus from real-time situational awareness and current operational reporting to a strategic function: the direct guidance of national or NATO recruiters on the cognitive qualification measurements essential for future leadership assignment. This predictive stance is necessary to select personnel who are equipped to manage prevalent cognitive challenges, specifically mitigating subjective systemic distortions in decision-making and sustaining high-level functional clarity when faced with contradictory demands.

The application of this Cognitive Framework of Assessing Leadership (CFAL) methodology will confirm the integration of these critical inputs, thereby establishing the key planning elements of this framework and articulating the CFAL's overall design as the institutional mechanism required for culturally aware and cognitively prepared NATO civil and military leadership.

1.1 Essential Principles

Despite the vast number of existing leadership assessment frameworks in the academic realm of leadership theories, they might not fully capture the complexities of a modern world that is constantly changing. Furthermore, a one-size-fits-all approach to assessing leadership may not be effective, considering that academic research continually generates new insights into leadership theories, and a fresh methodology can integrate these insights to offer a more comprehensive understanding.

More specifically, every organization has a distinct organizational culture, mission, and set of risks, challenges, and assumptions. Therefore, it is considered that a tailored methodology can better address exclusive needs, such as NATO Multi-Domain Operations (MDO). In this regard, a theoretical study can explore the cognitive dimension of MDO beyond its physical and virtual dimensions by developing a Cognitive Framework of Assessing Leadership that may focus on evaluating the cognitive processes and knowledge dimensions.

Culminating areas of focus in a cognitive framework can be categorized by cognitive complexity, cognitive skills, executive intelligence, metacognition, etc. The ultimate goal of the study is to create a cognitive framework that is not only theoretically predictive but also practically applicable for NATO civil & military leaders who can cope with the future environment's speed, complexity, and uncertainty.

1.2 Structuring the study with golden threads

The Cognitive Framework of Assessing Leadership is conceptualized as a cognitive process in the MDO context that requires products that can also include personality and motivational questionnaires, cognitive ability tests, interviews, 360-degree feedback, and simulation exercises, etc. In general,

the CFAL is a process of identifying the cognitive qualification measurements for a national or NATO recruiter's selection to provide the right people with the right skills before assigning them to NATO civil & military leadership vacancies. Because of that, the most important achievement of the CFAL is that NATO civil & military leadership is predictably assessed through a cognitive approach; more importantly, this predictive assessment contributes to assisting the Alliance in needing leaders and staff who can cope with the future environment's speed, complexity, and uncertainty.

CFAL intends to provide a robust analytical methodology that develops deeper insights into civil & military leadership theories. And yet, the author argues that the vast knowledge in civilian leadership academia and the sporadic knowledge in military NATO documentation cannot be easily disaggregated and directly integrated into a cognitive framework without a comprehensive planning approach¹. Therefore, this study detects and appraises the key planning elements that are critical to composing the CFAL. Moreover, the author intends to recombine ideas that, to his knowledge, have never been connected before by inspiring an Operations Assessment (OPSA) perspective², which is adopted in assessing key planning elements³. In this regard, the following four types of key planning elements' definitions describe the functional morphology of the key planning elements (KPE) within the CFAL.

1.2.1 Generative KPE

This type of KPE is responsible for actively translating high-level CFAL requirements (derived from the Multi-Domain Operations context and leadership challenges) into precise, measurable cognitive criteria for assessment. Their main function is to conceptually outline the exclusive goals, expressed as changes in the system's state within the framework. The overall function is to facilitate the creation of the measurable assessment standards

1 The author is partially inspired by the ref. [1] to create the CFAL.

2 The ref. [2] explains the OPSA perspective for assessing key planning elements.

3 The ref. [3] provides descriptions of some of these key planning elements in details in the Chapter 4 – Orchestrating fighting power, Section 1 – Doctrinal key tenets.

based on the cognitive needs of the framework and the prevailing operational environment. In the CFAL planning terminology, the Generative KPEs encompass the CFAL end state, CFAL objective, and CFAL lines of effort, as follows:

CFAL End State⁴ (CES). The end state is the statement of conditions that defines a summary of the acceptable situation. In the CFAL planning terminology, an end state is produced through the achievement of objectives.

CFAL Objective⁵ (CO). An objective is a clearly defined and feasible aim or target within the planning process. In the CFAL planning terminology, objectives are achieved by consolidating decisive conditions through effects and their underlying actions.

CFAL Lines of Efforts⁶ (LoE). The CFAL can be planned using LoEs to arrange the cognitive process of identifying the qualification measurements for a national or NATO recruiter's selection to provide the right people with the right skills before assigning them to NATO civil & military leadership vacancies. The determination of LoEs will shape both the development of the CFAL and the execution of the selection process. In the CFAL planning terminology, the default approach for LoE demands one LoE per one CO to visualize how best to articulate a national and or NATO recruiter's cognitive vision.

1.2.2 Illustrative KPE

This type of KPE is an auxiliary, doctrine-compliant component used within the CFAL framework to visualize and link the cognitive measurement values. Their main function is to clarify the logical linkages amongst the knowledge dimensions and cognitive domains⁷ to achieve the CFAL end state

and the CFAL objectives. Moreover, this type of KPEs is the tools that materialize the abstract cognitive measurements of KPEs by establishing visual and logical connections compliant with NATO OPSA logic and existing NATO planning doctrine. Therefore, they are to facilitate setting the assessment metrics for the framework. In the CFAL planning terminology, Illustrative KPE encompasses CFAL phases, CFAL decisive condition, CFAL effect, and CFAL action as follows:

CFAL Phases⁸ (CP). The CFAL Phases (CP) define distinct stages within the cognitive process of identifying overall qualification measurements. These stages lead to the establishment of the Decisive Conditions (DCs) necessary for the successful accomplishment of the CFAL objective. Phases are sequential but may converge. In the CFAL planning terminology, a national or NATO recruiter's cognitive vision may guide a main effort in each phase to provide the right people with the right skills before assigning them to NATO civil & military leadership vacancies.

CFAL Decisive Condition⁹ (CDC). A decisive condition refers to a specific state or circumstance that, when achieved, significantly influences the ability to make effective decisions within varying degrees of information certainty and uncertainty. In the CFAL planning terminology, CDCs can also be utilized as unique modules to further develop the CFAL.

CFAL Effect¹⁰ (CE). An effect is a change to the engagement space as a consequence of an action. Derived from objectives, effects bridge the gap between objectives and actions by describing what changes in the engagement space are required. There are three types of effects: physical, virtual, and cognitive¹¹. Although all physical and virtual

4 Adopted from the ref. [3], pp. 91.

5 Adopted from ref. [3], pp. 91.

6 Adopted from ref. [1], pp. 4-63.

7 Knowledge dimensions and cognitive domains are basically described in the ref. [4].

8 Adopted from ref. [1], pp. 4-65.

9 Adopted from ref. [1], pp. 4-63 (see also its footnote #55).

10 Adopted from ref. [3], pp. 91-93.

11 Although the ref. [3] refers to three types of effects, this study specifically explores the cognitive dimension of the MDO beyond its physical and virtual dimensions, as stated in the abstract.

effects will lead to some form of cognitive effect, their primary purpose will be to influence the capabilities of audiences, while cognitive effects are principally directed towards changing or maintaining audiences' knowledge dimensions and cognitive domains. In the CFAL planning terminology, effects must be measurable.

CFAL Action¹² (CA). An action is described as the process of doing something to create an effect on a system within the engagement space. A sequence of coordinated assessment actions can be grouped into comprehensive evaluation processes. Actions are expressed as tasks or assignments and are imposed by an appropriate authority or derived during an analysis for the CFAL. In the CFAL planning terminology, the primary examples of actions are personality and motivational questionnaires, cognitive ability tests, interviews, 360-degree feedback, and simulation exercises. The author does not delve into further detail on this point due to the scope and constraints of this study¹³. In the CFAL planning terminology, actions must also be measurable.

1.2.3 *Reconstructive KPE*

This type of KPE is a feedback element within the CFAL framework to improve and adjust the framework itself and its associated assessment metrics, utilizing data obtained from the leadership qualification assessment process. Their main function is to support the analysis, interpretation, and recommendations for the CFAL by ensuring the adaptability of KPE to the current CFAL requirements and enabling the framework to self-calibrate against constantly evolving challenges. Moreover, this type

of KPEs is a dynamic components that continuously test and adapt not only the leader but also the CFAL framework itself, using real-time performance data. In the CFAL planning terminology, Reconstructive KPE encompasses the CFAL decision point, CFAL assumption, CFAL critical information requirement, and CFAL branch plans as follows:

CFAL Decision Points¹⁴ (CDP). A decision point is a cognitive marker defined in time, at which a national or NATO recruiter is expected to make a decision to ensure the timely selection and synchronized development of civil & military leaders, thereby ensuring the right people with the right skills are assigned to NATO civil & military leadership vacancies. However, CDPs are distinct from regular national and/or NATO data collection requirements (e.g., personality and motivational questionnaires, cognitive ability tests, interviews, 360-degree feedback, and simulation exercises). While these activities provide inputs, CDPs represent specific points where decisions are made based on the analyzed information. In the CFAL planning terminology, a CDP can be linked to assumptions and CCIRs.

CFAL Assumption¹⁵. There will be some gaps between information and knowledge at the time of the analysis of the situation or planning for the CFAL. In this regard, the author finds it necessary to make certain assumptions as a basis for further analysis or planning. In the CFAL planning terminology, an assumption is to be valid, and must be logical, and realistic. Therefore, assumptions necessitate regular review of the CFAL planning and analysis, especially if they are changed, added to, or deleted, impacting the CFAL's validity.

12 Adopted from the ref. [3], pp. 93.

13 The final attempt of this paper does not propose such actions in detail. Instead, the examples are presented to the reader just to stimulate the idea of actions in the CFAL planning terminology.

14 Adopted from the ref. [1], pp. 4-66. Moreover, a CDP should lead to the requirement for branches and sequels. At this stage of development of the CFAL, the author intends to leave the requirement for branches and sequels as notional terms. Therefore, the author encourages the reader to just think about them.

15 Adopted from the ref. [1], pp. 3-37 and 3-38. See also footnote #49 on pp. 3-37: "In planning, a supposition made about the current situation and/or the future course of events to complete an estimate of the situation and decide on the course of action. (AAP-06). Note: An assumption can equate to a potential risk and, as such, needs to be planned against."

CFAL Critical Information Requirement¹⁶ (CCIR).

Based on the initial analysis, the national or NATO staff should advise the national or NATO recruiter on critical information that may be required for future assignment of NATO civil & military leaders. At this stage, CCIRs should focus on cognitive qualification measurements for a national or NATO recruiter's selection to provide the right people with the right skills before assigning them to NATO civil & military leadership vacancies. In the CFAL planning terminology, such insights might lead or contribute to assisting the Alliance in needing leaders and staff who can cope with the future environment's speed, complexity, and uncertainty.

CFAL Branch Plan¹⁷. Branches are options within a particular phase of the CFAL, which are planned and executed in response to anticipated opportunity or reversal within that phase, to provide the national or NATO recruiters with the flexibility to retain the initiative and ultimately achieve the original CFAL objective. Branches address the question of what if?

1.2.4 Other Critical System(s) Elements

These elements are associated with certain system states and institutional requirements necessary for the successful integration of the CFAL. They are the integration points for KPE-generated data, ensuring that the results move beyond a mere report to actively trigger and modify all other relevant critical institutional systems. In the CFAL planning terminology, Other Critical System(s) Elements encompass risks, opportunities, pre-conditions for success, including the assessment metrics as follows:

CFAL Risk¹⁸. During the analysis, the recruiter will identify some risks to the achievement of the CFAL

objective that result from the engagement space of the key planning actors. In the CFAL planning terminology, a risk is portrayed accordingly based on the likelihood of its occurrence and the gravity of its impact. Moreover, a risk can be associated with the CFAL Critical Information Requirement (CCIR) as the risk management is owned by the national recruiter in close cooperation with concerning NATO authorities. Nonetheless, at this stage of the study, the risk management issue will be under discussion for further development of the CFAL.

Opportunities¹⁹. In conjunction with the risk analysis for achieving the CFAL objective, the national or NATO recruiter will also identify corresponding opportunities that may arise from that achievement. Opportunities, similar to risks, involve several critical dimensions: the likelihood of success or failure, the subsequent risk of a missed opportunity, the potential impact on selection processes (to acquire the right people with the right skills), and the requirement to pre-plan mitigation measures for both successful and unsuccessful outcomes. In the CFAL planning terminology, opportunities relate to controlled risk-taking over time within the Military Instrument of Power (MloP)²⁰ based on the MDO context, specifically addressing the cognitive dimension of MDO.

Preconditions for Success²¹ (PCS). It will define requirements that should be in place or achieved before NATO civil & military leaders are assigned to NATO civil & military leadership vacancies. In the CFAL planning terminology, PCS may have associated assumptions, risks, or CCIR that national or NATO recruiters can consider in their planning as part of the leadership assessment process.

16 Adopted from the ref. [1], pp. 4-14.

17 Inspired from the ref. [1] and especially interpreted through pp. B-5-3 (its footnote #28).

18 Adopted from the ref. [1], pp. 4-57 and its footnote #52.

19 Adopted from the ref. [1], pp. 4-57.

20 The ref. [5, pp. 11] defines one of the functions of the MloP, which is "Out-Pace", as it is: "The future Alliance Military Instrument of Power (MloP) must be able to recognize risks, seize opportunities, decide and act faster than potential adversaries."

21 Adopted from the ref. [1], pp. 3-28, footnote #36.

CFAL Measures of Effectiveness²² (C-MOE). A C-MOE is a metric used to assess the current cognitive state of a system/system element related to the overall CFAL results (effects, decisive conditions, objectives, end state). It helps answer the crucial question, “Are we on track to achieve the intended new system state within the planned timescale?” Each C-MOE describes a snapshot of the system state at one particular moment in time. Therefore, to fully understand how a system is changing, C-MOEs must be repeatedly measured across time to determine trends in system states. The C-MOE is an integral component of the CFAL Data Collection Plan (DCP) that is specifically based on the planning logic defined by the KPE.

CFAL Measures of Performance²³ (C-MOP). A C-MOP is a metric used to measure the accomplishment and quality of the activities in an overall CFAL design. It is primarily intended to answer the practical question, “Are the actions being executed as planned?” If, during execution, progress towards the achievement of desired effects is not made as expected, one possibility is that actions are not being carried out as planned or at the required standard. The C-MOP is also an integral part of the CFAL DCP that is based on the granular planning elements identified by the KPEs.

In conclusion, the author highlights a critical challenge to leadership assessment by referencing Sophocles (496-406 BC), who stated, “*how dreadful it is when the right judge judges wrong.*”²⁴ This challenge, therefore, forms the main rationale for the CFAL, which is designed as a process of identifying the cognitive qualification measurements for national or NATO recruiters’ selections to provide the right people with the right skills. In essence, the framework addresses the challenge articulated by Italian astronomer and physicist Galileo Galilei (1564-1642), to “measure what is measurable and make measurable that which is not”. This cognitive approach to assessment is precisely

what contributes to the Alliance’s ability to appoint leaders who can execute robust C2 while coping with the future environment’s complexity and uncertainty. Following this conclusion, a summary of propositions (P=) is as follows:

- P1=Predictive Focus: CFAL shifts the assessment focus from real-time operational reporting to predictive cognitive qualification measurement for future assignments.
- P2=Institutional Mechanism: CFAL is designed to function as an institutional mechanism that directly guides national/NATO recruiters on essential cognitive traits.
- P3=OPSA Alignment: The framework achieves robustness by integrating its methodology with the proven NATO Operations Assessment (OPSA) logic based on the NATO comprehensive operations planning view.

22 Adopted from the ref. [2], pp. 17-18.

23 Adopted from the ref. [2], pp. 18.

24 The author uses a compelling concluding quote as a powerful and memorable way to summarize the critical importance of the framework. It resonates with the need for accurate and effective leadership assessment.

2.0 ANALYTICAL METHODOLOGY

2.1.1 Scope and Limitations

This study contains four lenses in terms of its scope. First lens is of creating a framework for NATO civil & military leadership assessment as a critical function of C2 in the MDO context. The second lens is the cognitive dimension of MDO beyond its physical and virtual dimensions. The third lens is conceptual approach of the study, which is not typically an empirical study. In other words, this study neither attempts real-time data collection for leadership theories nor evaluates the results of any leadership assessment reports. Fourth lens partially discusses leadership styles,²⁵ acknowledging that there is no leadership style that is superlative than the other, and leadership qualities, etc.

Herein, the discussion draws upon both civil & military leadership literature and theories in terms of cultural clusters and leadership behaviour across NATO²⁶. Finally, this study aims to suggest CFAL in a way that the success of NATO civil & military leadership – as a vital element of adaptive C2 structures – needs to be predictably assessed through a cognitive approach. Therefore, the given framework will assist the Alliance in needing leaders and staff who can cope with the future

environment's speed, complexity, and uncertainty.

This study is primarily limited to a conceptual cognitive approach that will not be tested or validated within the paper because it is unprecedented. Nonetheless, the effectiveness of the cognitive framework should be evaluated by at least empirical studies in the future. A constraint for the study is time and word limitations in terms of structuring an academic paper. Because of that, some additional, deeper discussions, such as the terms of branches and sequels, are left out of the content of the study, and the readers are encouraged to think about more details. Moreover, there is a specific constraint that this study cannot introduce the details about the implementation of the actions of the cognitive framework for assessing leadership. The primary examples of actions are personality and motivational questionnaires, cognitive ability tests, interviews, 360-degree feedback, and simulation exercises that are beyond the scope of the study.

Therefore, the framework focuses more on what and why questions than how questions, which is left for following studies in the future. Finally, a one-size-fits-all approach to assessing leadership is not the only or ultimate solution for civil & military leadership assessment in NATO. However, this study can generate new insights and offer a

25 The reader may visit the ref. [6] in terms of a comparative analysis of different leadership styles. It is however up to the leader to make use of the appropriate style that will motivate his or her team members taking into consideration the characteristics of team members as pointed out in the ref. [7].

26 The author of this paper also intends to initiate an intensive discussion spinning around the terms of "cultural clusters", "leadership behaviour" and "the MDO across NATO". For that reason, the author significantly refers to the study of Peter G. Northouse, the ref. [8]. According to the given study, the author of this paper makes a limited scale analysis for Culture Clusters and Desired Leadership Behaviours with a focus of NATO core member countries. The purpose of the analysis is to express uncertainty without being unclear about the terms. Therefore, the author put a series of numerical qualifiers (1-3) in the intersections of "cultural clusters", and "leadership behaviours" as an appropriate means of ensuring an interpretation. In conclusion; The analysis of this paper's author implies that the most dominant desired leadership behaviour is charismatic / value-based leadership which is briefly explained as is in the aforementioned reference: "Charismatic/value-based leadership reflects the ability to inspire, to motivate, and to expect high performance from others based on strongly held core values. This kind of leadership includes being visionary, inspirational, self-sacrificing, trustworthy, decisive, and performance oriented." Moreover, the most dominant cultural clusters implied from the analysis are Anglo and Nordic Europe. The last finding that the analysis presents is about cultural dimensions classified by cultural clusters. In this regard, there may be implied that there are five dominant cultural dimensions across NATO in the MDO context, such as "future orientation; gender egalitarianism; institutional collectivism; performance orientation and uncertainty avoidance". Last but not least, the author of this paper encourages the readers to structure deeper analysis for studying on "cultural clusters", "leadership behaviour" and "the MDO across NATO".

more comprehensive understanding of leadership theories, requiring constant adaptation and development.

2.1.2 *Defining the Problem Space & the Central Research Question, and Working Hypotheses:*

The contemporary operational environment, labelled by the appearance of the MDO, demonstrates significant complexities for NATO civil & military leadership. While physical and virtual dimensions of the MDO concept are widely discussed across NATO, the author strongly argues that the cognitive dimension of leadership in the MDO often remains quasi-tangible²⁷ as well as quasi-intangible²⁸ in terms of systematic leadership predictive assessment. This leads to the author's Central Research Question: How should a CAFL be configured in the context of MDO? This question addresses two contrasting realities. Firstly, there remains a sparse NATO doctrine or standardized procedures for assessing leadership success at the Alliance level. This may partly emanate from the traditional requirement that "formal leader development is a national responsibility."²⁹ However, the concept of MDO accentuates that "Allies must reconsider building their recruitment, selection, training, promotion, and retention systems to deliver the right people with the right skills."³⁰

This highlights a vital disconnect: the strategic need for specific leadership capabilities in MDO versus a perceived gap in the Alliance-wide civil & military leadership predictive assessment framework. Consequently, this study postulates that a cognitive approach to leadership assessment is vital to ensure the resilience of the C2 chain by predictably

assessing future NATO leaders' ability to steer the speed, complexity, and uncertainty of the MDO environment. Therefore, this study adopts a conceptual approach by analysing relevant literature and insights from leadership theories. Besides, the cognitive framework focuses on knowledge dimensions and cognitive processes. From this point of view, the working hypothesis³¹ is structured as follows: NATO civil & military leadership would be predictably assessed through a cognitive approach, thereby enhancing C2 effectiveness and decision-making³² in the MDO context.

2.1.3 *Literature Review & Collecting Information*

The literature review, designed to address the CRQ, is structured around three essential tenets for collecting the required information:

The first tenet frames the problem space concerning the central research question. In this regard, the literature review is based on two axes to collect the required information. The first one is related to the well-known leadership theories that include: the core set of leadership capabilities, factors that shape specific leadership requirements, leadership performance, leadership capability framework, leadership outcomes, understanding of executive leadership, preferred leadership style, business leadership, military leadership, executive intelligence, etc. The second one is related to existing terms across NATO, specifically focusing on the evolution of C2 Models versus Decision-Making, the MDO Concept, and Operational Planning, etc. The second tenet is the identification of the required analysis type, such as descriptive and prescriptive, which also includes the psychology of analysis.

27 The author anonymously considers "Quasi Tangible" as follows: This refers to an abstract concept, value, or right that, while not inherently physical, can be manifested, represented, or evidenced through a tangible medium (e.g., a document, record, or physical output). Thus, the abstract core has a demonstrable physical counterpart.

28 The author anonymously considers "Quasi Intangible" as follows: This denotes an intangible value or asset that, though not physical or directly perceptible, is inextricably linked to or inferred from a tangible entity or its observable effects (e.g., a product, document, or the observable performance of a company). Here, the intangible core is made evident through its tangible associations.

29 The ref. [9, pp. 2-18] discusses the given national responsibility.

30 The ref. [10, pp. 16-17] points out the requirements of the right people with the right skills as one of the enablers of MDO.

31 The author considers the working hypothesis, as state in the ref. [11, pp. 1709]: "The working hypothesis is first and foremost a hypothesis or a statement of expectation that is tested in action. The term "working" suggests that these hypotheses are subject to change, are provisional and the possibility of finding contradictory evidence is real.". Additionally, the author defends the judgment as stated in the ref. [12, pp. 128] that: "A hypothesis is accepted only when it cannot be rejected".

32 This paper has no intent to discuss and assess decision-making styles of leadership in the MDO context.

The third tenet is the constructive wisdom for the CAFL, which is based on a quotation: *«To create is to recombine' said the molecular biologist François Jacob...Recombination is about cross-pollination, reaching across the problem space, and bringing together ideas that have never been connected before. We might call this rebel recombination³³»*. The author considers that this constructive wisdom will cross-pollinate across the problem space and bring together ideas that have never been connected before, so that the study answers the CRQ, embraces the MDO vision, and contributes to the C2 literature by integrating cognitive-level functions known as «Classify, Examine, Verify, and Compose³⁴».

2.2 A summary of the CEV&C Approach in the MDO context

Herein, the author presents an illustrated cognitive configuration as the overarching term, including the CEV&C Approach, and serves as a mnemonic device (schema) that helps externalize the new systemic approach for OPSA in the context of MDO. In addition to the cognitive formulation of the approach, the schema also presents some major procedural actions, that are also associated with the approach, which facilitate the analytical OPSA process in MDO. On this basis, the CEV&C Approach

as part of the cognitive configuration in the context of MDO is illustrated in Figure 2-1.³⁶

An important remark for this section concerns the interdependency between the first two columns of the Cognitive Configuration. The author acknowledges that these horizontal and vertical relationships are not expected to be solely linear, but rather more complex within NATO operations planning and execution. However, the author would strongly argue that the projected area of interest (P-AOI) linked to the respective MDO principles will build integrity of the Cognitive Configuration in inductive reasoning. Besides, this inductive reasoning will pave the way for generating the cognitive functions for OPSA in the context of MDO. The aim of the horizontally linear illustration of the expected interdependency between the first two columns is to simplify the visualization of the Cognitive Configuration³⁷. More importantly, this is not for creating distinctions amongst the MDO principles but for turning complexity into simplicity to construct the CEV&C Approach.

33 The ref. [13, pp. 135-136] introduces the logical linkages between to create and to recombine.

34 This approach was created for the SAS-192 Research Symposium titled "18th NATO Operations Research and Analysis (OR&A) Conference 2024". Thereafter it was first presented on 05 NOV 2024 at the Symposium. Additionally, the CEV&C Approach was partially presented at the annual conference organized by the NATO Command and Control Centre of Excellence (NATO C2COE) in close collaboration with the Canadian Joint Warfare Centre, on 21 NOV 2024. Finally, this study was published by NATO C2COE. See the ref. [14, pp. 54-70] for more information.

35 The author acknowledges that "for brevity within the definition, all the non-military Instruments of Power (IoP), commercial entities and other stakeholders are described as 'non-military', and the importance of 'synchronisation' with these entities is stated but not meant to imply any form of military primacy. The coordination and synchronisation of military and non-military activities will enable delivery of 'converging effects' as stated in the ref. [10, pp. 7]. Therefore, it is clear that this synchronisation also influences leadership requirements.

36 The colours of the parts of the pyramid and the gradient colours of the cognitive functions are relatively matching. The legends for colours, from red to green, are simply based on the achievement status of the development of OPSA cognitive functions over time. In addition to that, the 5th section of this paper is prepared for Acronyms & Abbreviations to help readers recall the terminology used in the study.

37 The evolutionary MDO principles (Unity, Agility, Interconnectivity, Creativity) are subsequently narrowed down by the projected area of interest of the study. The P-AOI are in turn: Collaboration, Data Exploitation, Shared Understanding, and Turning Complexity into Simplicity. In light of the P-AOI, the study applies the key principles to the CEV&C Approach to generate the cognitive functions of OPSA in the context of MDO.




Figure 2-1: The CEV&C Approach as part of Cognitive Configuration in the context of MDO.³⁵ Copyright Mehmet Akif MORKAYA.

2.3 Implementation of the CEV&C Approach for a CFAL Leadership

As it was stated earlier, the CFAL is a process of identifying the cognitive qualification measurements for a national or NATO recruiter's selection to provide the right people with the right skills before assigning them to NATO civil & military leadership vacancies. Consequently, the expected assignment needs to be predictably assessed through a cognitive approach, and more importantly, this predictive assessment contributes to assisting the Alliance in needing leaders and staff who can cope with the future environment's speed, complexity, and uncertainty. Therefore, this section aims to present an externalisation of functions of the CEV&C approach by integrating the MDO concept. And yet, the author consistently references the relevant OPSA Step (1-4) to provide a strong procedural anchor for the CFAL while also identifying the 'significant issue' for each function to highlight its core purpose and challenge.

2.3.1 Implementation of the Cognitive Function - CLASSIFY

This function is linked to the MDO principle

of "Unity," focusing on collaboration, and it is implemented as it corresponds to the OPSA Step-1 (Assessment design and support to planning). The significant issue for this function is carefully crafting the key planning elements of the CFAL and their assessment elements so that the CFAL is measurable. Therefore, it requires the Formulation of Measures of Effectiveness (MOE) / Measures of Performance (MOP) in a Data Collection Plan (DCP) for the CFAL. The expected collaborative effort requires engagement of concerned stakeholders from the academic realm of leadership theories, predominantly across the Alliance. This effort aims to assure the CFAL's measurability and its ability to assist the Alliance in identifying leaders and staff who can cope with the future environment's speed, complexity, and uncertainty.

2.3.2 Implementation of the Cognitive Function - EXAMINE

This function is linked to the MDO principle of "Agility" focusing on data exploitation and it is implemented as it corresponds to the OPSA Step-2 (Developing the DCP). The significant issue for this function is to ensure the integrated consistency of three major formulas; such as the CFAL, Allies'

Leadership Synchronized Assessment Matrix (ALSAM)³⁸, and DCP for the CFAL. This is vital for the predictive accuracy³⁹ of assessing leadership, which requires a systematic approach to deliver the right people with the right skills, in the concept of MDO.

2.3.3 *Implementation of the Cognitive Function – VERIFY*

This function is linked to the MDO principle of “Interconnectivity” focusing on shared understanding, and it is implemented as it corresponds to the OPSA Step-3 (Data collection and treatment). The significant issue for this function is to formulate conclusions that address any aspects of the CFAL CIRs. Moreover, when formulating conclusions, the following illustrative questions could also be considered in the MDO context for the given vacancy: What should we expect from military/civil leadership? Are the assumptions for the CFAL still valid? What are the critical risks and challenges to the success of the leadership? What are the critical opportunities to accelerate the success of the leadership? Are the civil or military leaders able to perform effectively as desired in their operational business? What are the critical changes or lack of change observed in their operational business? What are the most likely reasons for these changes or lack of change? In addition to that, this function will need to be supported by an AI tool to verify the interconnectivity among the CFAL’s MOEs and CCIRs.

2.3.4 *Implementation of the Cognitive Function – COMPOSE*

This function is linked to the MDO principle of “Creativity” focusing on turning complexity into simplicity, and it is implemented as it corresponds to the OPSA Step-4 (Analysis, interpretation, and recommendations). The significant issue for this function is the formulation of cognitive requirements of civil/military leadership effectiveness and skills in the MDO context that streamlines the Alliance’s leaders to cope with the future environment’s speed, complexity, and uncertainty. Moreover, when formulating recommendations that support the estimated type of decision for the CFAL, the following illustrative questions could also be considered within the MDO context: Update, change, add, or remove critical assumptions for the CFAL; Transition between phases of the CFAL; Reallocation of resources or new resource requirements to ensure the right people with the right skills are available; Adjustment of the CFAL elements such as effects, decisive conditions, decision points, objectives, and end state; Adjustment of preconditions for success for the CFAL; Change of the Alliance C2 arrangements; Adaptation of the CFAL, ALSAM and its DCP.

38 Allies’ Leadership Synchronized Assessment Matrix (ALSAM): This term is conceptually introduced by the author as a projected matrix to be engineered to amalgamate and to synchronize various leadership assessment components within the CFAL for the Alliance. At this stage, it is presented as notional to accelerate discussion about the CFAL and contribute to improving creativity for its further development. From this point of view, it can be inferred that this matrix can be considered a sort of management tool that may also require an organizational change across NATO. “The estimated change can be complex but not chaotic. Complexity can be managed, scheduled, and coordinated with management tools” as stated by Engelberg, in the ref. [15, pp.5]. In conclusion, the author acknowledges the argument that nothing is more important to an organisation’s long-term health and success than choosing the right people to lead it, as is foremost stated by the ref. [16].

39 At this stage, the author intends to initiate a discussion about the terms of “predictions”, “consistency”, “complexity and cognition” to shed light on how vital predictive accuracy is. On the one hand, it quotes that «when predictability is poor, inconsistency is destructive of any predictive validity. The research suggests a surprising conclusion: to maximize predictive accuracy, final decisions should be left to formulas, especially in low validity environments» as stated by the ref. [17, pp. 225]. On the other hand, it quotes that “like most other forms of automatic responding, consistency offers a shortcut through the complexities of modern life” which is followed by a subsidiary comment as follows: “Because of the increasing tendency for cognitive overload in our society, the prevalence of shortcut decision-making is likely to increase proportionality” as stated by the ref. [18, pp. 295 and 446].

2.4 Other Identical Issues for the CFAL⁴⁰

2.4.1 Knowledge Development

This section addresses the integral process of Knowledge Development (KD) for the CFAL. The author posits that the KD management process is vital to a well-designed framework for leadership, as it enables predictable assessment through a cognitive approach to deliver the right people with the right skills. This study, therefore, explores the KD Pyramid⁴¹, which is also known by other names⁴², a fundamental concept for comprehending the progression from raw data to actionable insights within the CFAL process. This section seeks to answer: What would be the most effective approach to develop the required knowledge for the CFAL? To address this, the application of the KD Pyramid is conceptually discussed, as follows.

This study illustrates the KD pyramid⁴³ from bottom to top: data, information, knowledge, wisdom, and shared understanding achieved at each level. At each level, as the reader moves up the pyramid from its base, the initial data for the CFAL is aggregated, which improves its utility in answering higher-level leadership cognitive assessment requirements, such as success, effectiveness, and performance, etc. Given that this study delineates two descriptive cognitive areas – Customary and Anticipatory – corresponding to OPSA in MDO, a brief description of KD's application for the CFAL via cognitive functions helps to explain the benefits generated for the leadership assessment framework:

- Cognitive Functions CLASSIFY & EXAMINE practiced by exploiting the “customary datasets⁴⁴” based on the levels of data, information, and knowledge of the KD Pyramid in the Customary Cognitive Areas. At this level of the pyramid, the CFAL's key planning elements are examined in the MDO context through the “Know What and Know-How” questions.
- Cognitive Functions VERIFY & COMPOSE practiced by exploiting the “anticipatory datasets⁴⁵” based on the levels of knowledge, wisdom, and shared understanding of the KD Pyramid in the Anticipatory Cognitive Area. At this level of the pyramid, CFAL's key planning elements are examined in the MDO context through the “Know-How and Know Why” questions.

Practically, from the assessment perspective, the best possible way to create a holistic understanding of both customary and anticipatory datasets is to effectively manage a CFAL Data Collection Plan, which is left notional for the paper at this stage of the study. And yet, the DCP is to describe and define the purpose of the quantifiable measurements used by a national and or NATO recruiter's selection to provide the right people with the right skills to NATO civil & military leadership vacancies.

40 Reminder: This section is partially adopted from the ref. [14]. Therefore, the author encourages the reader to revisit the given reference, in order to better understand this section herein.

41 The author detects “the importance of the intersection between the DIKW pyramid and artificial intelligence in the context of the information age” as stated in the ref. [19].

42 Some of them can be listed as follows: DIKW hierarchy, wisdom hierarchy, knowledge hierarchy, information hierarchy, information pyramid, etc.

43 Reminder: The illustration of the pyramid is seen in Figure 2-1. The given illustration also presents the importance of the KD pyramid within the cognitive configuration.

44 This study evaluates three “customary datasets”: The overall design for Cognitive Framework of Assessing Leadership, Allies' Leadership Synchronized Assessment Matrix, and CFAL Data Collection Plan.

45 This study evaluates two “anticipatory datasets”: Formulation of Conclusions and Formulation of Recommendations which are to be made by national and or NATO recruiters for providing the right people with the right skills to NATO civil and or military leadership vacancies. These formulations are currently left notional.

2.4.2 Cognitive Challenges for the CFAL

At this stage, the study highlights two crucial terms – unconscious bias and paradoxical awareness – as the cognitive challenges for the CFAL. These two challenges create a critical analytical dilemma for national and or NATO recruiters in establishing cognitive analysis models⁴⁶ for qualification measurements to deliver the right people with the right skills. Providing the right personnel for NATO civil and military leadership vacancies requires a predictable⁴⁷ assessed cognitive approach to effectively mitigate these challenges, thereby enhancing decision-making in the MDO context. In this regard, the given cognitive challenges are described as follows:

Unconscious bias⁴⁸: Building on Ikiseh B.N.'s (2021) emphasis, it is evident that "unconscious bias" can heavily influence the effective decision-making processes of organizational leaders and managers. This study identifies the given influence as a fundamental cognitive challenge for leadership assessment. This challenge necessitates a systematic and predictably assessed cognitive approach to effectively identify and mitigate such pervasive, non-objective factors.

Paradoxical awareness: In a dynamic and complex operating environment, leaders cannot possess comprehensive oversight of every detail. This reality necessitates maximizing the leadership potential entrenched within teams and organizations. However, this required decentralization is a significant cognitive challenge, demanding that leaders overcome two persistent paradoxes⁴⁹: first, balancing formal hierarchical leadership against collective process, and second, reconciling maintaining control (based on standard operating procedures) against enabling scope for flexibility and creativity. The author contends that successfully managing this duality is fundamentally dependent on the candidate leader's paradoxical awareness. Therefore, a systematic and predictably assessed cognitive approach will serve as a vital mitigation strategy for this inherent challenge⁵⁰.

46 For the sake of the given analysis, the author draws the reader attention onto the ref. [20, pp. xxii]: "Heuer is critical of the orthodox prescription of "more and better information" to remedy unsatisfactory analytic performance. He urges that greater attention be paid instead to more intensive exploitation of information already on hand and that in so doing, analysts continuously challenge and revise their mental models".

47 Moreover, the author recites as stated in the ref. [21, pp. 57 (footnote)]: «Predictions also create rich datasets, and are wonderfully amenable to mathematical analysis»

48 At this stage, the reader is kindly invited by the author to discuss unconscious bias across their organizations to support decision making. In this regard, the ref. [22] could be a good starting point to elaborate the discussion.

49 The ref. [23, pp. 42–57] intensively discuss "The Paradoxes of Future Military Leadership".

50 The author accepts that the given challenges are crucial, but may not be limited to. For example, the emergence of artificial intelligence (AI). The ref. [24, pp. 03-05] discusses how the advent of AI introduces several challenges. Another example of this issue could be C2 Agility, which is described by the ref. [25, 1-1] as follows: "Thus, agility was seen as an essential survival capability to deal adequately with complexity, dynamics, and uncertainty when addressing 21st century military challenges. C2-related failures have been attributed to many causes but are actually traceable to a systemic inability to adequately deal with complexity, dynamics, and uncertainty: in other words, an agility deficit." In addition to that its Chapter 4 more specifically discusses "Assessing C2 Agility". At this stage of the study, the author leaves this issue to the reader, for their further discussion, by citing another challenge for C2 in the ref [26, pp. 29-30]: "The Alliance Concept seeks collaborative, agile and empowered Command and Control. (NATO, 2023, p.15) The mechanism to achieve this is through the Supported Supporting-Interrelationships (SSI) component of Allied Doctrine, a specific relation between equal commanders established, defined, and if necessary arbitrated, by a common superior commander through an establishing directive developed during the planning process for specific missions, and a network of Liaison Officers within the command staffs [see (Prescott, 2022) for supporting doctrinal references but also further clarifications and examples of usage of SSI].

3.0 CONCLUSION

The present study addressed the critical need identified in the abstract: the lack of a predictably assessed cognitive approach for NATO civil and military leadership. The analysis was anchored by the core argument that success in the future operating environment – characterized by speed, complexity, and uncertainty – requires a structural solution. Consequently, the central proposition, which posited that complex leadership selection requirements can be institutionally structured through a cognitive framework, has been conclusively affirmed.

The paper's findings offer a definitive synthesis based on the assumption that culture clusters and desired/preferred leadership behaviours or styles are now categorically listed in NATO. Thereby, the paper's findings demonstrate the linkage between the foundational analysis of the culture cluster and the required preferred leadership style, which fundamentally mandates the establishment of the CFAL Critical Information Requirement as a strategic tool, which is also a vital element of the key planning elements. In other words, the CCIR can be considered as a strategic cognitive qualification assessment filter. This filter is essential for enabling NATO leaders to manage latent cognitive challenges, specifically those related to mitigating unconscious bias and maintaining paradoxical awareness.

This successful synthesis firmly establishes the key planning elements of the framework and articulates the CFAL's overall design as a viable institutional instrument. The chief contribution of this work is CFAL itself. This conceptual approach, which focuses on knowledge dimensions and cognitive processes, provides a predictable assessment mechanism to assist the Alliance in selecting leaders and staff who can cope with the future environment's speed, complexity, and uncertainty. For subsequent research, future efforts must be dedicated to developing and standardizing the precise cognitive metrics required for these assessments and validating the CFAL's overall design with its key planning elements through practice and simulation within established NATO operations planning principles.

5.0 ACRONYMS & ABBREVIATIONS

AI	Artificial Intelligence
ALSAM	Allies' Leadership Synchronized Assessment Matrix
CEV&E	Classify, Examine, Verify, and Compose
CFAL	Cognitive Framework of Assessing Leadership
DCP	Data Collection Plan
DM	Decision-Making
EBTA in 3D	Effect-Based Trend Analysis in 3 Dimensions (Virtual / Cognitive / Physical)
KPE	Key Planning Elements
MDO	Multi-Domain Operations
MDC2	MDO Command and Control
MOE	Measures of Effectiveness
MOP	Measures of Performance
NOAH	NATO Operations Assessment Handbook
OPSA	Operations Assessment
P-AOI	Projected Area of Interest
WH	Working Hypothesis

References

- [1] NATO Comprehensive Operations Planning Directive Version 3.1, 13 October 2023.
- [2] NATO Operations Assessment Handbook (NOAH) V.4.0, 05 January 2023.
- [3] NATO Standard AJP-01 Allied Joint Doctrine, Edition F Version 1, 19 December 2022.
- [4] Leslie Owen Wilson (2016, 2013, 2005, 2001), *Anderson and Krathwohl Bloom's Taxonomy Revised Understanding the New Version of Bloom's Taxonomy*. Internet (The latest access date: 03.06.26). https://quincycollge.edu/wp-content/uploads/Anderson-and-Krathwohl_Revised-Blooms-Taxonomy.pdf
- [5] NATO Warfighting Capstone Concept (NWCC). Internet (The latest access date: 03.06.26). <https://www.act.nato.int/wp-content/uploads/2023/06/NWCC-Glossy-18-MAY.pdf>
- [6] Maria Edwige dos Santos Pires, Sandra Maria dos Santos, Genivaldo Marcilio Frez, Marcelo Rabelo Henrique (2021). Leaders: a comparative analysis of the various leadership styles. *Research, Society and Development*, v. 10, n. 15. Internet (The latest access date: 03.06.26). DOI: <http://dx.doi.org/10.33448/rsd-v10i15.22908>.
- [7] Bans-Akutey, A., Tiimub, B.M. (2021). *The Effect of Leadership Styles on the Dynamics of a Group*. *Academia Letters*, Article 2061. Internet (The latest access date: 03.06.26). DOI: <https://doi.org/10.20935/AL2061>.
- [8] Peter G. Northouse (2016). *Leadership: Theory, and Practice*, Seventh Edition, *SAGE Publications*. Internet (The latest access date: 03.06.26): https://yourhomeworksolutions.com/wp-content/uploads/edd/2017/12/Leadership_theory_and_practice_7e____.pdf
- [9] NATO ATP-3.2.2 Command and Control of Allied Land Forces, April 2015.
- [10] NATO Alliance Concept for Multi-Domain Operations (MDO), 10 March 2023.
- [11] M. Casula, N. Rangarajan, P. Shields (2021). The potential of working hypotheses for deductive exploratory research. *Quality & Quantity* (2021) 55: 1703–1725. Internet (The latest access date: 03.06.26): DOI: <https://doi.org/10.1007/s11135-020-01072-9>
- [12] R. J. Heuer, Jr. (1999). *Psychology of Intelligence Analysis*. *Center for the Study of Intelligence*.
- [13] M. Syed (2019). *Rebel Ideas, The Power of Diverse Thinking*. *Publishers: John Murray*.
- [14] M.A. Morkaya (2025). A New Systemic Approach for OPSA: The Cognitive Formulation (CEV&C) in the Context of MDO. *NATO C2 COE, Annals of Command and Control, Vol 2, No 1*.
- [15] Engelberg, S. (2021). Decision-making, Innovation, and Organizational Change: The Need for New Paradigms. *Academia Letters*, Article 324. Internet (The latest access date: 03.06.26). DOI: <https://doi.org/10.20935/AL324>
- [16] Hudson (2019). A capability framework to identify leaders. *Hudson*. Internet (The latest access date: 03.06.26). A capability framework to identify leaders.
- [17] D. Kahneman (2012). *Thinking Fast and Slow* (Winner of the Nobel Prize). *Penguin Random House UK*.
- [18] R.B. Cialdini (2021). *Influence: The Psychology of Persuasion*. *Harper Business*.
- [19] D. Sampaio (2024). The Intersection of Artificial Intelligence and the DIKW Pyramid: A Comprehensive Analysis. Internet (The latest access date: 03.06.26): <https://www.linkedin.com/pulse/intersection-artificial-intelligence-dikw-pyramid-analysis-sampaio-ifxlf/>
- [20] Ikiseh, B.N. (2021). Unconscious Bias: The Unconscious Conscious Affect. *Academia Letters*, Article 2291. Internet (The latest access date: 03.06.26). DOI: <https://doi.org/10.20935/AL2291>

[21] Martijn W. van Eetveldt, Richard G. Oppelaar, and Peter Olsthoorn (2021). Leadership Undefined: The Paradoxes of Future Military Leadership. *Journal of Peace and War Studies, International Symposium of Military Academies (ISOMA) Special Edition: 42-57*. Internet (The latest access date: 03.06.26): (PDF) Leadership Undefined -The Paradoxes of Future Military Leadership - van Eetveldt et al - 2021

[22] Deppe C. and Schaal GS. (2024). Cognitive warfare: a conceptual analysis of the NATO ACT cognitive warfare exploratory concept. *Front. Big Data* 7:1452129. <https://doi.org/10.3389/fdata.2024.1452129>. Internet (The latest access date: 03.06.26): Cognitive warfare: a conceptual analysis of the NATO ACT cognitive warfare exploratory concept - PubMed

[23] STO Technical Report (2018). STO-TR-SAS-104 AC/323(SAS-104) TP/845, *C2 Agility: Next Steps*. Internet (The latest access date: 03.06.26): [https://publications.sto.nato.int/publications/STO%20Technical%20Reports/STO-TR-SAS-104/\\$\\$TR-SAS-104-ALL.pdf](https://publications.sto.nato.int/publications/STO%20Technical%20Reports/STO-TR-SAS-104/$$TR-SAS-104-ALL.pdf)

[24] A. Kalloniatis and C. Bowles (2024). Cyber and the Cube: More options for the C2 of multi-Domain. NATO C2 COE, *Annals of Command and Control, Vol 1, No 1*. Internet (The latest access date: 03.06.26): 20240614 Annals of C2 Vol 1 (2024-1).pdf

Author biographies

Mehmet Akif Morkaya

Mr. Morkaya is a retired Turkish Army OF-5. Between 1994 and 2019, he undertook a variety of national and international staff and command duties. In the last three years before his retirement, he served as the Deputy Chief of Staff Plans and the Political Advisor of HQ NRDC-Türkiye, simultaneously with his roles as Joint EVAL Team Lead and Joint Evaluator (Operations Assessment Subject Matter Expert-OPSA SME) for SHAPE. Since 2020, Mr. Morkaya's principal engagements have been as the Director of International Business Development at CNK Global Consultancy, the SME of Multi Domain Operations-Land supporting HQ JWC, and a Member/Mentor of the System Analysis and Studies Panel (SAS-192) for the Science & Technology Organization (STO). He is the author of the book "European Union Managerial Military Capacity and an Indefinite Negotiation Chapter: Common Foreign, Security, and Defense Policy," published in 2014. He is a native speaker of Turkish, fluent in English, has advanced knowledge of Spanish, French, and Italian, and possesses basic skills (writing and reading) in Russian.



MASKED SPEECH RECOGNITION AMONG MULTINATIONAL TALKER AND LISTENER COHORTS

Rebecca E Bieber, Matthew J Makashay, Benjamin M Sheffield, & Douglas S Brungart*



Objective. Large multinational operations such as those conducted by the North Atlantic Treaty Organization (NATO) require participants who come from a variety of language backgrounds to communicate with team members in a common language that they speak with varying levels of proficiency. The present study used operators at the NATO Trident Jupiter Exercise to evaluate the intelligibility of English sentences in noise as a function of the NATO Standard Agreement 6001 Language Proficiency Levels of the talkers and listeners.

Design. Stimuli included American English Matrix Test sentences produced by 20 NATO personnel, including 6 native and 14 non-native speakers of English, which were mixed in either multitalker babble or speech-shaped noise. Recognition accuracy and response times were recorded and were analyzed using mixed-effects regression modeling.

Study Sample. One hundred and seventeen personnel in attendance at a NATO joint exercise completed the study, including 28 native speakers of English and 89 native speakers of 11 other languages. The limitedly available demographic data were included in the statistical analyses to control for these variables to the extent possible.

Results. Talker and listener proficiency interacted with masker to predict recognition accuracy, with more prominent effects observed in babble than in noise. Sentences presented in babble also elicited longer response times than those in noise, but only when recognition accuracy was low.

Conclusion. Non-native talker and listener proficiency impacts communication dynamically with changes to masker type. Results are discussed in the context of the NATO standard language profile descriptors.

Keywords:

Speech in noise; non-native speech; speaking proficiency; listening proficiency

Introduction

In military operations, speech recognition in noise is critically important. However, recognition of speech in noise may be complicated by many features of the military environment including varying background noise types, multi-talker conversations, signal distortion in radio communications, or linguistic mismatch between talkers and listeners. Many complex multinational collaborative projects require participating personnel to understand speech produced by talkers from different language backgrounds than their own and, to complicate matters further, this speech may be spoken in a third language that both talker and listener learned after childhood. The impacts of linguistic mismatch on recognition performance may hold implications for multinational military environments, in which speech-based communications may be critical for accomplishing the mission.

English language proficiency and NATO

A prominent example of a multinational military alliance is the North Atlantic Treaty Organization (NATO). At the time of writing, NATO comprises 32 member countries, which have a range of official and commonly-spoken languages (NATO, 2024). The official languages of NATO are English and French. NATO personnel who are not native speakers of English may be required to take a test of their English language abilities, as certain occupational requirements may include English language proficiency. The term 'proficiency' can encompass many aspects of an individual's abilities in a non-native target language and assessments of proficiency can range widely (de Bruin, 2019; Harsch, 2014; Hulstijn, 2011). Measuring proficiency can be highly complicated, and there are additional factors to consider when testing language for specific uses such as military or other professional needs (Bullock & Westbrook, 2021; Green & Wall, 2005). The language proficiency scale explored in this study is NATO's Standardization Agreement 6001, or STANAG 6001 (NATO, 2014). The STANAG 6001 is an established standard to guide testing and rating of English proficiency that is shared among member nations (NATO, 2016).

STANAG 6001 is sponsored and maintained by the NATO Bureau of International Language Coordination (BILC), which serves as a resource regarding language teaching and language proficiency among participating member countries. The STANAG 6001 criteria serve as the standard that informs the development and implementation of language proficiency measures among NATO countries, and includes proficiency level recommendations for 4 domains of language skills: listening ("L"), speaking ("S"), reading ("R"), and writing ("W"). Following evaluation, each skill is assigned a proficiency level rating from 0 to 5, categorized as follows: 0 = No proficiency; 1 =

Survival; 2 = Functional; 3 = Professional; 4 = Expert; 5 = Highly-articulate native (here, termed "Native-Equivalent"). These language profiles are influential in determining eligibility for various job roles within NATO (Green & Wall, 2005; Skilleås & Grande, 2024). At present, each individual member country can design and implement its own language proficiency tests, though BILC provides recommendations and best practices for administering testing, scoring results, and aligning practices with other countries that also provide STANAG 6001 language profiles.

The recommended language proficiency testing does not include any measure of speech recognition in noise. In fact, the BILC explicitly states in their *Best Practices in STANAG 6001 Testing* document that "Listening texts [sic] should be clear without interference or background noise." (NATO BILC, 2023). However, the STANAG-6001 documentation describing listening proficiency makes only limited mention of communicating in acoustically degraded environments, suggesting that Level 5 (Native-Equivalent) listeners may be unimpacted by noise: *"Able to fully understand all forms and styles of speech intelligible to the well-educated native listener, including a number of regional dialects, highly colloquial speech, and language distorted by marked interference from other noise."* (NATO, 2016). While it is logical to test proficiency under optimum listening conditions, it is also important from a readiness standpoint to know how individuals can perform under acoustically degraded conditions, which are not uncommon in multinational operational settings, and which can interact with features of the talker and listener to impact communication outcomes.

The wording used in the STANAG 6001 level descriptors indicates that "native"-like performance is the standard, and implies that the goal of successful communication is with a native speaker (e.g., Level 3 [Professional]: *"....occasional errors*

1 The authors acknowledge that the concept of "nativeness" with respect to language experience can be complex and at times problematic (see Brown et al., 2023; de Bruin, 2019; Strand et al., 2024). However, this study references groups of "native" speakers and listeners of English, defined as personnel from English-speaking nations such as the United States and Canada who self-report native speaker status and who therefore may not be required to report STANAG 6001 language proficiency scores in order to fill occupational requirements. Results of this study should be interpreted with the understanding that there exists variability in language outcomes among these "native" talkers which may not be captured by the groupings used in this study.

in pronunciation, grammar or vocabulary are not serious enough to distort meaning, and rarely disturb the native speaker”; (NATO, 2016). However, BILC states that a proficiency level of 3 (Professional) is sufficient for most roles (BILC Working Group on Level 4 Proficiency, 2013). In this study, talkers and listeners are grouped based on STANAG 6001 Speaking and Listening levels, respectively, to examine patterns of speech recognition performance through the lens of this proficiency scale.

Impact of Noise in Military Operations

The presence of a masker during speech recognition will have an impact on speech recognition outcomes (Mattys et al., 2012). Noise can impact operational outcomes across a range of military contexts. For example, in a study of U.S. Navy Command and Control personnel in experimental tactical scenarios, Keller et al. (2017) observed reduced performance, poorer communication, and decreased situation awareness as participants experienced stimulated hearing loss in noise. Negative impacts of noise on military task performance have been documented by several others (Garinther et al., 1995; Whitaker, 1991; Whitaker et al., 1990). It is well established that speech-based informational maskers have a greater impact on recognition than noise-based energetic maskers; this effect is exacerbated for listeners who are not native speakers of the target language (Cooke et al., 2008; Lecumberri et al., 2010). There is evidence to suggest that, for bilingual listeners, proficiency in the non-native target language can influence susceptibility to informational masking effects, with higher proficiency providing a relative benefit for listening (Kilman et al., 2014; Warzybok et al., 2015).

Interplay of noise, language experience, and communication outcomes

There are several factors that could influence communication outcomes under degraded conditions in multinational operations. While bilingualism confers many benefits, bilingual adults do display reduced recognition (e.g., decreased accuracy, increased effort) for masked and/or acoustically degraded speech as compared to their monolingual peers (Krizman et al., 2017; Shi, 2012). This effect has been found even in bilingual adults

who are highly proficient or even dominant in the language that is being used for testing (Bieber, Phillips, et al., 2024). For example, Nakashima et al. (2018) found that bilingual adults who reported high fluency in English with respect to speaking, reading, and writing performed more poorly on a speech recognition task as compared to their monolingual peers when listening in military noise or via military communications equipment. Similarly, Giguère et al. (2020) found that lower English proficiency was associated with lower word recognition performance when listening over a radio.

While self-report metrics of proficiency have several inherent limitations (de Bruin, 2019), self-rated language proficiency has been documented as a moderating factor in predicting speech-in-noise outcomes for bilingual adults (Giguère et al., 2020; Nakashima et al., 2018; Shi, 2011, 2015). To the authors’ knowledge, proficiency measured with the STANAG 6001 scale has not been examined for its ability to predict masked speech recognition abilities. While the lack of standardization for specific testing creates some limitations for the STANAG 6001’s utility as a metric for comparison across countries, its relevance in terms of fitness for duty and occupational qualifications suggests that personnel are likely to have high accuracy for recall of their STANAG 6001 language profiles. For these practical reasons and implications, it is of interest to examine the relationship between STANAG 6001 scores and speech recognition outcomes in a military population.

NATO Research Task Group 285

The NATO Human Factors and Medicine Panel established Research Task Group 285 (HFM 285) in 2018, the objectives of which included the development of test materials to investigate the impact that different levels of talker and listener proficiency in English might have on the ability to communicate effectively in multinational military operations, with a particular emphasis on the interaction between English language proficiency and hearing loss (STO-Activities - HFM-285, 2024). Under the provision of HFM 285, research task group members developed a corpus of speech recordings produced by NATO personnel. These recordings were subsequently used as stimuli for studies of speech recognition (Bieber et al., 2023;

Bieber, Makashay, et al., 2024) among military service members and civilians.

In the present study, which was conducted during the development of the corpus, a cohort of NATO personnel in attendance at a joint exercise were recruited to both provide voice recordings to add to the corpus, and to complete a speech recognition task with the stimuli which had been created to date. Thus, both the stimulus talkers and the participants for the data reported herein were NATO personnel. Accuracy and response times corresponding to the speech recognition task were analyzed in order to answer the following research questions:

1. Do talkers' STANAG-Speaking and listeners' STANAG-Listening scores interact to predict masked speech recognition performance?
2. Do talkers' STANAG-Speaking and listeners' STANAG-Listening scores have a similar impact on masked speech recognition whether the masker is speech-based or noise-based?

This study is of particular relevance to Command and Control (C2), which impacts military operations at all levels. Effective communication underlies essentially all C2 functions, and is critical for mission success. Currently, objective standards in place for evaluating C2 procedures, training, and equipment are developed within NATO member nations (e.g., MIL-STD-1474E; Department of Defense, 2015) and may not account for any additional difficulties that may present during communication among individuals who do not share a native language. Understanding the interplay of complex factors (e.g., linguistic differences, masking of the speech signal) that can impact the accuracy and speed with which speech is recognized could ultimately inform the design of and decision-making in C2 centers.

Materials and Methods

Participants

The participants for this study included 117 operators who were recruited during the NATO Trident Jupiter joint exercise in November 2019 in Stavanger, Norway. These included 28 native speakers of English (Canada, Great Britain, United States), and 89 native speakers of other languages including Danish (Denmark), Dutch (Belgium, Netherlands), Estonian (Estonia), Finnish (Finland), French (Canada, France, Belgium), German (Germany), Greek (Greece), Hungarian (Hungary), Italian (Italy), Latvian (Latvia), Lithuanian (Lithuania), Norwegian (Norway), Polish (Poland), Romanian (Romania), Slovak (Slovakia), Spanish (Spain), Swedish (Finland, Sweden), and Turkish (Türkiye). Due to limitations in the type and amount of data that could be collected from exercise participants, participant age and sex were not queried, though years of military service was documented.

The remaining 89 participants who were not native speakers of English reported their STANAG 6001 English language proficiency scores. Limitations on data collection timing and resources precluded the collection of any additional metrics of English language proficiency. Twenty-eight participants had a STANAG-Listening score of 2 (Functional), 46 had a score of 3 (Professional), 12 had a score of 4 (Expert), and three had a score of 5 (Native-Equivalent). For analysis purposes, the listeners with scores of 4 and 5 were grouped together to create a group of $n=15$ (Expert/Native-Equivalent). Native speakers of English, who are not necessarily required to complete English language testing and thus may not have STANAG 6001 proficiency scores were grouped together for analysis ("Native"). Table 1 displays characteristics of the participants grouped by STANAG 6001 English listening score.

All participants provided voice recordings and completed a listening task. The voice recordings were processed and utilized as stimuli in subsequent research studies; further details of the recordings and subsequent stimuli can be found in prior documentation (Bieber et al., 2023; Bieber, Makashay, et al., 2024).

Stimuli

Stimuli for this study included American English Oldenburg Matrix Test (OMT) sentences (Kollmeier et al., 2015) recorded by both native and non-native speakers of English. In the present study, stimuli were produced by a total of 20 talkers, who had reported the following when queried for their native language and native region: Belgian Dutch (5 talkers), Netherlands Dutch (5 talkers), Canadian French (4 talkers), and Canadian English (6 talkers). These talkers were all NATO personnel who had been recorded at the Defence Research and Development Canada – Toronto Research Centre or from participants at a C2 meeting sponsored by the NATO Command and Control Centre of Excellence (Utrecht, NL). The recordings were made with an aviation headset with a close-talking microphone that was modified to work with the line-in recording microphone of an Android device (Google Nexus 7). For each phrase recorded, participants were shown a written version of the sentence and also played an example of the sentence spoken by a native English speaker. They were then asked to say the sentence into the microphone and press a button on the screen when they were done. Once this recording was completed, it was played back over the headphones and participants were given the opportunity to re-record the sentence if they wished. Otherwise, the recorded sentence was saved and the next recording prompt was played.

In addition to native country and language, each talker self-reported their proficiency for spoken English, using the STANAG 6001 English Language Proficiency Level scale. Talkers who were native speakers of English were labeled as such in a separate category (“Native”, $n = 6$). Of the non-native talkers, the majority had STANAG-Speaking scores of 3 (Professional, $n = 5$) or 4 (Expert, $n = 5$), with two talkers each having scores of 2 (Functional) or 5 (Native-Equivalent). A preliminary generalized linear regression modelling keyword recognition as a function of talkers’ STANAG speaking scores indicated no significant differences between the Functional and Professional talkers, nor between the Expert and Native-Equivalent talkers. Thus, for the full analysis, the non-native talkers were sorted into two categories: “Functional/Professional” (2 or 3) and “Expert/Native-Equivalent” (4 or 5), each comprising 7 talkers.

Each participant listened to a total of 16 sentences, including six practice trials and ten experimental trials. The practice trials were presented at a favorable signal-to-noise ratio (SNR) of +10 dB SNR and were produced by a female native speaker of American English who was not one of the speakers included in the experimental corpus. The experimental trials were mixed in two masker types (speech-shaped noise, 4-talker babble), and presented at four different signal-to-noise ratios (SNRs; noise: -9, -6 dB SNR; babble: -4, -1 dB SNR) for an overall level of 78 dB SPL. The multitalker babble included four talkers who were native speakers of English (3 male, 1 female) reading passages from the Wealth of Nations (Smith, 1975). These noise types and SNRs were selected in order to represent a range of listening environments and of difficulty levels that might be encountered by NATO personnel during the course of their operational duties. Different SNRs were selected for the two masker types due to the differential masking effects of a babble versus a noise masker. For each participant’s experimental trials, ten sentences were quasi-randomly selected from the overall corpus to allow for approximately 25% of sentences to be produced by a native speaker of English, and 75% to be produced by a non-native speaker of English. The distributions of talker, noise type and SNR were randomly represented within and across listeners. Across the listener group, each masker type/SNR level configuration represented approximately 25% of all stimuli heard. No sentences were repeated within-listener.

Figure 1 shows the overall intelligibility of the stimuli as a function of masker type and SNR for the participants in this study, separated by their status as native/nonnative speakers of English. Stimuli produced by non-native talkers are collapsed across talker proficiency group. Note that performance is roughly similar at the -4 (babble) and -6 (noise) dB SNR levels. A similar figure can be found in Bieber et al. (2024) which presented an expanded version of this stimulus set to a large cohort of United States service members.

Four participants (1 native, 3 non-native) had scores of under 70% on the practice trials at +10 dB SNR; their data in the masked conditions were excluded from the following analyses.

Procedure

Hearing thresholds. Pure tone hearing thresholds in each ear at 4 and 6 kHz were collected for each participant using a modified Hughson-Westlake technique on a Wireless Automated Hearing Test System (WAHTS) boothless audiometer (WAHTS Hearing LLC, Lebanon NH). The average of the two thresholds was calculated in each ear and

the average of the better-hearing ear was used for subsequent analysis.

Questionnaire. Questionnaire and listening task data collection was carried out using Android touch screen devices and TabSINT hearing research software (Shapiro et al., 2020). Each participant reported their years of military service, country of origin and native language, and self-rated proficiencies in English listening comprehension and speaking using the STANAG 6001 rating scale.

Listening task. Stimuli were presented binaurally via a Rugged Air supra-aural aviation headset. In each trial, listeners heard one sentence and were asked to identify what they heard. On the touch screen, the participants saw a 5 x 10 grid containing the full closed-set matrix of the 5-word sentences with 10 options per word. They were instructed to use the touch screen to select the words they heard, and submit their response. No correct answer feedback was provided. Each response was scored for the total number of words correctly identified. In addition, the TabSINT software logged the response time of the 'submit' button press for each trial in milliseconds.




Figure 1. Keyword identification performance as a function of talker nativeness, listener nativeness, noise type, and signal-to-noise ratio (SNR). Error bars show the 95% confidence interval around each data point. dB = decibel

Data Analysis

Data were analyzed using general(ized) linear mixed effects regression modelling. Two outcome measures were examined for the influence of listening proficiency (listeners) and speaking proficiency (stimulus talkers) on intelligibility scores: response accuracy and response time. For both outcomes, the *buildmer* package (Voeten, 2022) in R (v. 4.1.2; R Core Team, 2022) was used to establish the best fitting model. Response times were log-transformed for analysis. Continuous predictors including years of military service and better-ear hearing thresholds were centered and z-transformed before being entered into the models.

For response accuracy, the maximal model entered for evaluation comprised fixed effects including a four-way interaction of talkers' speaking proficiency (three levels, reference: Functional/Professional), listeners' listening proficiency (four levels, reference: Functional), masker type (babble/noise), and SNR level (high/low), and three-way interactions of talker and listener proficiency with high frequency hearing and with years of service. Lower-order effects and interactions were included for all of the above-mentioned evaluated effects. The maximal random effects structure included by-participant intercepts and slopes for speaking proficiency, by-token intercepts and slopes for high frequency hearing, and by-trial number interacting intercepts and slopes for speaking proficiency and listening proficiency.

For response times, the maximal model evaluated comprised fixed effects including three-way interactions of speaking proficiency, listening proficiency, and each of the following: high frequency hearing; masker type; years of service; and SNR level, and a four-way interaction of speaking proficiency, listening proficiency, masker type, and recognition score (6 levels, reference: 0), as well as all lower-order interactions. The maximal random effects structure included by-participant intercepts and slopes for speaking proficiency, by-token intercepts and slopes for high frequency hearing, and by-trial number interacting intercepts and slopes for speaking proficiency and listening proficiency.

Results

Speech Recognition Accuracy

The final model selected to describe the recognition accuracy data, shown in Table 2, included: a three-way interaction of speaking proficiency, listening proficiency, and masker type, a two-way interaction of listening proficiency and SNR level, and non-interacting fixed effects of high frequency hearing and years of service. The random effects structure contained by-subject and by-trial number random intercepts, and a by-token random intercept and slope for high frequency hearing. These interactions and effects are visualized and discussed in the following paragraphs.

The model revealed a significant three-way interaction of speaking proficiency, listening proficiency, and masker type. Examination of this three-way interaction revealed the following: regardless of masker type, listeners with the lowest listening proficiency rating of "Functional" (Figure 2, far left symbols in each panel) did not show any significant differences in performance as a function of talkers' speaking proficiency ($p > .05$, all comparisons).

For all other listener groups (Professional, Expert/Native-Equivalent, Native), there were differences in performance level as a function of talker proficiency, which differed depending on the type of masker. When comparing performance with stimuli produced by the two non-native talker groups (Expert/Native-Equivalent non-native talkers relative to the Functional/Professional talkers; blue circles and green squares, respectively), the two stimulus types result in similar performance when mixed in noise (right-hand panel; $p > .05$, all comparisons), but intelligibility of the Functional/Professional stimuli is significantly lower than that of the Expert/Native-Equivalent stimuli when mixed with multitalker babble (left-hand panel; $p < .05$, all comparisons). In both babble and noise, the stimuli produced by Expert/Native-Equivalent talkers resulted in similar performance to the natively-produced stimuli (blue circles and gray triangles, respectively; $p > .05$, all comparisons).




Figure 2. Model-predicted speech recognition scores as a function of the talkers' and listeners' English proficiency, and masker type. Error bars = 95% CI. Babble maskers were presented at -1 and -4 dB SNR; noise maskers were presented at -6 and -9 dB SNR.

Examining the differences across listener groups: in babble, the native listener group outperformed all non-native listener groups when listening to native stimuli or Expert/Native-Equivalent stimuli ($p < .05$, all comparisons). When listening to the Functional/Professional level stimuli, the native listener group outperformed the Expert/Native-Equivalent level listeners ($\beta = -1.08$, $SE = 0.34$, $t = -3.05$, $p < .01$), but not the Functional or Professional level listeners ($p > .05$, both comparisons), likely driven by the relatively lower performance of the native listener group when listening to the Functional/Professional talkers in babble. In noise, differences across groups were less pronounced. The native listener group outperformed all non-native groups when listening to the lowest proficiency Functional/Professional talkers ($p < .01$, all comparisons), but performed similarly to all non-native listeners in nearly all other talker conditions in noise, with only two exceptions (Native listeners performed better than Functional level listeners with a native talker: $\beta = -0.89$, $SE = 0.32$, $t = -2.79$, $p < .01$; Native listeners performed better than Expert/Native-Equivalent level listeners with an Expert/Native-Equivalent talker: $\beta = -0.69$, $SE = 0.33$, $t = -2.09$, p

$< .05$). In summary, non-native talker proficiency had a stronger impact on intelligibility of non-native English speech when listening in babble than when listening in noise.

The two-way interaction of SNR and listener proficiency, visualized in Figure 3, indicates that at the less favorable SNRs (orange triangles), there is an increase in performance with increasing listener proficiency for the lower proficiency non-native listeners (Functional vs Professional: $\beta = 0.77$, $SE = 0.3$, $t = 2.77$, $p < .0001$), and that native listeners outperform both the Functional ($\beta = -1.63$, $SE = 0.4$, $t = -4.05$, $p < .0001$) and Professional ($\beta = -0.86$, $se = 0.38$, $t = -2.26$, $p < .05$) listeners.

Listeners with Expert/Native-Equivalent listening proficiency scores do not show significantly different performance as compared to either the Professional-level listeners or the native listeners ($P > .05$, both comparisons). At the more favorable SNRs (gray circles), native listeners outperform all non-native listener groups ($P < .05$, all comparisons), but listeners with Functional and Professional proficiency scores do not differ from listeners with




Figure 3. Model-predicted speech recognition scores as a function of SNR level. Error bars = 95% CI

Expert/Native-Equivalent proficiency ($P > .05$, both comparisons).

The model also revealed a main effect of high frequency hearing sensitivity: listeners with poorer hearing thresholds showed poorer speech recognition scores ($\beta = 0.77$, $SE = 0.3$, $t = 2.77$, $p < .0001$). Lastly, the model showed a main effect of years of service: listeners with more years of service showed poorer speech recognition scores ($\beta = -0.15$, $SE = 0.06$, $t = -2.53$, $p < .05$). It is worth noting that participant ages were not recorded as a part of this data collection effort, but it is likely that age would correlate strongly with years of service: those with more years of service are likely older than those with fewer years of service.

In summary, both the English listening proficiency of the participant and the English speaking proficiency of the talker impact speech intelligibility, with talker proficiency effects emerging more prominently when stimuli are mixed with a babble masker as compared to a noise masker. Speech intelligibility is negatively impacted both by declines in listener high frequency hearing sensitivity and by increased years of military service.

Response Times

The final model selected to describe the response time data, shown in Table 3, included a two-way interaction of talkers' speaking proficiency and listeners' listening proficiency scores, a two-way interaction of masker type and recognition score, and a non-interacting fixed effect of SNR level. The random effects included by-subject, by-token, and by-trial number random intercepts. The interactions and effects are visualized and discussed in the following paragraphs.

The two-way interaction of speaking proficiency and listening proficiency scores, shown in Figure 4, is driven by the significantly slower response times displayed by low proficiency listeners (Functional) when listening to higher proficiency non-native talkers (Expert/Native-Equivalent, blue circles), as compared to low-proficiency non-native talkers (Functional/Professional, green squares; $\beta = -.15$, $SE = 0.05$, $t = -3.14$, $p < .01$), or native talkers (gray triangles; $\beta = -.16$, $SE = 0.05$, $T = -3.15$, $P < .01$). All other listener groups showed no significant differences in response time across talker types ($p > .05$, all comparisons).




Figure 4. Model-predicted response times, displayed as a function of talker and listener English proficiency. Error bars reflect the 95% CI.

Stimuli produced by native English talkers elicited faster response times in native listeners as compared to lower proficiency listeners (gray triangles, Native vs Professional listeners: $\beta = .13$, $SE = 0.06$, $t = 2.31$, $p < .05$; Native vs Functional listeners: $\beta = .17$, $SE = 0.06$, $T = 2.72$, $p < .01$), though response times to native talkers did not differ between native listeners and high-proficiency Expert/Native-Equivalent non-native listeners ($p > .05$). The listener groups did not differ from one another in their response times to low-proficiency stimuli (green square symbols, $p > .05$, all comparisons).

Examination of the two-way interaction between masker type and score (visualized in Figure 5) revealed that, at scores of 0, 1, and 2, there was a significant effect of masker type such that response times were slower for sentences presented in babble as compared to those presented in noise ($p < .05$, all comparisons). At scores of 3, 4 and 5, there was no significant effect of masker: response times did not differ between responses to sentences in babble vs those in noise ($p > .05$, all comparisons). In sum, masker type had a larger impact on response times at lower levels of recognition performance.

The model also revealed a significant main effect of relative SNR level, such that response times were slower when responding to sentences in more favorable SNRs as compared to the less favorable SNRs, relative to each masker type ($\beta = -.07$, $SE = 0.02$, $t = -3.1$, $p < .01$).

Discussion

The goal of this study was to examine English speech recognition in noise performance in a cohort of multinational military personnel, focusing on changes in keyword recognition scores and response times with different levels of talker and listener English language proficiency as indexed by their self-reported STANAG 6001 language proficiency levels. Two masker conditions were evaluated—speech-shaped noise, and multitalker babble—allowing for the examination of performance under the demands of energetic-only or both energetic and informational sources of masking. Of particular interest for the present study were the interactions of masker type, listener proficiency, and talker proficiency on response accuracy and response times.




Figure 5. Model-predicted response times as a function of masker type and keyword recognition score. Error bars = 95% CI.

Review of main findings

The analysis of accuracy scores showed that STANAG 6001 proficiency levels had different impacts on recognition when speech was presented in a noise masker as compared to a babble masker. One area in which this distinction was most notable was in the relative intelligibility of the non-native talkers. When speech was presented in a noise masker, scores from listening to non-native talkers with the lower English speaking proficiency (Functional/Professional) levels were similar to those from listening to non-native talkers who had higher English speaking proficiency (Expert/Native-Equivalent) levels. In babble, however, intelligibility scores from listening to the lower proficiency non-native talkers were significantly poorer than those from listening to the higher proficiency non-native talkers. For listeners with the lowest English listening proficiency scores, there were no significant differences in performance as a function of the stimulus talkers' spoken proficiency in either babble or noise.

The analysis of response times revealed no interaction of masker type with talker or listener proficiency, but there was an interaction of masker type with intelligibility score in predicting response

times: when response accuracy was lower, response times were longer for stimuli presented in babble as compared to noise, even though the SNR levels were higher in the babble. This finding is consistent with that observed by Bieber et al. (2024) in a cohort of United States service members listening to the same stimuli.

Masked speech recognition and language

One factor likely contributing to the differences in intelligibility and response times between the masker conditions is the presence of linguistic content in the babble masker. The differential impacts of an informational versus a primarily energetic form of masking have been well documented (Brungart et al., 2001; Cooke et al., 2008), with linguistic (mis)match between the language of target and masker speech, as well as their relation to the native language of the listener, modulating the degree of masker-related interference experienced by a listener (Brouwer et al., 2012; Calandruccio et al., 2013; Smith et al., 2024). In the present study, the interactions between proficiency and masker type could relate to differential weighting of cues and differences in the cognitive demands associated with listening under informational vs energetic sources of masking.

Mattys et al. (2009) suggest that when listening to speech in an energetic masker, a listener will depend relatively more heavily on acoustic detail. Informational sources of masking are understood to be associated with increased cognitive demands of listening (Koelewijn et al., 2012; Zhang et al., 2016). With these considerations, the present findings might suggest that any acoustic manifestations of 'non-nativeness' may have been relatively more salient and contributed to the higher proficiency non-native talker groups patterning more poorly than the native talkers when their speech was presented in noise, as compared to the babble conditions, in which higher-proficiency non-native talkers patterned similarly to the native talkers. While sentence-level stimuli do typically present a demand on higher-level cognitive processes (e.g., working memory), it is unlikely that listeners in this study would have relied as heavily upon lexico-semantic information to support recognition of these closed-set matrix sentences as they might with an open-set stimulus or with real-world discourse comprehension in which contextual information is available and will support recognition (Fereczkowski & Neher, 2023; Pichora-Fuller, 2008). While increased cognitive resources are necessary for extracting the stimuli from a babble masker, in noise the cognitive demand of listening is relatively lower. While additional work is needed in order to tease out the relative weighting of low-level and high-level cues in nonnative-on-nonnative speech recognition in a babble masker, the results of this study suggest that lower-level acoustic detail may have contributed more substantially to differences in intelligibility between talkers with different English speaking proficiency levels. It should also be noted that the SNRs used for testing differed across masker type, which may have contributed to this effect.

Language Proficiency and the STANAG 6001

In this study, stimulus talkers were grouped by self-reported STANAG 6001 speaking proficiency, and participants (i.e., listeners) were grouped by self-reported STANAG 6001 listening proficiency. The present study examines recognition of speech by both native and non-native speakers and listeners, which is a communication situation highly likely to occur during multinational operations, and indicates

that speaker proficiency as indexed by the STANAG 6001 profile may contribute differentially to intelligibility under masked recognition conditions.

Regarding spoken proficiency, the STANAG 6001 descriptors reference pronunciation as a factor contributing to proficiency level, with most levels described as involving some degree of non-native accent in their production of speech. An individual's accentedness in their non-native language in and of itself is not necessarily associated with the intelligibility of their speech in that language (Derwing & Munro, 1997; Munro & Derwing, 1995a, 1995b); there are a great many factors which contribute to individual intelligibility (for a review, see Chan, 2021). Other factors described as contributing to speaking proficiency in the STANAG 6001, such as fluency and grammatical structure, are less relevant to the recording of the stimuli for this study, but would also be expected to contribute to differences in intelligibility in real-world communication contexts (Baese-Berk & Morrill, 2015; Chau & Huensch, 2025).

The limited mention of acoustic degradation in the STANAG-6001 documentation suggests that listeners with the highest levels of listening proficiency should be unimpacted by factors such as noise or regional dialects. While this study did not explicitly include a comparison condition of recognition performance in quiet, the present findings indicate that, relative to a native listener, even a Level 5 (Native-Equivalent) non-native listener may experience greater speech recognition difficulties when listening in background noise, though the specifics of the masker will modulate this effect. Similarly, this study suggests that individuals with Level 5 listening proficiency may still experience reductions in English speech recognition accuracy when listening to a talker who is not a native speaker of the language.

The NATO BILC also states that a Level 3 (Professional) score is the highest level required for an individual to fulfill most job requirements (BILC Working Group on Level 4 Proficiency, 2013, p. 4). In this study we observed different patterns of recognition performance between Level 2 (Functional) and Level 3 (Professional). Listeners with Functional-level listening proficiency were

unimpacted by differences in talker proficiency, while listeners with Professional-level listening proficiency showed not only differential effects of proficiency as a function of masker type, but overall higher recognition of and faster response times when listening to native talkers as compared to the Functional-level listeners. Professional-level listeners patterned similarly with the Expert/Native-like listeners, supporting the BILC's description of listeners at this level. Unfortunately, the present sample of talkers only contained two talkers who were rated as Functional, limiting the ability to examine them as a separate group in this study.

In exploring and extending the characteristics of the STANAG 6001 standard language profile descriptors with respect to speech recognition in noise, this study provides information that may expand considerations for the C2 aspect of military operations. The success and timing of communications can be key for operational outcomes, but there are numerous ways through which communication can be impacted in these environments. The findings of this study allow for an expanded understanding of how the STANAG 6001 language profile descriptors may indicate the expected functional performance in English of personnel for whom English is not a native language, depending on the level and type of any background noise. However, it is worth noting that while the BILC language level descriptors are used for all NATO nations, the testing procedures used to determine STANAG 6001 profile levels vary from nation to nation. Therefore, there was likely some variability in proficiency levels within the groups as examined in the present study. Collectively, the study findings can provide a foundation for further detailed studies including standardized proficiency measures, which can be carried out with the goal of translating results into operational recommendations.

Limitations

To our knowledge, this study is novel in examining the impacts of both talker and listener proficiency in a shared non-native language on speech intelligibility. However, there are some limitations which should be acknowledged and may motivate future work in this area. A primary limitation is

the significantly constrained demographic detail available for the participants sampled in this study, which excluded their ages and any information regarding low and mid frequency hearing sensitivity, all factors which are known to contribute to speech recognition in noise (Dubno et al., 1984; Humes & Roberts, 1990). However, 4-6 kHz thresholds were available, which include the regions most likely to be impacted by both age-related and noise-related changes in sensitivity and which do contribute to speech recognition. Years of military service, which we have observed to correlate with age in a separate cohort of US military personnel tested by our lab, was included in the present analyses; additional study is required to elucidate any distinct and/or interacting contributions of age and years of military service on speech recognition in noise.

A second limitation of the study is that the noise conditions of the study were conducted at an unfavorable signal-to-noise level that was challenging for even the native participants and may have been at or near chance performance for many of the non-native participants. The unusual logistics of running the experiment during an active NATO exercise required the SNR levels to be set prior to the initiation of data collection, so no preliminary data was available to set these SNR levels. Although there was overlap in performance between the -4 dB SNR babble and -6 dB SNR noise conditions, it is possible that the inclusion of a noise conditions with a higher SNR would have revealed a larger difference between native and non-native talkers, similar to the one seen for the babble masker in Figure 2.

A further limitation to this study is that participants self-reported their STANAG 6001 proficiency scores and that STANAG 6001 language testing procedures vary across nations. Any further detailed language histories or proficiency measures standardized across nations were unavailable. Additionally, this study included no explicit rating of proficiency among the participants and talkers who reported English as their native language. To the authors' knowledge, there is no explicit waiver for language testing for personnel from English speaking nations. However, in descriptions for certain NATO occupational roles, requirements for STANAG 6001 scores are only listed for individuals who are

not native speakers of English, suggesting that in some cases native speaker status is considered as acceptable in lieu of testing. There can be substantial variation in language outcomes among “native” speakers of a language; the use of this category in the present study does limit findings in this sense while simultaneously reflecting what appears to be current practice in some cases. While self-report may be inherently more susceptible to variations in accuracy, it is likely that this population would have highly accurate recall for this value which holds critical importance for determining occupational qualifications. While the STANAG 6001 language profile certainly has many limitations, including inconsistent test specifics across member nations, its status as the metric used across nations to index language proficiency and determine eligibility for occupational specialties makes it one worth investigating. That said, the present results should be interpreted with this lack of test standardization in mind; the results are suggestive of expected performance levels per STANAG 6001 profile and cannot necessarily be generalized to be reflective of different proficiency metrics and levels.

The speech recognition task employed in this study, which involved manual identification of words heard in the context of closed-set, grammatically constrained sentences, should also be considered when interpreting the findings. The nature of the task is such that individual factors such as vision, manual dexterity, and working memory likely influence response accuracy and timing. The use of a closed-set task, while beneficial for avoiding tester bias or error in scoring responses, minimizes other factors such as vocabulary knowledge which can contribute to recognition ability in more naturalistic tasks. Therefore, the findings seen here likely represent the minimal expected effects of STANAG 6001 proficiency on recognition, with group differences expected to be increased in open-set or more linguistically complex listening tasks.

Conclusion

In conclusion, a sample of NATO operators from multiple member countries demonstrated higher recognition accuracy for sentences produced by native talkers as compared to non-native talkers. The talkers’ STANAG 6001 speaking score had a stronger impact on talker intelligibility when sentences were presented in multitalker babble than when they were when presented in noise; with mid-level proficiency talkers patterning similar to native talkers in babble, but similar to low proficiency non-native talkers in noise. An analysis of response times to the matrix sentences showed that, at lower intelligibility levels, masker type had a stronger impact on response times than at higher intelligibility levels. The findings of this study generally align with the STANAG 6001 descriptors laid out by NATO BILC and extend them to provide information regarding patterns of masked speech recognition performance.

References

- ATrainP-5 Language Proficiency Levels, Edition A Version 2 (Version 2). (2016). <https://www.natobilc.org/files/ATrainP-5%20EDA%20V2%20E.pdf>
- Baese-Berk, M. M., Bradlow, A. R., & Wright, B. A. (2013). Accent-independent adaptation to foreign accented speech. *The Journal of the Acoustical Society of America*, 133(3), EL174–EL180. <https://doi.org/10.1121/1.4789864>
- Baese-Berk, M. M., McLaughlin, D. J., & McGowan, K. B. (2020). Perception of non-native speech. *Language and Linguistics Compass*, 14(7). <https://doi.org/10.1111/lnc3.12375>
- Baese-Berk, M. M., & Morrill, T. H. (2015). Speaking rate consistency in native and non-native speakers of English. *The Journal of the Acoustical Society of America*, 138(3), EL223–EL228. <https://doi.org/10.1121/1.4929622>
- Bieber, R. E., Makashay, M. J., Sheffield, B. M., & Brungart, D. S. (2024). Intelligibility of Natively and Nonnatively Produced English Speech Presented in Noise to a Large Cohort of United States Service Members. *Journal of Speech, Language, and Hearing Research*, 1–19. https://doi.org/10.1044/2024_JSLHR-23-00312
- Bieber, R. E., Makashay, M. J., Simpson, B., Sheffield, B. M., & Brungart, D. S. (2023). Short-term retention of learning after rapid adaptation to native and non-native speech. *The Journal of the Acoustical Society of America*, 153(6), 3362. <https://doi.org/10.1121/10.0019749>
- Bieber, R. E., Phillips, I., Ellis, G. M., & Brungart, D. S. (2024). Current age and language use impact speech-in-noise differently for monolingual and bilingual adults. *OSF*. <https://doi.org/10.31234/osf.io/7jxvc>
- Brouwer, S., Van Engen, K. J., Calandruccio, L., & Bradlow, A. R. (2012). Linguistic contributions to speech-on-speech masking for native and non-native listeners: Language familiarity and semantic content. *The Journal of the Acoustical Society of America*, 131(2), 1449–1464. <https://doi.org/10.1121/1.3675943>
- Brown, V. A., McLaughlin, D. J., Strand, J. F., & Van Engen, K. J. (2020). Rapid adaptation to fully intelligible nonnative-accented speech reduces listening effort. *Quarterly Journal of Experimental Psychology*, 174702182091672. <https://doi.org/10.1177/1747021820916726>
- Brungart, D. S., Simpson, B. D., Ericson, M. A., & Scott, K. R. (2001). Informational and energetic masking effects in the perception of multiple simultaneous talkers. *The Journal of the Acoustical Society of America*, 110(5), 2527–2538. <https://doi.org/10.1121/1.1408946>
- Bullock, N., & Westbrook, C. (2021). Testing in ESP: Approaches and Challenges in Aviation and Maritime English. In B. Lantaigne, C. Coombe, & J. D. Brown (Eds.), *Challenges in Language Testing Around the World* (pp. 67–77). Springer Singapore. https://doi.org/10.1007/978-981-33-4232-3_7
- Calandruccio, L., Brouwer, S., Van Engen, K. J., Dhar, S., & Bradlow, A. R. (2013). Masking release due to linguistic and phonetic dissimilarity between the target and masker speech. *American Journal of Audiology*, 22(1), 157–164. [https://doi.org/10.1044/1059-0889\(2013/12-0072\)](https://doi.org/10.1044/1059-0889(2013/12-0072))
- Chan, V. (2021). Factors Influencing Intelligibility and Comprehensibility: A Critical Review of Research on Second Language English Speakers. 12.
- Chau, T., & Huensch, A. (2025). The relationships among L2 fluency, intelligibility, comprehensibility, and accentedness: A meta-analysis. *Studies in Second Language Acquisition*, 47(1), 282–307. <https://doi.org/10.1017/S0272263125000014>
- Clarke, C. M., & Garrett, M. F. (2004). Rapid adaptation to foreign-accented English. *The Journal of the Acoustical Society of America*, 116(6), 3647–3658. <https://doi.org/10.1121/1.1815131>
- Cooke, M., Garcia Lecumberri, M. L., & Barker, J. (2008). The foreign language cocktail party problem: Energetic and informational masking effects in non-native speech perception. *The Journal of the Acoustical Society of America*, 123(1), 414–427. <https://doi.org/10.1121/1.2804952>

- de Bruin, A. (2019). Not All Bilinguals Are the Same: A Call for More Detailed Assessments and Descriptions of Bilingual Experiences. *Behavioral Sciences*, 9(3), Article 3. <https://doi.org/10.3390/bs9030033>
- Department of Defense. (2015). Design Criteria Standard Noise Limits (No. MIL-STD-1474E).
- Derwing, T. M., & Munro, M. J. (1997). Accent, Intelligibility, and Comprehensibility. Evidence from Four L1s. *Studies in Second Language Acquisition*, 19(1), 1–16.
- Dubno, J. R., Dirks, D. D., & Morgan, D. E. (1984). Effects of age and mild hearing loss on speech recognition in noise. 76(July 1984), 87–96.
- Fereczkowski, M., & Neher, T. (2023). Semantic Context Can Mask Intelligibility Declines at Above-Conversational Speech Levels in Normal-Hearing Listeners. *Journal of Speech, Language, and Hearing Research*, 66(6), 2177–2183. https://doi.org/10.1044/2023_JSLHR-22-00506
- Garinther, G., Whitaker, L., & Peters, L. (1995). The effects of speech intelligibility on military performance. *Proceedings of the Symposium on Speech Communication Metrics and Human Performance*, 72–82.
- Giguère, C., Vaillancourt, V., & Laroche, C. (2020). Effects of hearing loss and language proficiency on speech intelligibility over radio transmission with tactical communication devices. *International Journal of Audiology*, 59(sup1), S31–S39. <https://doi.org/10.1080/14992027.2019.1682691>
- Green, R., & Wall, D. (2005). Language testing in the military: Problems, politics and progress. *Language Testing*, 22(3), 379–398. <https://doi.org/10.1191/0265532205lt314oa>
- Harsch, C. (2014). General Language Proficiency Revisited: Current and Future Issues. *Language Assessment Quarterly*, 11(2), 152–169. <https://doi.org/10.1080/15434303.2014.902059>
- Hulstijn, J. H. (2011). Language Proficiency in Native and Nonnative Speakers: An Agenda for Research and Suggestions for Second-Language Assessment. *Language Assessment Quarterly*, 8(3), 229–249. <https://doi.org/10.1080/15434303.2011.565844>
- Humes, L. E., & Roberts, L. (1990). Speech-Recognition Difficulties of the Hearing-Impaired Elderly. *Journal of Speech, Language, and Hearing Research*, 33(4), 726–735. <https://doi.org/10.1044/jshr.3304.726>
- Keller, M. D., Ziriach, J. M., Barns, W., Sheffield, B., Brungart, D., Thomas, T., Jaeger, B., & Yankaskas, K. (2017). Performance in noise: Impact of reduced speech intelligibility on Sailor performance in a Navy command and control environment. *Hearing Research, Noise in the Military*, 349, 55–66. <https://doi.org/10.1016/j.heares.2016.10.007>
- Kilman, L., Zekveld, A., Hällgren, M., & Rönnerberg, J. (2014). The influence of non-native language proficiency on speech perception performance. *Frontiers in Psychology*, 5. <https://doi.org/10.3389/fpsyg.2014.00651>
- Koelewijn, T., Zekveld, A. A., Festen, J. M., Rönnerberg, J., & Kramer, S. E. (2012). Processing Load Induced by Informational Masking Is Related to Linguistic Abilities. *International Journal of Otolaryngology*, 2012(1), 865731. <https://doi.org/10.1155/2012/865731>
- Kollmeier, B., Warzybok, A., Hochmuth, S., Zokoll, M. A., Uslar, V., Brand, T., & Wagener, K. C. (2015). The multilingual matrix test: Principles, applications, and comparison across languages: A review. *International Journal of Audiology*, 54, 3–16. <https://doi.org/10.3109/14992027.2015.1020971>
- Krizman, J., Bradlow, A. R., Lam, S. S.-Y., & Kraus, N. (2017). How bilinguals listen in noise: Linguistic and non-linguistic factors. *Bilingualism: Language and Cognition*, 20(4), 834–843. <https://doi.org/10.1017/S1366728916000444>
- Lecumberri, M. L. G., Cooke, M., & Cutler, A. (2010). Non-native speech perception in adverse conditions: A review. *Speech Communication*, 52(11), 864–886. <https://doi.org/10.1016/j.specom.2010.08.014>

- Mattys, S. L., Brooks, J., & Cooke, M. (2009). Recognizing speech under a processing load: Dissociating energetic from informational factors. *Cognitive Psychology*, 59(3), 203–243. <https://doi.org/10.1016/j.cogpsych.2009.04.001>
- Mattys, S. L., Davis, M. H., Bradlow, A. R., & Scott, S. K. (2012). Speech recognition in adverse conditions: A review. *Language and Cognitive Processes*, 27(7–8), 953–978. <https://doi.org/10.1080/01690965.2012.705006>
- Munro, M. J., & Derwing, T. M. (1995a). Foreign Accent, Comprehensibility, and Intelligibility in the Speech of Second Language Learners. *Language Learning*, 45(1), 73–97. <https://doi.org/10.1111/j.1467-1770.1995.tb00963.x>
- Munro, M. J., & Derwing, T. M. (1995b). Processing Time, Accent, and Comprehensibility in the Perception of Native and Foreign-Accented Speech. *Language and Speech*, 38(3), 289–306. <https://doi.org/10.1177/002383099503800305>
- Nakashima, A., Abel, S. M., & Smith, I. (2018). Communication in military environments: Influence of noise, hearing protection and language proficiency. *Applied Acoustics*, 131, 38–44. <https://doi.org/10.1016/j.apacoust.2017.10.010>
- NATO BILC. (2023, January). Best Practices in STANAG 6001 Testing. <https://www.natobilc.org/Best%20Practices%20in%20STANAG%206001.pdf>
- Phillips, I., Bieber, R. E., Dirks, C., Grant, K. W., & Brungart, D. S. (2024). Age Impacts Speech-in-Noise Recognition Differently for Nonnative and Native Listeners. *Journal of Speech, Language, and Hearing Research*, 67(5), 1602–1623. https://doi.org/10.1044/2024_JSLHR-23-00470
- Pichora-Fuller, M. K. (2008). Use of supportive context by younger and older adult listeners: Balancing bottom-up and top-down information processing. *International Journal of Audiology* (ISSNOnline) *Journal International Journal of Audiology*, 47, 1499–2027. <https://doi.org/10.1080/14992020802307404>
- Shapiro, M. L., Norris, J. A., Wilbur, J. C., Brungart, D. S., & Clavier, O. H. (2020). TabSINT: Open-source mobile software for distributed studies of hearing. *International Journal of Audiology*, 59(sup1), S12–S19. <https://doi.org/10.1080/14992027.2019.1698776>
- Shi, L.-F. (2011). How “proficient” is proficient? Subjective proficiency as a predictor of bilingual listeners’ recognition of English words. *American Journal of Audiology*, 20(1), 19–32. [https://doi.org/10.1044/1059-0889\(2011/10-0013\)](https://doi.org/10.1044/1059-0889(2011/10-0013))
- Shi, L.-F. (2012). Contribution of linguistic variables to bilingual listeners’ perception of degraded English sentences. *Journal of Speech, Language, and Hearing Research: JSLHR*, 55(1), 219–234. [https://doi.org/10.1044/1092-4388\(2011/10-0240\)](https://doi.org/10.1044/1092-4388(2011/10-0240))
- Shi, L.-F. (2015). How “Proficient” Is Proficient? Bilingual Listeners’ Recognition of English Words in Noise. *American Journal of Audiology*, 24(1), 53–65. https://doi.org/10.1044/2014_AJA-14-0041
- Skilleås, H. K., & Grande, B. (2024). Language Lessons Learned? NATO Requirements and Norwegian Officer Cadets. *Scandinavian Journal of Military Studies*, 7(1). <https://doi.org/10.31374/sjms.221>
- Smith, E. D., Holt, L. L., & Dick, F. (2024). A one-man bilingual cocktail party: Linguistic and non-linguistic effects on bilinguals’ speech recognition in Mandarin and English. *Cognitive Research: Principles and Implications*, 9(1), 35. <https://doi.org/10.1186/s41235-024-00562-w>
- STANAG 6001 Language Proficiency Levels Edition 5. (2014). <https://www.natobilc.org/files/file/6001EFed05.pdf>
- STO-Activities—HFM-285. (2024). NATO Science & Technology Organization.
- Voeten, C. (2022). buildmer: Stepwise elimination and term reordering for mixed-effects regression (Version 2.7) [Computer software]. <https://CRAN.R-project.org/package=buildmer>

Warzybok, A., Brand, T., Wagener, K. C., & Kollmeier, B. (2015). How much does language proficiency by non-native listeners influence speech audiometric tests in noise? *International Journal of Audiology*, 54(sup2), 88–99. <https://doi.org/10.3109/14992027.2015.1063715>

Whitaker, L. A. (1991). Performance as a Function of Communication in Military Vehicle Simulators. *Proceedings of the Human Factors Society Annual Meeting*, 35(10), 614–617. <https://doi.org/10.1177/154193129103501003>

Whitaker, L. A., Peters, L., & Garinther, G. (1990). Effects of Speech Intelligibility among Bradley Fighting Vehicle Crew Members: Simnet Performance and Subjective Workload. *Proceedings of the Human Factors Society Annual Meeting*, 34(3), 186–188. <https://doi.org/10.1177/154193129003400301>

Zhang, C., Arnott, S. R., Rabaglia, C., Avivi-Reich, M., Qi, J., Wu, X., Li, L., & Schneider, B. A. (2016). Attentional modulation of informational masking on early cortical representations of speech signals. *Hearing Research*, 331, 119–130.

Voeten, C. (2022). *buildmer: Stepwise elimination and term reordering for mixed-effects regression (Version 2.7)* [Computer software]. <https://CRAN.R-project.org/package=buildmer>

Warzybok, A., Brand, T., Wagener, K. C., & Kollmeier, B. (2015). How much does language proficiency by non-native listeners influence speech audiometric tests in noise? *International Journal of Audiology*, 54(sup2), 88–99. <https://doi.org/10.3109/14992027.2015.1063715>

Whitaker, L. A. (1991). Performance as a Function of Communication in Military Vehicle Simulators. *Proceedings of the Human Factors Society Annual Meeting*, 35(10), 614–617. <https://doi.org/10.1177/154193129103501003>

Whitaker, L. A., Peters, L., & Garinther, G. (1990). Effects of Speech Intelligibility among Bradley Fighting Vehicle Crew Members: Simnet Performance and Subjective Workload. *Proceedings of the Human Factors Society Annual Meeting*, 34(3), 186–188. <https://doi.org/10.1177/154193129003400301>

Zhang, C., Arnott, S. R., Rabaglia, C., Avivi-Reich, M., Qi, J., Wu, X., Li, L., & Schneider, B. A. (2016). Attentional modulation of informational masking on early cortical representations of speech signals. *Hearing Research*, 331, 119–130. <https://doi.org/10.1016/j.heares.2015.11.002>

Author's note

The identification of specific products or scientific instrumentation is considered an integral part of the scientific endeavor and does not constitute endorsement or implied endorsement on the part of the author(s), Department of War, or any component agency. The views expressed in this article are those of the author(s) and do not necessarily reflect the official policy of the Department of War, Defense Health Agency, the U.S. Government, or the Henry M. Jackson Foundation for the Advancement of Military Medicine, Inc. Approved for public release, distribution unlimited.

This work was supported by the Office of the Assistant Secretary of War for Health Affairs through the Clinical and Rehabilitative Medicine Research Program under Award W81XWH-18-2-0014 (Douglas S. Brungart).

Acknowledgments

The authors would like to acknowledge Hector Galloza for his work in developing the data collection protocol, as well as Ann Nakashima and Senne Goris for their roles in the stimulus recording procedures. The authors would also like to thank Commander Christian Johanssen for facilitating access to North Atlantic Treaty Organization operators and Nandini Iyer and Frederick Gregory for assisting with data collection on the experiment.

Table 1. Participant demographics

Participant Group (STANAG 6001: n Listening)		Years of military service Mean (sd)	Average of 4-6 kHz in the better ear; dB HL Mean (sd)
2 (Functional)	28	17.8 (9.35)	16.7 (11.0)
3 (Professional)	46	20.5 (9.56)	20.1 (13.3)
4 or 5 (Expert or Native-Equivalent)	15	21.7 (11.1)	17.2 (9.02)
Native	28	16.4 (9.04)	15.2 (9.23)

Table 2. Speech Recognition Accuracy. Model: Accuracy ~ 1+ betterPTA_46_z + Yearsserv_z + (STANAG_L * SNR_highlow) + (STANAG_S * STANAG_L * MaskerType) + (1 | Participant) + (1 + betterPTA_46_z | item) + (1 | trialNumber). Reference levels: MaskerType = Babble, STANAG_S = Native, STANAG_L = Native, SNR = High. Number of obs.: 1130, Participants: 113, Items: 274, Trials = 10.

Predictors	Odds Ratios	std. Error	Statistic	p
(Intercept)	34.01	13.51	8.88	<0.001
MaskerType [N]	0.08	0.03	-5.87	<0.001
snr highlow [Low]	0.17	0.03	-9.01	<0.001
STANAG_S [2_or_3]	0.14	0.06	-4.67	<0.001
STANAG_S [4_or_5]	0.48	0.21	-1.72	0.086
STANAG_L [2]	0.11	0.05	-5.17	<0.001
STANAG_L [3]	0.29	0.12	-3.05	0.002
STANAG_L [4_or_5]	0.24	0.12	-2.90	0.004
betterPTA 46 z	0.77	0.06	-3.41	0.001
yearsserv z	0.86	0.05	-2.53	0.012
snr highlow [Low] × STANAG_L [2]	1.72	0.38	2.43	0.015
snr highlow [Low] × STANAG_L [3]	1.46	0.29	1.87	0.062
snr highlow [Low] × STANAG_L [4_or_5]	2.44	0.62	3.54	<0.001
STANAG_S [2_or_3] × STANAG_L [2]	4.98	2.22	3.59	<0.001
STANAG_S [4_or_5] × STANAG_L [2]	2.63	1.20	2.13	0.033
STANAG_S [2_or_3] × STANAG_L [3]	2.91	1.22	2.54	0.011
STANAG_S [4_or_5] × STANAG_L [3]	1.46	0.63	0.88	0.376
STANAG_S [2_or_3] × STANAG_L [4_or_5]	1.51	0.79	0.79	0.432

Predictors	Odds Ratios	std. Error	Statistic	p
STANAG_S [4_or_5] × STANAG_L [4_or_5]	0.94	0.50	-0.12	0.908
MaskerType [N] × STANAG_S [2_or_3]	4.78	2.54	2.94	0.003
MaskerType [N] × STANAG_S [4_or_5]	1.15	0.63	0.26	0.798
MaskerType [N] × STANAG_L [2]	3.60	1.68	2.74	0.006
MaskerType [N] × STANAG_L [3]	2.14	0.95	1.71	0.087
MaskerType [N] × STANAG_L [4_or_5]	2.28	1.22	1.55	0.122
(MaskerType [N] × STANAG_S [2_or_3]) × STANAG_L [2]	0.18	0.11	-2.94	0.003
(MaskerType [N] × STANAG_S [4_or_5]) × STANAG_L [2]	0.57	0.34	-0.93	0.350
(MaskerType [N] × STANAG_S [2_or_3]) × STANAG_L [3]	0.22	0.12	-2.82	0.005
(MaskerType [N] × STANAG_S [4_or_5]) × STANAG_L [3]	0.79	0.44	-0.42	0.671
(MaskerType [N] × STANAG_S [2_or_3]) × STANAG_L [4_or_5]	0.43	0.29	-1.23	0.219
(MaskerType [N] × STANAG_S [4_or_5]) × STANAG_L [4_or_5]	0.99	0.66	-0.02	0.986

Random Effects

σ^2	3.29
$\tau_{00 \text{ Item}}$	0.53
$\tau_{00 \text{ Participant}}$	0.19
$\tau_{00 \text{ trialNumber}}$	0.04
$\tau_{11 \text{ wav.betterPTA_46_z}}$	0.24
$\rho_{01 \text{ Item}}$	-0.35
ICC	0.23
$N_{\text{Participant}}$	113
N_{Item}	160
$N_{\text{trialNumber}}$	10
Observations	1130
Marginal R^2 / Conditional R^2	0.274 / 0.442

Table 3. Response times. Model: $\text{LogRT} \sim 1 + \text{SNR_highlow} + (\text{STANAG_T} * \text{STANAG_L}) + (\text{Score} * \text{MaskerType}) + (1 | \text{Participant}) + (1 | \text{Item}) + (1 | \text{trialNumber})$. Reference levels: MaskerType = Babble, STANAG_S = Native, STANAG_L = Native, SNR = High, Score = 0. Number of obs.: 1130, Participants: 113, Items: 274. Trials = 10.

Predictors	Estimates	std. Error	Statistic	p
(Intercept)	9.85	0.07	147.14	<0.001
SNR_highlow [Low]	-0.07	0.02	-3.10	0.002
MaskerType [N]	-0.12	0.05	-2.15	0.032
Score [1]	0.06	0.07	0.97	0.330
Score [2]	0.14	0.06	2.29	0.022
Score [3]	0.08	0.06	1.41	0.159
Score [4]	-0.02	0.05	-0.29	0.770
Score [5]	-0.22	0.05	-4.25	<0.001
STANAG_L [3]	-0.06	0.05	-1.17	0.241
STANAG_L [4_or_5]	-0.11	0.07	-1.47	0.141
STANAG_L [Native]	-0.10	0.06	-1.75	0.080
STANAG_S [4_or_5]	0.15	0.05	3.14	0.002
STANAG_S [Native]	-0.01	0.05	-0.13	0.901
STANAG_L [3] × STANAG_S [4_or_5]	-0.18	0.06	-3.13	0.002
STANAG_L [4_or_5] × STANAG_S [4_or_5]	-0.14	0.08	-1.85	0.064
STANAG_L [Native] × STANAG_S [4_or_5]	-0.18	0.06	-2.76	0.006
STANAG_L [3] × STANAG_S [Native]	0.02	0.06	0.38	0.707
STANAG_L [4_or_5] × STANAG_S [Native]	0.07	0.08	0.84	0.401
STANAG_L [Native] × STANAG_S [Native]	-0.07	0.06	-1.06	0.288
MaskerType [N] × Score [1]	-0.05	0.08	-0.64	0.522
MaskerType [N] × Score [2]	-0.03	0.08	-0.40	0.690
MaskerType [N] × Score [3]	0.09	0.07	1.22	0.222
MaskerType [N] × Score [4]	0.03	0.07	0.41	0.679
MaskerType [N] × Score [5]	0.17	0.08	2.23	0.026

Random Effects

σ_2	0.08
$\tau_{00 \text{ Item}}$	0.00
$\tau_{00 \text{ Participant}}$	0.02
$\tau_{00 \text{ trialNumber}}$	0.00
ICC	0.26
$N_{\text{Participant}}$	113
N_{Item}	160
$N_{\text{trialNumber}}$	10
Observations	1108
Marginal R^2 / Conditional R^2	0.147 / 0.369

Author biographies

Rebecca E Bieber

Matthew J Makashay

Benjamin M Sheffield

Douglas S Brungart

Rebecca E Bieber is a Research Audiologist at the National Military Audiology and Speech Pathology Center at Walter Reed National Military Medical Center in Bethesda, Maryland. She received both her doctorate in Clinical Audiology (AuD) and her PhD in Hearing and Speech Sciences from the University of Maryland in College Park. Rebecca's research interests include the study of how humans are able to perceive speech under challenging conditions, and how individual factors such as aging, hearing loss, and language background impact this ability.

Matthew J. Makashay is a Hearing Program Studies Investigator in the Hearing Conservation and Readiness Branch of the Defense Centers for Public Health – Aberdeen and holds a PhD in Linguistics. His current studies in the National Military Audiology and Speech Pathology Center at the Walter Reed National Military Medical Center are in areas involving psycholinguistic interactions with hearing impairment, individual and dialectal differences in speech perception, the development of validated fitness-for-duty standards for speech intelligibility and comprehension, and hearing protection device fit testing.

Benjamin M. Sheffield is an Auditory Studies Investigator in the Hearing Conservation and Readiness Branch of the Defense Centers for Public Health – Aberdeen. He holds a MS in Mechanical and Aerospace Engineering and a BS in Electrical Engineering. After a four-year stint in the defense industry, he was introduced to auditory research as a volunteer in a hearing and speech lab during graduate school at UC Irvine. He has been a member of the National Military Audiology and Speech Pathology Center at Walter Reed since 2010, primarily focusing on the relationship between hearing acuity and operational performance, as well as the application of boothless audiometry in hearing conservation programs.

Douglas S. Brungart, PhD, is the Chief Scientist of the National Military Audiology and Speech Pathology Center at Walter Reed. From 1993 to 2009, he was a research engineer at the Air Force Research Laboratory with research focusing on advanced auditory displays for spatial and speech information. Since 2009, he has been at Walter Reed focusing on the application of advanced technology to improve the prevention, diagnosis, and treatment of hearing loss and other hearing and speech disorders. He holds a MS and PHD in Electrical Engineering from the Massachusetts Institute of Technology and a BS in Computer Engineering from Wright State University.



NATO
C2COE

SET A GOAL, ALLOCATE RESOURCES, DON'T GET IN THE WAY

7 Principles for effective acquisition of C2
support systems

The acquisition of technical command and control support systems (C2 systems) is characterised by high complexity and can be understood as a 'wicked problem'. Although the systems themselves are technical, their acquisition involves interrelated organisational, human, and technical factors, as well as significant uncertainty. Consequently, understanding of the problem must evolve in parallel with the acquisition process.

The objective of this study was to identify principles capable of enhancing the effectiveness and efficiency of C2 system acquisition, which often suffers from budget overruns, delays, and suboptimal outcomes. The study employed an iterative research design. Initially, an exploratory literature survey identified success factors for effective C2 system acquisition, which were synthesised into seven principles. Semi-structured interviews with 14 participants were then conducted to validate, refine, and expand these principles. Finally, a case study involving focus groups, and additional interviews provided insights into the practical application of the principles and their organisational effects.

The seven principles for effective and efficient C2 system acquisition derived from this study are: (1) recognition of acquisition as a complex and wicked problem; (2) development of a supportive organisational culture; (3) adoption of goal orientation, a holistic perspective, and tailored processes; (4) implementation of learning-oriented follow-ups; (5) maintenance of competence over time; (6) promotion of collaboration with stakeholders; and (7) management of uncertainty through adaptability.

The results of the case study show that these principles were successfully applied at team level, leading to improvements in efficiency and effectiveness. The study also challenges the need for comprehensive early requirements, advocating for evolving requirements that respond to dynamic project conditions.

Keywords:

Acquisition, C2 systems, Complexity, Principles

Introduction

The deteriorating security situation in Europe following Russia's annexation of Crimea and the subsequent invasion of Ukraine has prompted many countries to strengthen their defence forces (European Defence Agency, 2023). A key factor in enhancing military capability is the ability to conduct effective command and control (C2) (U.S. Marine Corps, 2017). Central to this development is the increasing use of modern information and communication technologies, which may reduce some of the friction traditionally associated with military operations. By enabling the direct exchange of information and facilitating decision-making at the most appropriate level and at the right time, digital technologies can improve information flows and enhance coordination across organisational boundaries (Alberts et al., 2002; Claverie & Desclaux, 2022). At the same time, the rapid development of information and communication systems means

that conditions for C2 are constantly evolving, and traditional approaches often lack the agility required in contemporary conflicts (Alberts & Hayes, 2006; Granåsen et al., 2023).

Organisations that can leverage recent technological advances thus hold a significant advantage over those that cannot. For example, these organisations can transform C2 by replacing fragmented, slow, and siloed communication with integrated, real-time information sharing that provides a clearer and shared operational picture, thereby enabling faster and better-informed decisions (Davidson, 2025; Hoehn, 2020; de Leon, 2021). Newer concepts for C2, such as the Swedish C2 Concept 2045 and the United Kingdom's Command and Control in the Future, assume access to modern information and communication systems for their realisation and regard C2 as a capability that must be continuously developed (Ellis et al., 2024; Granåsen et al., 2023). Rapid technological advancements therefore

necessitate that the acquisition of information systems to support military C2 (hereinafter referred to as C2 systems) be conducted more swiftly in order to maintain their relevance.

Acquiring C2 systems remains a challenging task as projects responsible for their acquisition are commonly associated with delays and budget overruns. In addition, acquired C2 systems are often poorly adapted to user needs and inadequately integrated with existing systems, and discussions on improving acquisition processes have been ongoing for years (Berg & Ritschel, 2023; Wong et al., 2022). These discussions have often focused narrowly on process modifications, organisational restructuring, or the introduction of new controls to ensure that projects are delivered on time, within budget, and according to specifications. This limited focus has overshadowed broader considerations of fundamentally transforming acquisition practices to achieve more substantial and sustainable improvements. Brady (1990) states that, reorganisation and restructuring can improve efficiency to some extent, but more profound changes are required to address the root causes of the problem.

The objective of this study is to identify principles capable of genuinely enhancing the effectiveness and efficiency of C2 system acquisitions. In pursuing this objective, the study adopts a complexity- and systems-theory perspective.

Complexity and Wicked Problems

The nature of problems is often overlooked when deciding on approaches to solving them. This lack of awareness significantly affects the chosen approach and, consequently, the outcome – especially when failing to understand the difference between complicated and complex problems (Nelson & Stolterman, 2012; Poli, 2013). Although there is no uniform definition of complexity, the literature offers a wide array of similar suggestions. Complicated problems can be analysed in terms of separate causes and effects, which can subsequently be addressed in a stepwise manner. Complex problems, on the other hand, result from multiple interrelated causes interacting in such a way that they cannot be separated from each other

(Hasan, 2014; Poli, 2013). As a result, the idea of a correct solution is irrelevant in complex contexts, and alternatives should be evaluated in terms of better or worse based on their effects.

The concepts of complicated and complex problems are similar to the notion of tame and wicked problems (Rittel & Webber, 1973). Tame problems are characterised by the fact that it is relatively easy to identify relationships between cause and effect. Thus, tame problems resemble complicated problems. Wicked problems, on the other hand, are dynamic, complex, and poorly structured, making it extremely difficult to identify or simulate causes and effects (Batie, 2008). However, there is no clear line between tame and wicked problems; rather, problems vary in their degree of wickedness (Farrell & Hooker, 2013).

Socio-technical problems in areas where technology changes rapidly, and military problems within the domain of systems of systems, are often characterised by a high degree of wickedness (Albinsson, 2004; Cook & Pratt, 2014). Additionally, Ardagna and Francalanci (2005) argue that the development of IT infrastructures involves a complex design process. Nelson and Stolterman (2012) state that failing to identify a wicked problem can overwhelm those involved, leading to various types of paralysis, such as analysis, value, holistic, or wicked problem paralysis.

According to the Cynefin framework, decision-makers or problem-solvers should adopt different approaches depending on whether a problem is complicated or complex in order to achieve effective outcomes (Snowden and Boone, 2007). In complicated contexts, expert analysis and structured methods can be applied, as these problems have identifiable relationships between cause and effect. In contrast, complex problems require more creative and adaptive approaches, such as experimentation and iterative learning, which allow patterns to emerge and guide subsequent actions.

Research Design

Improving the effectiveness of C2 system acquisition is inherently a complex problem, involving multiple interrelated socio-technical factors, e.g. doctrine, organisation, training, materiel, leadership, personnel, facilities and interoperability. To address this complexity, a stepwise, iterative approach was adopted for this study, beginning with an exploratory literature survey to establish a foundational understanding of the field. Subsequent research steps were adapted based on the findings obtained at each stage, allowing the methodology to evolve in response to emerging insights and opportunities. The final research design comprised four sequential phases. The study began with an exploratory literature review, ending with the formulation of preliminary principles, and a supplementary literature survey which continued in parallel with the subsequent phases: an interview study and a case study (Figure 1). Each phase built upon the previous one, ensuring that the methods employed were informed by the evidence gathered and tailored to the specific characteristics and challenges identified in the context of C2 system acquisition.

The empirical data from the interview study and the case study were collected in Swedish, where the word 'effektiv' is used in a broad, context-dependent manner that does not consistently distinguish between the analytically separate concepts of 'efficiency' and 'effectiveness' as defined in the English-language. Since the empirical material does not allow for a reliable separation of these dimensions, we treat 'effektiv' as an emic evaluative category encompassing both. Accordingly, the combined expression 'efficiency and effectiveness'

are used to reflect this conceptual breadth in the source material, rather than to imply a clear empirical distinction between the two.

Literature survey

The aim of the literature survey was to identify success factors for effective and efficient acquisition. Since the acquisition of C2 systems in military contexts involves processes both preceding and following procurement—such as needs identification and analysis, development, and integration of the new product into the socio-technical C2 system—the literature study was not confined to success factors exclusively related to procurement. Moreover, the study was not limited to military organisations or scholarly papers. Rather, it encompassed success factors related to product and software development, project management, acquisition, and integration in both civilian and military contexts, drawing on both scholarly and more practice-oriented literature to provide a broad foundation for improving C2 system acquisition.

Using a snowballing method (Badampudi, 2015; Jacobsen, 2002), the exploratory literature survey produced in total approximately 130 reports, articles, papers, books, and other written sources, all of which were reviewed. During the review phase, factors that the original authors had identified as important to effectiveness or efficiency in acquisition, project work, or development—success factors—were noted.

A process of categorising and compiling the success factors began once 129 success factors had been identified through the literature survey. In this process, success factors with similar meanings




Figure 1: The research approach.

were identified, grouped together and synthesised to form seven principles that were deemed to increase effectiveness and efficiency when acquiring C2 systems. An additional 199 success factors were identified during the supplementary literature review. Since all new success factors could be categorised into the seven previously formulated principles, data saturation was deemed to have been achieved, and the literature study was concluded.

Interview study

The aim of the interview study was to explore whether experiences from existing procurement projects within the Swedish Armed Forces supported or contradicted the notion that the seven principles from the literature study could enhance acquisition effectiveness and efficiency.

The interview guide for the interview study was developed with reference to the preliminary principles

emerging from the literature review. Rather than being translated directly into interview questions, the principles informed the selection of themes and areas of inquiry addressed during the interviews. Consequently, the interview questions focused on participants' experiences and practices relating to issues such as requirements management, stakeholder involvement, risk and uncertainty management, communication, decision-making, competence, and incremental development, rather than on the principles themselves. This approach enabled an open exploration of acquisition and integration practices while ensuring coverage of topics identified as important in the literature.

In order to identify suitable interview candidates, several relevant acquisition projects were selected, and individuals working on these projects were contacted. In total, 14 interviews were conducted with people employed at the Swedish Armed Forces (SwAF) or the Swedish Defence Materiel Administration (SDMA) during 2020 (Table 1).

Table 1: Conducted interviews

Interviewee	Interview length	Organisation	Role
1	109 min	SDMA	Project manager
2	101 min	SwAF	User representative
3	102 min	SwAF	User representative
4	165 min	SwAF	User representative
5	81 min	SwAF	User representative
6	100 min	SwAF	User representative
7	111 min	SDMA	Project manager
8*	120 min	SwAF	Overarching***
8*	≈ 90 min**	SwAF	User representative
9	101 min	SwAF	Overarching***
10	70 min	SDMA	Overarching***
11	74 min	SwAF	Overarching***
12	67 min	SwAF	Overarching***
13	55 min	SwAF	Overarching***

*The same person was interviewed twice, as the individual in question had relevant experience from different roles in the acquisition process.

** Due to technical problems, this interview was not recorded in its entirety; therefore, it was not possible to determine the exact length of the interview.

***E.g., senior officers, chief engineers, or senior managers

The interviews were analysed in two steps using the software NVivo. In the first step (Figure 2), it was evaluated whether a statement from the interviewees supported or contradicted a hypothesis that adhering to a principle would lead to more effective and efficient acquisition and integration. If the statement aligned with or contradicted the hypothesis, it was transcribed and incorporated into the analysis.




Figure 2: The first step in the post-interview analysis process.

In the second step, the findings from the first step were compared with each of the principles from the literature survey to ascertain whether the interview statements supported or contradicted the hypothesis that the principle in question contributed to a more efficient acquisition and integration of C2 systems (Figure 3).




Figure 3: The second step in the post-interview analysis process.

Case study

The case study was carried out with two teams (Modern Approach Group 1 and 2) within the Swedish Armed Forces Communication and Information Systems Command (SwAF CISCOT). In order to improve the effectiveness of receiving and integrating two IT systems into the organisation's operations, the two teams had changed their way of working and implemented what SwAF CISCOT referred to as Modern Working Methods. These methods were inspired by mission command, service-oriented leadership, trust-based management, and methodologies from the civil production industry, such as agile methods, DevOps, Scrum, and the Scaled Agile Framework (SAFe). The development of Modern Working Methods was also informed by early findings (the seven principles) from the initial exploratory literature survey reported in this article.

The aim of the case study was to:

1. evaluate the extent to which Modern Working Methods adhere to the seven principles identified in the literature review
2. assess whether the seven principles indeed lead to increased effectiveness
3. gather insights and experiences to support the further development and refinement of the seven principles.

The case study encompassed the two Modern Approach Groups and one Traditional Working Group that followed established traditional processes.

Two focus group interviews, conducted in 2022 and 2024, were held with each of the Modern Approach Groups, and one focus group interview was conducted with the Traditional Working Group in 2024.

The interview data were analysed independently by three of the authors. Two analysts applied the procedure described in Figures 2 and 3, using NVivo to support the coding and analysis. This analysis was based on the notes taken during the group interviews. Subsequently, a more detailed examination was conducted focusing on attitudes towards Modern Ways of Working, issues related to trust, and divergent perceptions of processes.

The third analyst reviewed the audio recordings of the interviews and transcribed selected passages that either corroborated or challenged the preliminary principles identified through the literature survey.

To verify some of the claims made by the Modern Approach Groups during the focus groups, three additional interviews were conducted in 2025 with:

- the officer responsible for the materiel domain to which the IT systems in focus for the two Modern Approach Groups belonged
- one representative of the client organisation for each of the two IT systems, with each system having been the subject of a separate Modern Approach Group.

The additional interviews were analysed primarily to assess whether the Modern Approach Groups' claims of greater effectiveness compared with other groups could be substantiated, and to examine how collaboration between the working groups and other stakeholders was perceived and experienced.

The Seven Principles

This chapter presents the seven principles for effective and efficient acquisition of C2 systems. Each principle is presented with a brief description and the rationale for the principle alongside the relevant findings from the interviews and case study. The seven principles are:

1. Recognition of Acquisition as a Complex and Wicked Problem
2. Development of a Supportive Organisational Culture
3. Adoption of Goal Orientation, a Holistic Perspective, and Tailored Processes
4. Implementation of Learning-oriented Follow-ups
5. Building and Sustaining Competence
6. Promotion of Collaboration with Stakeholders
7. Management of Uncertainty through Adaptability

Recognition of Acquisition as a Complex and Wicked Problem

The principle of *Recognition of Acquisition as a Complex and Wicked Problem*, is an important starting point, as it determines how acquisition work should be organised to achieve success.

Underlying rationale

A clear indication that C2 system acquisition is a complex issue is the need to consider factors within the socio-technical C2 system that extend beyond the system itself. Furthermore, the acquisition of C2 systems exhibits all the characteristics of a wicked problem, as outlined by Rittel and Webber (1973, p. 161). For example: (1) there is no definitive formulation of a wicked problem—the full scope of users' needs for C2 systems cannot be fully determined in advance, as these needs depend on operational contexts, emerging threats, and technological developments; (2) wicked problems have no stopping rule—C2 systems can always be further improved, and developers can never claim that the system is entirely complete; and (3) solutions to wicked problems are not true-or-false, but rather good-or-bad—there is no single correct design for a C2 system, and stakeholders may hold differing views on the system's design.

Command and Control (C2) is characterised by a high degree of complexity (Stanton & Baber, 2026). In accordance with Ashby's (1956) Law of Requisite Variety, this complexity also shapes the process of acquiring C2 systems. A complex C2 environment requires systems that can provide sufficient variety to respond effectively to diverse and dynamic conditions. Consequently, the development process itself must be capable of generating the necessary variety to meet these demands. This perspective aligns with Sommerville et al. (2012), who argue that software cannot be understood independently of the socio-technical environment in which it operates.

The complexity of military capability development is further heightened by the need to address not only technical system characteristics, but also broader considerations such as doctrine, organisation, training, materiel, leadership, personnel, facilities,

and interoperability. As a result, C2 system acquisition involves a wide range of stakeholders with differing perspectives, priorities, and requirements. This diversity further contributes to the inherently complex and wicked nature of the problem.

Albinsson (2004) similarly concludes that the development of C2 systems shares many characteristics of a wicked problem. Moreover, socio-technical problems in rapidly evolving technological domains frequently display traits of wicked problems, a feature also evident in military problems within the system-of-systems context (Cook & Pratt, 2014). Statements from the interview study further support the conclusion that the acquisition of C2 systems is a complex problem:

Then there are always those who think that competitive tendering will result in the lowest price. I don't believe that for a second in this type of complex environment. /.../ It is too hard to write everything down and then it's too dynamic. Even if the specification is perfect at the time, it'll be completely different when the product is ready. It doesn't matter if the project, project members and the project leader are at the top of their game if the conditions aren't right. (Interviewee 7, 2020).

Development of a Supportive Organisational Culture

The principle of *Development of a Supportive Organisational Culture* involves fostering ways of working that facilitate the resolution of complex problems and rewarding behaviours that increase the likelihood of successful and efficient acquisition.

This includes providing those tasked with solving a problem with the best possible conditions for doing so, trusting them to complete the work, offering whatever assistance is required, and promoting innovation. At the same time, the organisational culture should be shaped so that appropriate behaviour is rewarded, thereby conferring prestige on running small, simple, and fast projects that deliver what users actually need, rather than on large projects, which face a greater risk of delay and cost overruns.

Underlying rationale

Organisational culture can be understood as the shared values, norms, and practices that shape how individuals and groups behave within an organisation (Schein, 2010). In the context of C2 system acquisition, this culture is manifested through several interrelated factors, including trust, approaches to decision-making, risk aversion, the design and use of processes, and the encouragement of small, simple and fast projects. The following section therefore examines these aspects in more detail.

Trust-based management has been shown to provide conditions suitable for the solving of complex problems (Bringselius, 2019a) and it is also important for the efficiency in large organisations (La Porta et al., 1996). Servant leadership—that is, that higher hierarchical levels (managers) are prepared to support and assist lower levels when needed—is an important part of trust-based management (Bringselius, 2019b). However, trust cannot be limited to within one's own organisation, it needs to be built across the different organisations that are involved in the acquisition:

We have to work TOGETHER, that's the thing. This whole thing about sending bits of paperwork between two agencies and then saying that one of them doesn't do its job properly in some way, that's no good. (Interviewee 12, 2020)

According to Johnson (2022), the single most important factor for successful IT projects is the ability to make rapid decisions. The capacity for swift decision-making is so crucial that it is often more cost-effective to make a poor decision early, which may later require correction, than to delay a decision in an attempt to get it right the first time. A key factor in achieving rapid decision-making is a broad allocation of decision rights, allowing decisions to be made as far down the organisation as possible (Baccarini & Collins, 2003; Breiding & Brilla, 2016; Pinto & Slevin, 1989).

Several interviewees mentioned experiencing an excessive focus on established processes, decision gates, and achieving set project milestones, which they reported leads to inefficiency:

The process has created bottlenecks; I mean I shouldn't just blame the process, there are people involved as well. But all these different roles, decision gates and part processes make it too long to get through, so the people who are supposed to do it [the work] are sitting idle a large part of the time. (Interviewee 8, 2020)

Consequently, most decisions should be taken at the team level, consistent with a mission command philosophy, unless there is a need for higher-level direction or coordination. In addition to a broad allocation of decision rights, the ability to make rapid decisions also depends on employees' willingness to take initiative, which is influenced by the level of risk aversion within the organisation. Achieving swift decision-making therefore requires that risk aversion is not rewarded, which in turn necessitates a general acceptance within the organisation that mistakes will sometimes occur. During one of the focus groups in the case study, a participant from one of the Modern Approach Groups commented on the importance of accepting mistakes, so that work can proceed efficiently without being stalled by excessive risk minimisation:

What I hope will come more naturally is that risk avoidance isn't the main focus anymore... That it's okay to make mistakes, it's expected, that's how it works. But that the focus shifts to actually getting things done, rather than planning and trying to minimise risks right up to the very last moment. (Interviewee from Modern Approach Group 1, 2024)

These behaviours may also reflect implicit or tacit organisational objectives—such as risk avoidance, adherence to established processes, or the preservation of existing structures—which may, at times, be in tension with formally stated goals.

Risk aversion may also undermine an organisation's innovative capacity, which is crucial for effectively addressing complex problems (Blackham, 2007; Cranfield, 2007), as it is easier to control risks when established solutions and working methods are chosen than when new approaches are tested.

This emphasis on reducing risk aversion should be distinguished from the concept of risk mitigation. Reducing risk aversion does not imply an absence of risk management. Rather, it concerns creating conditions in which individuals are willing to act despite uncertainty. In this context, structuring work into smaller and more manageable projects can be seen as a form of risk mitigation, as it reduces complexity and uncertainty in each individual initiative while still enabling experimentation and learning.

Small, simple, and fast projects should be encouraged since these pose a lower risk of failure (Berg & Ritschel, 2023; Sauer et al., 2007; Ward, 2014). However, measures should be taken to prevent the deliberate or unintentional underestimation of project time or cost in order to secure approval to start, or as one participant stated in the interview study:

There are probably some [people] who know right from the start that it [the project] won't be finished on time. They know that there is a certain budget and deadline that they need to agree to, or the decision [to start the project] will not be made. And then they think, based on how things have worked previously, that once they're up and running they'll make it work. The money won't dry up. Maybe we haven't been clear enough. (Interviewee 8, 2020)

Examples of such measures could include terminating projects whose costs exceed a certain proportion of the original budget, thereby reducing incentives for individuals to underestimate project costs (Ward, 2014).

Adoption of Goal Orientation, a Holistic Perspective, and Tailored Processes

The principle of *Adoption of Goal Orientation, a Holistic Perspective, and Tailored Processes* holds that acquisition activities should be guided by clearly defined objectives that reflect users' operational needs while simultaneously taking into account the system's role within the wider socio-technical context. In this sense, goal orientation provides the direction for acquisition by ensuring that activities are aligned with intended operational outcomes, whereas the holistic perspective ensures that these goals remain appropriate within the broader environment.

However, as both operational needs and contextual conditions evolve over time, these objectives cannot be treated as static. Instead, they may need to be continuously adjusted in response to changes in requirements, technology, and operational context, thereby emphasising the need for flexibility in both planning and execution.

To support this adaptability, acquisition processes should be designed to enable and not constrain the work. This implies that processes must be sufficiently flexible to accommodate changing conditions and differing project contexts, rather than enforcing a uniform approach across all types of projects. Consequently, tailored processes are not a separate consideration but a necessary mechanism for enabling both goal-oriented action and context-sensitive adaptation.

Underlying rationale

Several sources identified in the literature survey emphasise the formulation of explicit, user-centred objectives, grounded in operational needs, as essential for successful IT projects (Berg & Ritschel, 2023; Lawless, 2019; Ward, 2014). The formulation of such objectives should begin with an understanding of the users' needs to address their operational tasks and with a holistic perspective that recognises that the C2 system will form part of a wider socio-technical system (Arvidsson et al., 2012; Baskerville, 1999; Currie & Guah, 2007; Lytinen et al., 2004; Modigliani, 2015; Salwin, 1998).

Experience has shown that knowledge about users' needs, and consequently the objectives for an acquisition, often needs to be revised during an ongoing acquisition or after a new system has been introduced (Albinsson, 2004; Kasser, 2002; Laufer & Hoffman, 1998). Such revisions are, however, constrained by project limitations, such as time and cost. One reason for the need to revise objectives is that users often do not know what they actually need at the outset of an acquisition, owing to the complexity of the socio-technical C2 system:

The root of the [Swedish] Armed Forces' problem is an overconfidence in planning and an overconfidence in 'point based', meaning that they know what they want from the beginning, although, in reality, they actually don't know what they [need]. (Interviewee 8, 2020)

Clearly formulated objectives for acquisition projects also enable a movement away from process-focused milestones. This allows for a more efficient use of time, as time is not spent on follow-up activities conducted solely for the purpose of formally reporting what has been done. By maintaining a stronger focus on the project objectives rather than the processes, project teams can prioritise delivering results that meet user needs, even if this requires departures from standard routines or the adaptation of existing processes to the specific context of the acquisition. As one participant in the interview study explained:

Our goal has been a good end result. I think that differs among [different] people, whether you're just interested in the delivery or if you actually want to produce something useful for the end users. I think we've prioritised that [a good end result for the end users] in this project; that's been the goal for everyone involved. It's not the processes and the deliverables [that matter]; they're secondary to a good end result for the end customer. (Interviewee 1, 2020)

With regard to processes and routines, these should be primarily tailored to the acquisition team. This can be achieved either by having the teams develop them themselves or by adapting them from one of the organisation's existing processes (Karlsson & Ågerfalk, 2009; Ward, 2014). What is important is

that the processes do not impose constraints but instead support the work of the acquisition teams, as the following dialogue from one of the focus groups illustrates:

Interviewee A: [The processes] we write down and approve are things we've actually tried out and stand behind. /.../ We work from experience; it's not like we plan how something should be done, write it down, approve it, and only then start doing it.

Interviewee B: We don't really have any [governing] documents that are normative, do we? They're all just descriptive: –This is how we do it. We don't have anything that says: –This is how we must do it. Right?

Interviewee A: No, exactly. The only thing that says: –This is how we should do it, is... well, that's our objective. (Dialogue between interviewees from Modern Approach Group 1, 2024).

Focusing on achieving the current objectives for the acquisition also means limiting documentation to what is absolutely necessary (Beck et al., 2001; Breiding & Brilla, 2016; Lawless, 2019; Pinto & Slevin, 1989; Sauer & Reich, 2009; Ward, 2014). This does not imply that documentation is unnecessary, but rather that its level and timing should be aligned with the development process. In the empirical material, challenges were described where extensive upfront documentation was produced and formalised before implementation, only to require subsequent revision when it proved misaligned with actual development outcomes:

Since it's an important function we've developed, whenever there's been a discussion about the documentation, I've always been able to weigh it against getting the actual work done: –Do you want this product or not? –Yes, we do.

–Right, then we don't bother with all that documentation now; you'll get it when you get the product. We simply don't have the energy to produce all the documentation first, only to have a meeting where it all gets torn apart and then we have to go back and redo it, while no code gets written in the meantime. (Interviewee from Modern Approach Group 2, 2022)

Implementation of Learning-Oriented Follow-Ups

The principle of *Implementation of Learning-Oriented Follow-Ups* means that follow-ups should primarily support learning and improvements rather than focus solely on control. Follow-ups should be conducted in a collegial manner and support the identification of:

- which actions and activities need to be taken to ensure that the objectives for the acquisition are achieved
- which measures should be implemented to make subsequent work more efficient
- what lessons have been learned that can contribute to more effective work in the future

This means that follow-ups should primarily focus on the future, i.e., what work needs to be carried out to achieve acquisition success, and not on what the work has achieved so far, while still fulfilling necessary functions of control and accountability.

Underlying rationale

The manner in which performance is monitored can have far-reaching consequences within an organisation, since monitoring and follow-ups influence the work conducted in various ways and can affect the outcomes (Bringselius, 2019b; Muller, 2018). It is, therefore, important that follow-ups focus on factors essential for the success of the acquisition. The interview study supports this view, with interviewees describing how a predominant focus on so-called milestones has had a negative impact on product quality and interoperability. This problem is a result of project managers chasing green milestones (checking boxes), often linked to time or cost, rather than focusing on activities that ensure the quality of the final product:

The green milestones were what mattered, whether you actually delivered something of value or something that worked wasn't important. That was a bit of a shame, I think. (Interviewee 12, 2020)

On the other hand, follow-ups centred on addressing the problems that need to be solved to meet users' and stakeholders' needs contribute positively to effectiveness and efficiency:

Production is the focus [during the weekly meeting]. And each team /.../ bring up what we need help with, where we've got stuck, and it's always nice to share a few wins as well, just to get a bit of positivity. It's not a long status update or anything, you just raise the things you need help with and everything else is business as usual. (Interviewee from Modern Approach Group 2, 2024)

Moreover, follow-ups should be based on actual deliverables rather than mere reports and be conducted in cooperation with users and other stakeholders. It is essential to avoid follow-ups that serve only to instil a sense of control among senior managers and hold little value for those involved in the acquisition, bearing in mind that excessive control increases the likelihood of inefficiency and malpractice (Laufer & Hoffman, 1998). Claims supporting these results from the literature survey were also found in the case study:

Interviewee A: But before [Modern Working Methods], it was just so much... There were all these G2 documents that had to be produced...

Interviewee B: No, that was just time wasted.

Interviewee A: ...and the list kept changing, which documents needed to be done and which templates applied? I mean, we spent—don't know the exact percentage—but way too much of our working time on documentation that just had to be redone constantly. And trying to figure out which documentation it was at any given time... /.../

Interviewee B: I'd also like to add that it's not like—we've never claimed this—we've stopped doing documentation. But what I want to say is that we still have full traceability of what we've done and the results. So if someone asks—what have you done—we can go back and show: this was what we were supposed to do, this is what we developed, and this is what we validated and

delivered. So we do have that documentation.

Interviewee A: And I'd say—even more than many other projects, who have to go through their CABs [Change Advisory Boards] left, right, and centre, sitting there half-lying to get things approved. So those who they push stuff for—now I'm saying a lot here on record—but they push it because they know, we just have to get through this CAB. It doesn't always even match the documentation. (Dialogue between interviewees from Modern Approach Group 2, 2024)

Since the principle of Learning-Oriented Follow-Ups is closely related to the formulation of clear objectives, it is important to emphasise that information obtained from follow-ups with end users and stakeholders may require updates to the acquisition objectives. Consequently, the objectives set at the beginning of a project are not static but are updated continuously as new insights emerge.

Building and Sustaining Competence

The principle of *Building and Sustaining Competence* is necessary to enable work guided by clear, result-focused goals rather than by strict processes. In this context, competence is understood as a multi-dimensional construct that includes both individual and collective capacities. Access to sufficient competence is ensured through:

- careful selection of acquisition project members, taking into account factors other than purely knowledge, i.e., affective factors, personality traits, and social skills
- long-term work developing new competencies in the development, acquisition, and integration of C2 systems
- low staff turnover

It is important to remember that competence does not consist solely of specialist knowledge; aspects such as commitment, values, collaboration, communication skills, and leadership all need to be considered when selecting team members.

Underlying rationale

The selection of project team members is crucial for the success of acquisition projects, and it is essential to ensure that team members possess the competence required to carry out the acquisition project effectively, bearing in mind that competence encompasses more than knowledge within a specific field (Baccarini & Collins, 2003; Beck et al., 2001; Crowstone et al., 2004; Delano, 1999; Laufer & Hoffman, 1998; Pinto & Slevin, 1989). Competence is therefore understood as a multi-dimensional construct that includes cognitive and technical knowledge, as well as affective and personality-related factors such as attitudes, values, motivation, and self-confidence, together with social, communication, and cooperative skills at the individual level. It also encompasses collective competence at the team level (Ellström & Kock, 2008). Accordingly, these aspects should be considered when selecting team members.

Taken together, this implies that ensuring sufficient competence in acquisition projects requires attention not only to individual capabilities, but also to how competence is developed over time and sustained within teams and organisational structures.

One team leader stated:

The attitudes and initiatives of the individuals involved in the work are pretty central for agile work to function well – that people tackle challenges themselves, or if they can't solve something on their own, they go to a colleague or a manager and get support, guidance, or help with making a decision. /.../ Having people in the team who are not only solution-focused but also open to dealing with things like that and taking initiative really makes a huge difference for a group like this to function properly. (Interviewee from Modern Approach Group 1, 2024)

In addition to the individual team members' competence, it is important that the team as a whole exhibit diversity of perspectives, competencies and experiences, as this facilitates the resolution of complex problems (Snowden & Boone, 2007).

Proactivity is important to ensure access to the competence required for the acquisition of C2 systems. This means that long-term investments are required in modelling, simulation, experimentation, and the development of the capability for rapid testing, along with a continuous focus on high-quality technology and good design. This entails sustained efforts to evaluate both existing and emerging technologies (Beck et al., 2001; Laufer & Hoffman, 1998; Modigliani, 2015; Nylén & Holmström, 2015; Ward, 2014).

Efficiency losses occur in acquisition projects when team members leave, as this leads to a direct loss of competence, and also due to the time it takes to introduce a new team member to the work. Interviewees emphasised the challenge of retaining the available competence within acquisition projects by minimising staff turnover. High staff turnover carries the risk of other negative effects, in addition to the loss of valuable competence, such as disrupting collaborative dynamics and team spirit:

It [staff turnover] is one heck of a problem. /.../ We know it's bad [replacing staff during a project]. So much is lost when we replace staff. How do we avoid replacing staff during a project? This has to mean that the projects are as short as possible. /.../ The shorter the project, the less risk of staff turnover, the less risk for the entire project, and the more likely that we deliver. /.../ If there are projects that run over several years, with a lot of staff replacements, you can bet your bottom dollar that it'll turn into a complete money pit and end up going left, right and centre. In the worst-case scenario, it'll completely miss the original goal because it'll have died in this staff replacement game of Chinese whispers throughout the course of the project. (Interviewee 8, 2020)

By making acquisition projects short, the problem of staff turnover in specific projects can be minimised. However, this does not address the general problem of competence loss within an organisation. The implementation of Modern Working Methods, though, seems to have a positive effect on work satisfaction:

Even with all the pushback in the organisation, it's still more enjoyable [working according to Modern Working Methods]—and I say that as someone who used to be just another cog in that very traditional wheel. It's so much more fun to actually get on with the work. (Interviewee from Modern Approach Group 1, 2024).

Promotion of Collaboration with Stakeholders

The principle of *Promotion of Collaboration with Stakeholders* concerns developing cooperation with various stakeholders—such as users, system managers, suppliers of the C2 system, and information security specialists—at an early stage of C2 system acquisitions in complex environments. Developing the necessary level of collaboration requires effective communication, particularly during the initial phases of the project, when it is essential to establish and nurture trust among the different stakeholders. While competence includes the ability to communicate and cooperate, effective collaboration with stakeholders depends on more than individual or team capabilities. It emerges through interactions between actors and is shaped by organisational arrangements, relationships, and forms of coordination.

Underlying rationale

To ensure that each member contributes effectively to the acquisition work, project teams should generally be kept relatively small, both in absolute terms and in relation to the mission (Ward, 2014). However, due to the complexity of C2 system acquisition, such teams cannot encompass all necessary domain knowledge and the use context for which the system is being acquired. This makes collaboration with external stakeholders essential in order to access the full range of expertise required (Baccarini & Collins, 2003; Currie & Guha, 2007; Gareis et al., 2013; Lawless, 2019; Šmite et al., 2010). This view is also supported by the interviews conducted:

We also communicate with those involved. That's the model—talk to each other, both the users, the monitoring authorities, the information security officers and the Armed Forces Intelligence and Security Service, etc. You talk to each other. You understand each other. That's what it's about. /.../ We interact with 30 different projects, it's all about collaboration. (Interviewee 7, 2020)

The relationship between supplier and client should, for an effective acquisition, be characterised by collaboration and flexibility rather than risk allocation (Baccarini & Collins, 2003; Beck et al.,

2001; Laufer & Hoffman, 1998; Reagan & Rico, 2010). The future users of the C2 system are an especially important group to collaborate with, since involving the users at an early stage reduces the need for future changes. In general, there should be daily collaboration between developers and domain experts (Beck et al., 2001; Laufer & Hoffman, 1988; Marple et al., 2001).

Choosing the right means and forms of interaction between stakeholders is imperative for achieving close collaboration. At the beginning of an acquisition project, when trust needs to be developed, face-to-face communication should be the primary choice, as it is the most effective form of communication. Face-to-face communication should also be the primary choice when misunderstandings or similar issues need to be resolved throughout the acquisition project (Bannister & Remanyi, 2009; Beck et al., 2001; Cockburn, 2000; Crowstone et al., 2007; Laufer & Hoffman, 1998).

Management of Uncertainty through Adaptability

The principle of *Management of Uncertainty through Adaptability* includes the ability to respond flexibly to unforeseen developments, adjust plans when necessary, and continuously learn from emerging situations in order to achieve the project objectives. Consequently, organisations involved in the acquisition of C2 systems should, according to this principle:

- ensure that plans are not more detailed than the prevailing uncertainties allow
- adapt the organisation, processes and the working practices of the project team so that swift decisions can be made, this includes a broad allocation of decision rights
- apply an agile and stepwise—incremental or evolutionary—acquisition process
- support the adaptability of the C2 system under acquisition to unforeseen future needs
 - during the acquisition process
 - after the acquisition has been completed

In addition, acquisition teams should, in order to limit uncertainty (Ward, 2014):

- begin by working on what is most important for the success of the project—not on what is easiest
- primarily use mature and proven technologies, and limit innovation to a maximum of one novel technology at a time—acquisitions that rely on the implementation of more than one novel technology carry a greater risk of failure.

Underlying rationale

The consequence of C2 system acquisition being a complex and wicked problem is that it is characterised by high uncertainty, interdependent factors, and the absence of clear-cut solutions, with understanding of the problem evolving alongside potential solutions (Rittel & Webber, 1973; Snowden

& Boone, 2007). This implies that detailed planning carried out at an early stage will likely need to be revised as new insights are gained or new factors are identified or clarified. The more detailed a plan is, the more work will have been carried out in vain if it subsequently needs revision. Since efficiency is not compatible with redundant work, detailed planning should therefore be avoided early in projects involving uncertainty (Laufer & Hoffman, 1998). The problem with early planning was also highlighted during the case study interviews:

Asking us for recurring activities, planned activities, budgeting and so on, for 2026—I'd say it's completely pointless. There's no value in it at all—if anything, it's the opposite. You're committing to do something you don't have the slightest ability to plan at that level of detail. (Interviewee from Modern Approach Group 1, 2024)

Consequently, continuous planning is required throughout an acquisition process as new knowledge is obtained and uncertainties are resolved.

No matter how good a requirements specification [is], it still won't be good enough. There needs to be an extra margin for changes that might have to be implemented throughout the project. (Interviewee 5, 2020)

Given that planning is ongoing, it is also necessary to be able to make swift decisions in order to prevent work from stagnating while waiting for decisions at higher management levels, thus avoiding inefficiency. Johnson (2022) describes the ability to make rapid decisions as the single most important factor in successful IT projects.

In addition, complex systems generally change over time, meaning that a solution to a problem within a complex system will not remain effective indefinitely (Ashby, 1958). This implies that C2 systems will require modifications both during development and after delivery. For example, changes may result from the operational environment, but also because users were initially unaware of their requirements, could not articulate implicit or intuitive needs, or were unable to express their needs in sufficient

detail (Albinsson, 2004; de Mesquita et al., 2023). The need for continuous development and updating of IT and C2 systems throughout their lifespan distinguishes these systems from other, simpler materiel systems, such as vehicles:

And something that we /.../ are starting to understand is that this product never really ends. We're not done in Q4 2023 and then like, 'now we're completely finished'. This is something that's going to be an ongoing, iterative process; it's going to keep developing, rolling out, getting updates all the time—it's a system that could last, like, 25 years. (Interviewee from Modern Approach Group 2, 2022)

Since projects are intended to have a defined end, this requirement for ongoing updates makes the project form a less suitable working approach for the acquisition of C2 systems.

According to Ashby's (1956) Law of Requisite Variety, effective control of a system requires that the controlling mechanism possesses at least as much variety as the system itself. In the context of C2 system acquisition, this implies that rigid, predetermined processes may be insufficient to manage the complexity and unpredictability inherent in C2 system acquisition. Instead, acquisition processes need to be adaptable, allowing decision-makers to respond flexibly to emerging challenges and unforeseen requirements

This supports the case for using more agile working methodologies during the acquisition of C2 systems. According to Johnson (2022), such methodologies have been shown to lead to successful IT projects three times more often than traditional methods. However, it is important to remember that deciding on one or a few specific 'agile' methods to use means that the organisation is no longer truly agile. It must be possible to choose from a variety of methods and even tailor a method to suit the ongoing acquisition.

Asked how effectiveness and efficiency could be achieved when acquiring C2 systems, one interviewee answered: 'Set a Goal, Allocate Resources, Don't Get in the Way.' This should be interpreted to mean that senior managers are to set

the goal at a strategic level, allocate the necessary resources, and then allow the acquisition teams to operate independently, without imposing processes that constrain their autonomy or delay decision-making.

Implementation Experiences and Impacts of the Seven Principles

The case study showed that, at the team level, the two Modern Approach Groups had incorporated the seven principles. However, at higher organisational levels, these principles had not been implemented. This created friction with other parts of the organisation and reduced the efficiency and effectiveness of the two Modern Approach Groups' work.

Despite this friction, interviewees argued that a return to traditional ways of working would lead to reduced efficiency:

[Interviewer]: If Modern Working Methods were to be discontinued and you were forced to work within the Swedish Armed Forces' processes, what would happen to the efficiency of this object?

[Interviewee 1]: It would die instantly.

[Interviewee 2]: No, it would probably take a few days. But no. It would become a matter of capability at a strategic level. (Dialogue between interviewees from Modern Approach Group 1, 2024)

When asked why Modern Working Methods were more efficient and effective, the following response was given:

This is largely because we are able to spend a greater proportion of our time actively addressing problems rather than waiting for permission to do so. Moreover, this approach fosters a culture in which individuals develop solutions and put forward proposals instead of first asking whether they may proceed and then waiting for a response.

As a result, rather than creating a queue of pending approvals, people are able to begin implementing and testing ideas immediately. (Interviewee from Modern Approach Group 1, 2022)

Results from additional interviews conducted in 2025 further suggested that the two Modern Approach Groups were perceived to work more effectively and efficiently than other teams. One stakeholder argued that the groups working according to Modern Working Methods are the only ones actually delivering anything that can be regarded as 'on time' within the current materiel domain, and that they also do so 'very well'. The other two stakeholders, each commenting on one of the two Modern approach groups, also stated that these groups were performing well and were effective.

Discussion

By moving beyond a strict focus on processes, this study adopts a broader perspective on improving the effectiveness and efficiency of C2 system acquisition. In doing so, seven principles aimed at enhancing acquisition effectiveness and efficiency were developed. Among these principles, one is overarching—Recognition of Acquisition as a Complex and Wicked Problem. This principle provides a foundation for the remaining six principles, which can be understood as specific recommendations for establishing an environment suitable for managing the complexity and wickedness associated with C2 system acquisition. The seven principles are method-neutral, meaning that any method that fulfils these principles may be used, provided that the chosen method is suitable for the acquisition at hand.

By emphasising adaptability, flexibility, tailored processes, and rapid decision-making, the principles make it possible, in accordance with the Law of Requisite Variety, to maintain sufficient variety in management and control to cope with the uncertainty, complexity, and wickedness associated with C2 system acquisition. In contrast, developing a detailed, prescriptive process intended to be applied to all C2 acquisitions is inconsistent with the Law of Requisite Variety. The idea of designing

such an 'optimal' process should therefore be abandoned.

Historically, complete and comprehensive requirements specifications early in a project have often been highlighted as a success factor in the acquisition of information systems (e.g., Hofmann & Lehner, 2001; The Standish Group, 1995). The results of this study, however, suggest the opposite, thereby supporting claims by Albinsson (2004), Brooks (2002), Fraser et al. (2007), Johnson (2022), and Modigliani (2015).

Early-developed requirements specifications involve considerable uncertainty, as future users rarely understand their needs at the outset. In addition, the external environment may change over time. As a result, requirements documents frequently need to be revised. This phenomenon is commonly referred to as requirements creep and is often portrayed as a source of delays and increased costs. However, such a view overlooks the fact that changes to requirements may reflect improved understanding, learning, and adaptation to emerging conditions.

Rather than being avoided altogether, requirements changes should therefore be regarded as both inevitable and necessary for an effective acquisition, provided that they remain within the constraints of project objectives such as time and cost. Attempts to prevent all changes to requirements may instead result in a final product that is poorly aligned with operational needs. By explicitly recognising and accommodating the need for evolving requirements, the proposed principles support the progressive emergence of requirements throughout the acquisition process. This enables the acquisition team to respond to unforeseen developments in the external environment and to integrate new insights—such as evolving operational needs—without causing significant negative effects on efficiency.

The implications of the proposed principles extend beyond the choice of development methodology and have important consequences for how acquisition activities are organised and governed. By advocating an incremental and evolutionary approach, acquisition is no longer viewed as a finite project culminating in the delivery of a final

system. Instead, C2 systems are understood as evolving socio-technical systems operating in a dynamic environment. Consequently, any delivered solution should be regarded as temporary and subject to continuous adaptation, improvement and replacement. This perspective blurs the traditional distinction between acquisition and sustainment. Once a system has been fielded, development activities continue through upgrades, technology insertion, replacement of obsolete components and adaptation to emerging operational needs.

An important consequence is that acquisition may be less suited to governance through projects with predefined end states. Evolutionary development is characterised by continuous change and moving targets rather than clearly defined completion criteria. Under such conditions, acquisition may also be considered as an enduring capability development endeavour within the line organisation, thereby reducing the risk of competence loss between acquisition and sustainment phases and enabling continuous adaptation throughout the system life cycle.

Limitations and scope of findings

This study combines a broad literature-based synthesis with an empirical exploration of the application of the resulting principles. The seven principles were derived from an extensive literature survey covering a wide range of domains, including military and civilian acquisition, software development, and project management. As such, the principles themselves rest on a relatively broad analytical foundation and may be considered theoretically transferable across contexts involving complex socio-technical systems.

However, the empirical component of the study—comprising interviews and a case study involving a limited number of groups within a single organisational context—provides only an initial indication of how these principles may influence effectiveness and efficiency in practice. The small sample size and context-specific nature of the case study limits the ability to draw statistically generalisable conclusions regarding the effects of the principles.

Consequently, while the principles can be understood as analytically grounded propositions, the evidence regarding their practical impact should be interpreted as exploratory and context-dependent. Further empirical research, involving a larger number of cases and more diverse organisational settings, is required to assess the extent to which the observed effects can be generalised.

Conclusion

The most important recommendation for enhancing efficiency and effectiveness in C2 system acquisition is to approach such acquisitions from a complexity perspective. Applying the seven principles outlined in this paper enables organisations to navigate the inherent uncertainties, interdependencies, and diverse stakeholder needs associated with C2 system acquisitions, thereby improving acquisition efficiency and effectiveness. This study challenges the traditional reliance on early, comprehensive requirements and the implicit radical rationalism that assumes complex systems can be fully controlled through detailed planning and linear reasoning. Instead, developing requirements incrementally allows insights gained during earlier stages to inform subsequent decisions, supporting adaptive learning and flexibility throughout the project. Such an approach also demands trust in those tasked with conducting the acquisition, providing them with a broad allocation of decision rights and the autonomy to organise their work in the way they deem most effective to achieve results. This approach also shortens the preparatory phase, allocating more time to system development. The study concludes that abandoning rigid rationalist approaches in favour of an incremental, adaptive methodology—permitting requirements and decisions to evolve as new knowledge emerges—is essential for acquiring C2 systems in complex, dynamic, and unpredictable operational environments.

References

- Alberts, D. S., & Hayes, R. E. (2006). *Understanding command and control*. Department of Defense Command and Control Research Program.
- Alberts, D., Bailey, T., Choinard, P., Tolk, A., Wheatley, G., & Wilder, J. (2002). *NATO code of best practice for C2 assessment*. Computational Modeling and Simulation Engineering Books, 13.
- Albinsson, P.-A. (2004). *Interacting with command and control systems: Tools for operators and designers* (Licentiate thesis). Linköping University, Sweden.
- Ardagna, D., & Francalanci, C. (2005). A cost-oriented approach for the design of IT architectures. *Journal of Information Technology*, 20(1), 32–51.
- Arvidsson, V., Holmström, J., & Lyytinen, K. (2012, June 10). Information systems strategy-as-practice. In *International Workshop on IT Artefact Design & Work Practice Intervention*, Barcelona, Spain.
- Ashby, W. R. (1956). *An introduction to cybernetics*. Chapman & Hall.
- Ashby, W. R. (1958). Requisite variety and its implications for the control of complex systems. *Cybernetica*, 1(2), 83–99.
- Baccarini, D., & Collins, A. (2003, December). Critical success factors for projects. In *Proceedings of the 17th ANZAM Conference* (pp. 2–5).
- Badampudi, D., Wohlin, C., & Petersen, K. (2015, April). Experiences from using snowballing and database searches in systematic literature studies. In *Proceedings of the 19th international conference on evaluation and assessment in software engineering* (pp. 1–10).
- Bannister, F., & Remenyi, D. (2009). Multitasking: The uncertain impact of technology on knowledge workers and managers. *Electronic Journal of Information Systems Evaluation*, 12(1), 1–12.
- Baskerville, R. (1999). Investigating information systems with action research. *Communications of the Association for Information Systems*, 2(19), 1–32.
- Batie, S. S. (2008). Wicked problems and applied economics. *American Journal of Agricultural Economics*, 90(5), 1176–1191.
- Beck, K., Beedle, M., van Bennekum, A., Cockburn, A., Cunningham, W., Fowler, M., ... Thomas, D. (2001, February). *Manifesto for agile software development*. <https://agilemanifesto.org>
- Berg, H., & Ritschel, J. D. (2023). The characteristics of successful military IT projects: A cross-country empirical study. *International Journal of Information Systems and Project Management*, 11(2), 25–44.
- Blackham, J. (2007). Dealing with ‘wicked problems’. *The RUSI Journal*, 152(4), 36–38.
- Brady, L. J. (1990). *Streamlining the DoD acquisition process—One more time*. Air University, USA.
- Breiding, R. N., & Brilla, E. K. (2016). *Deliberate agility: Streamlining airborne ISR sensor acquisition*. Air University, USA.
- Bringselius, L. (2019a). *Kultur och värdegrundarbete ur ett tillitsperspektiv*. Ministry of Finance, Sweden.
- Bringselius, L. (2019b). *Uppföljning ur ett tillitsperspektiv*. Ministry of Finance, Sweden.
- Claverie, B., & Desclaux, G. (2022). C2 – Command and Control: A Systems of Systems to Control Complexity. *American Journal of Management*, 22(2), 45–63.
- Brooks, F. P., Jr. (2002). No silver bullet: Essence and accidents of software engineering. *IEEE Computer*, 20(4), 10–19.
- Cockburn, A. (2000). Selecting a project’s methodology. *IEEE Software*, 17(4), 64–71.
- Cook, S. C., & Pratt, J. M. (2014). Towards designing innovative SoSE approaches for the Australian Defence Force. In *2014 9th International Conference on System of Systems Engineering (SOSE)* (pp. 295–300). IEEE.

- Cranfield, A. (2007). Some further thoughts on addressing 'wicked problems'. *The RUSI Journal*, 152(6), 72–75.
- Crowston, K., Annabi, H., Howison, J., & Masango, C. (2004). Effective work practices for software engineering: Free/libre open source software development. In *Proceedings of the 2004 ACM workshop on Interdisciplinary software engineering research* (pp. 18-26).
- Crowston, K., Howison, J., Masango, C., & Eseryel, U. Y. (2007). The role of face-to-face meetings in technology-supported self-organizing distributed teams. *IEEE Transactions on Professional Communication*, 50(3), 185–203.
- Currie, W. L., & Guah, M. W. (2007). Conflicting institutional logics: A national programme for IT in the organisational field of healthcare. *Journal of Information Technology*, 22(3), 235–247.
- Davidson, A. S. (2025, October 2). NITMRE: How the Air Force transformed command and control for air mobility. *Mattermost Blog*. Retrieved 9th December 2025 from <https://mattermost.com/blog/nitmre-air-force-command-control/>
- Delano, K. J. (1999). *Critical success factors for DoD program managers*. Air University, USA.
- Ellis, C., Lucas, R., Fawkes, B., Robson, M., Brown, A., Keedwell, E., & Black, J. (2024). *Command and control in the future: Concept paper 3—Conceptualising C2 as a capability*. RAND Corporation.
- Ellström, P.-E., & Kock, H. (2008). Competence development in the workplace: Concepts, strategies and effects. *Asia Pacific Education Review*, 9(1), 5–20.
- European Defence Agency. (2023). *Defence data 2022: Key findings and analysis*.
- Farrell, R., & Hooker, C. (2013). Design, science and wicked problems. *Design Studies*, 34, 681–705.
- Fraser, S. D., Brooks, F. P., Jr., Fowler, M., Lopez, R., Namioka, A., Northrop, L., ... Thomas, D. (2007). "No silver bullet" Reloaded: A Retrospective on "Essence and Accidents of Software Engineering". In *Companion to the 22nd ACM SIGPLAN conference on Object-oriented programming systems and applications companion* (pp. 1026–1030).
- Gareis, R., Huemann, M., & Marinuzzi, A. (2013). *Project management and sustainable development principles*. Project Management Institute.
- Granåsen, M., Hallberg, N., Josefsson, A., & Ivori, J. (2023, December). A future C2 concept designed to guide capability development. In *28th International Command and Control Research & Technology Symposium Proceedings*.
- Hasan, H. (2014). Complexity theory. In H. Hasan (Ed.), *Being practical with theory* (pp. 49–54). University of Wollongong.
- Hoehn, J. R. (2020). *Joint all-domain command and control (JADC2)*. Congressional Research Service.
- Hofmann, H. F., & Lehner, F. (2001). Requirements engineering as a success factor in software projects. *IEEE Software*, 18(4), 58–66.
- Jacobsen, D. I. (2002). *Vad, hur och varför?* Lund: Studentlitteratur.
- The Standish Group. (1995). *Chaos*. The Standish Group.
- Karlsson, F., & Ågerfalk, P. J. (2011). Towards structured flexibility in information systems development. *Journal of Database Management*, 20(3), 51–75.
- Kasser, J. E. (2002). The cataract methodology for systems and software acquisitions. In *Systems Engineering, Test & Evaluation Conference* (pp. 32–39).
- La Porta, R., Lopez-de-Silanes, F., Shleifer, A., & Vishny, R. W. (1996). *Trust in large organizations*. National Bureau of Economic Research.

- Laufer, A., & Hoffman, E. J. (1998). *Ninety-nine rules for managing "faster, better, cheaper" projects*. NASA. Retrieved 9th December 2025 from https://appel.nasa.gov/wp-content/uploads/2013/08/47990main_47442main_ninety-nine_rules.7.pdf
- Lawless, S. (2023, April 12). *What are the 9 guiding principles of ITIL® v3?*. Purple Griffon. <https://purplegriffon.com/blog/9-guiding-principles-ITIL>
- de Leon, J. D. (2021). Understanding multi-domain operations in NATO. *The Three Swords Magazine*, 37, 92–94.
- Lyytinen, K., Yoo, Y., Varshney, U., Ackerman, M. S., Davis, G., Avital, M., Robey, D., Sawyer, S., & Sørensen, C. (2004). Surfing the next wave: Design and implementation challenges of ubiquitous computing environments. *Communications of the Association for Information Systems*, 13, 697–716.
- Marple, J., Clark, B., Jones, C., & Zubrow, D. (2001). *Measures in support of evolutionary acquisition*. Booz Allen Hamilton.
- Modigliani, P. (2015). Portfolio acquisition: How the DoD can leverage the commercial product line model. In *Proceedings of the Fourteenth Annual Acquisition Research Symposium* (pp. 249–266).
- Muller, J. Z. (2018). *The tyranny of metrics*. Princeton University Press.
- Nelson, H. G., & Stolterman, E. (2012). *The design way: Intentional change in an unpredictable world* (2nd ed.). MIT Press.
- Nylén, D., & Holmström, J. (2015). Digital innovation strategy: A framework for diagnosing and improving digital product and service innovation. *Business Horizons*, 58(1), 57–67.
- Pinto, J. K., & Slevin, D. P. (1989). Critical success factors in R&D projects. *Research-Technology Management*, 32(1), 31–35.
- Poli, R. (2013). A note on the difference between complicated and complex social systems. *Cadmus*, 2(1), 142–147.
- Reagan, R. B., & Rico, D. F. (2010). Lean and agile acquisition and systems engineering: A paradigm whose time has come. *Defense AT&L*, November–December, 48.
- Rittel, H. W. J., & Webber, M. M. (1973). Dilemmas in a general theory of planning. *Policy Sciences*, 4(2), 155–169.
- Salwin, A. E. (1998). *The road to successful ITS software acquisition - Volume II: Software Acquisition Process Reference Guide* [FHWA-JPO-98-036]. Mitretek Systems.
- Sauer, C., Gemino, A., & Reich, B. H. (2007). The impact of size and volatility on IT project performance: Studying the factors influencing project risk. *Communications of the ACM*, 50(11), 79–84.
- Sauer, C., & Reich, B. H. (2009). Rethinking IT project management: Evidence of a new mindset and its implications. *International Journal of Project Management*, 27(2), 182–193.
- Schein, E. H. (2010). *Organizational Culture and Leadership* (4th ed.). John Wiley & Sons.
- Šmite, D., Moe, N. B., & Ågerfalk, P. J. (2010). Fundamentals of agile distributed software development. In D. Šmite, N. B. Moe, & P. J. Ågerfalk (Eds.), *Agility across time and space: Implementing agile methods in global software projects* (pp. 3–7). Springer.
- Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, November, 1–12.
- Sommerville, I., Cliff, D., Calinescu, R., Keen, J., Kelly, T., Kwiatkowska, M., McDermid, J., & Paige, R. (2012). Large-Scale Complex IT Systems. *Communications of the ACM*, 55(7), 71–77.
- Stanton, N. A., & Baber, C. (2007). The ergonomics of command and control. *Ergonomics*, 49(12–13), 1131–1138.
- The Standish Group. (1995). The Standish Group.

U.S. Marine Corps. (2018). *Command and control*. Department of the Navy.

Ward, D. (2014). *F.I.R.E.: How fast, inexpensive, restrained and elegant methods ignite innovation*. Harper Business.

Wong, J. P., Younossi, O., Lacoste, C. K., Anton, P. S., Vick, A. J., Weichenberg, G., & Whitmore, T. C. (2022). *Improving defense acquisition: Insights from three decades of RAND research*. RAND Corporation.

Author biography

Mr Johan Nordström

Dr Susanna Nilsson

Mr Jonas Herkevall

Ms Moamia Gideskog

Johan Nordström is a Senior Scientist and project manager at the Swedish Defence Research Agency (FOI). His research focuses on command and control support system acquisition and assessments of socio-technical command and control systems. He also manages practical projects related to organisational adaptation during the introduction of new materiel. Prior to joining FOI in 2017, he served for over 16 years as a fire officer in the Swedish rescue services, including senior command roles. He holds an MSc in Ergonomics from KTH Royal Institute of Technology, as well as degrees in Fire Protection Engineering and Education. In addition, he serves as a field engineer in a Swedish Home Guard battalion headquarters and has operational experience from the national crisis response during the 2018 forest fires in Sweden.

Susanna Nilsson is a Senior Scientist at the Swedish Defence Research Agency (FOI). She holds a PhD in Cognitive Systems Engineering from Linköping University, a MSc in Cognitive Science from Umeå University and has worked as a researcher at FOI since 2010. Previously she has also worked as a researcher and lecturer at Linköping University. At FOI she has led and worked in research projects in areas such as user-centered system development and human computer interaction. She has also worked in projects related to C2 system acquisition, development of sociotechnical C2 systems, and user interfaces in different domains.

Jonas Herkevall is a Senior Scientist at the Swedish Defence Research Agency (FOI). He holds a MSc in Cognitive Science from Linköping University. Since joining FOI in 2021 he has been part of research and concept development projects focused on future C2, multi-stakeholder cooperation in crisis management, and organisational adjustments following the introduction of new materiel. Since 2023 he has been part of establishing a research program for Swedish civil defence at FOI and led research projects focused on the intersection of governance, defence and security.

Moamia Gideskog is a Research Engineer at the Swedish Defence Research Agency, Division of Cyber Defence and C2 Technology. She holds an MSc in Mechanical Engineering, specialising in Product Development and Design, from Linköping University. Since joining FOI in 2023, she has taken part in multiple research and concept development projects. These projects have focused on user-centred methods for assessing socio-technical C2 systems, enhancing C2 system acquisition, and field testing for the introduction and development of new materiel for the Swedish Armed Forces.

Authors Note

Declaration of conflicting interests: The authors declare no potential conflicts of interest with respect to the research, authorship, or publication of this article.

Funding: This work was supported by the Swedish Armed Forces (AT.9220216, AT.9220217, AM.4192020).

Acknowledgements: All quotes from the interview study in 2020 have been translated from Swedish into English by Sofia McGarvey, Swedish Defence Research Agency. Professor Niklas Hallberg, Swedish Defence Research Agency has reviewed the paper critically for important intellectual content. Ms Mari Olsén and Ms Maria Wikström, Swedish Defence Research Agency has assisted in the acquisition and analysis of data.

Use of AI: Quotes from the focus groups conducted in 2023 and 2024 have been translated from Swedish into English using ChatGPT. These translations were subsequently reviewed by the authors of this paper.

Corresponding:

Johan Nordström

Postal address: FOI, Totalförsvarets forskningsinstitut, 164 90 Stockholm, SWEDEN

Telephone: +46 (0)13 378 331

E-mail address: johan.nordstrom@foi.se

Authors:

Editor in Chief: COL Mietta Groeneveld (NLD A)
Deputy-Editor in Chief: LtCol Rabia SAYLAM, PhD (TUR AF)
Managing Editor: LTC Hendrik ten Hoeve (NLD A)
Academic Editors: Dragos Calitoiu, PhD (CAN CIV), Ahmed Ghanmi, PhD (CAN CIV),
Mr. Frank Gubbels (NLD CIV)
Formatting: SM Marvin Schaap (NLD A), Max van Rijn (NLD CIV)

NATO Command and Control Centre of Excellence

Kromhout Barracks, Building K1,
Herculeslaan 1, 3584 AB Utrecht
The Netherlands c2coe.org

This product is co-sponsored by the NATO Command and Control Centre of Excellence (NATO C2COE) and the Canadian Joint Warfare Centre (CJWC). It does not necessarily represent opinions or policies of NATO and/or the Sponsoring Nations of the Centre and it is designed to provide an independent opinion

©2026 NATO C2COE. All rights reserved. This publication may not be reproduced or used in any manner whatsoever without reference to the NATO C2COE.

