

The Need for a European Dual-Use Cognitive Intelligence Centre: Why the EU Needs Vanguard Leaders

Dražen Kapusta¹ and Damir Stručić²

¹ COTRUGLI Business School · ² Croatian Defence Academy “Dr. Franjo Tuđman”

ABSTRACT

Europe's security architecture faces a cognitive frontier: adversaries now wage conflict by degrading institutional sensemaking, manufacturing social consensus, and exploiting the seams of democratic deliberation. The EU has advanced important policy responses (the Strategic Compass, hybrid toolbox instruments, and dual-use technology frameworks), yet a structural gap persists: no pan-European institution systematically trains, measures, and improves decision superiority under cognitive attack across civil-military boundaries. This article proposes a European Dual-Use Cognitive Intelligence Centre (EU-DUCCI) to close that gap. Drawing on NATO's scientific findings on cognitive warfare, the Cotruglian tradition of merchant ethics (1458), the Vanguard Leadership framework, and the operational instruments HAI5 and VIS, the article argues that cognitive resilience is an institutional capability that can be trained and measured, rather than a cultural predisposition. The article offers a conceptual framework grounded in a mixed methodology combining Vanguard Intelligence Summary analysis, comparative institutional analysis, and Cotruglian historical analysis; outlines a Decision Theatre architecture; proposes metrics and governance requirements; and maps a pragmatic eighteen-month implementation path.

KEYWORDS: *cognitive warfare; hybrid threats; FIMI; decision superiority; mission command; human-AI collaboration; Vanguard Leadership; dual-use; EU strategic autonomy; resilience*

1 Introduction: The EU's Cognitive Vulnerability as a Strategic Constraint

The European Union's capacity to act, politically, economically, and militarily, rests on the integrity of its decision-making systems. In contemporary strategic competition, adversaries increasingly seek advantage not only through kinetic force, but by weaponizing perception, institutional trust, and collective attention. This is the logic of cognitive conflict: shaping how publics and leaders interpret events, what they consider achievable, and how quickly they can coordinate a collective response. When cognitive operations succeed, they do more than confuse individual decision-makers; they fragment the shared epistemic foundation that institutional coordination requires.

A recent NATO scientific paper formalizes this challenge, arguing that cognitive warfare targets cognition itself, seeking to affect or disrupt human decision-making and behavior through military

and non-military means, amplified by advanced technologies and hybrid influence operations (Du Cluzel, 2025). The dual nature of the threat is built into that definition: cognitive attacks cut across civilian and military populations simultaneously and therefore demand whole-of-government responses spanning research, policy, training, and resilience planning.

The EU recognizes the significance of this challenge in its policies. The Strategic Compass commits to strengthening resilience against hybrid threats and foreign information manipulation and interference (FIMI) and to improving situational awareness and response capabilities (Council of the European Union, 2022). Legal and industrial frameworks for managing dual-use technologies have been updated to manage the synergies of civil-military research (European Commission, 2024a, 2024b). These are genuine achievements.

Yet a specific gap in institutional capability remains that existing instruments do not address. Europe produces analyses, policies, and legal architecture, but it lacks a pan-European dual-use platform for repeatedly training decision superiority at scale, under realistic conditions of cognitive threat, with measurable performance outcomes and post-action learning at a research-grade level. This is not a hardware gap. It is a human-institutional gap: the absence of a dedicated centre for developing and sustaining the trained cognitive resilience that democratic governance and credible defense require.

This article advances a specific proposal: the European Dual-Use Cognitive Intelligence Centre (EU-DUCCI), a platform combining accredited education, immersive training and exercises through a Decision Theatre, applied research and innovation, and a pan-European community of practice. The centre's doctrine is Vanguard Leadership (VL), operationalized through two instruments: the Vanguard Intelligence Summary (VIS) and the HAI5 framework for human-AI integration. The intellectual lineage underpinning the proposal reaches back to Benedetto Cotrugli's merchant philosophy of 1458: an institution that cannot maintain clear epistemic accounts of what it knows, what it assumes, and what it has decided will fail precisely at the moment when strategic performance is most needed.

The article is structured as follows. Section 2 examines the mechanism of cognitive warfare as an industrialized weapon of social proof. Section 3 reviews the EU policy baseline and characterizes the remaining institutional gap. Section 4 argues that mission command logic now extends into the civilian domain. Section 5 presents Vanguard Leadership as a dual-use doctrine for cognitive conflict. Section 6 presents the mixed methodology. Section 7 sets out the EU-DUCCI proposal, including the Decision Theatre architecture. Section 8 proposes an evaluation framework. Section 9 addresses governance and democratic oversight. Section 10 examines strategic urgency and funding pathways. Sections 11 and 12 provide the discussion and conclusions.

2 Cognitive Warfare as Industrialized Social Proof

2.1 The Social Psychology of Collective Belief

To understand why cognitive conflict constitutes a structural vulnerability for democratic institutions, the analysis must be grounded in two converging bodies of knowledge: the social psychology of collective belief formation and the operational mechanics of contemporary influence operations.

The foundational insight comes from crowd psychology. Le Bon (1895/2002) observed that group context reduces individual critical thinking, amplifies impulsivity, and increases susceptibility to emotional contagion. Under conditions of ambiguity or threat, individuals tend to seek epistemic cues from others, a tendency formalized as social proof (Cialdini, 1984). This is not a defect of weak minds; it is an adaptive cognitive heuristic that operates under conditions of information scarcity and time pressure, which are precisely the conditions that adversarial cognitive operations seek to manufacture.

Contemporary disinformation research deepens this analysis. Lewandowsky et al. (2017) show that corrective information often fails to displace false beliefs because it competes with identity, emotional investment, and repetition effects. Kahneman's (2011) distinction between fast associative processing and slow, effortful analysis illuminates the mechanism: cognitive operations activate System 1 (fast, emotionally resonant, susceptible to social cues) while suppressing the System 2 processing that would reveal their artifice. The problem is not merely that false beliefs circulate, but that repeated exposure to manufactured uncertainty progressively erodes institutional epistemic discipline.

2.2 The Claque as Historical Prototype

The strategic exploitation of social proof has a longer history than is commonly recognized. An instructive analogue, important for the Cotruglian methodological thread of this article discussed in Section 6.3, is the claque system in eighteenth- and nineteenth-century European theatre. Professional clagues were organized groups of paid applauders, systematically distributed to produce emotional reactions in the audience and, crucially, to shape quality judgments among other audience members. The mechanism was not deception in the usual sense; it was the engineering of social context to manipulate individual inference.

This is not a historical curiosity. The claque is the operational ancestor of contemporary coordinated inauthentic behavior on digital platforms: bot networks, like storms, coordinated amplification campaigns, and algorithmically distributed narrative seeding. The scale differs; the cognitive mechanism is identical. Contemporary cognitive operations industrialize social proof at network speed, targeting not individual minds but the epistemic infrastructure of institutions. What Cotrugli (1458/2017) would have recognized as the corruption of a merchant's trusted market signal, manufactured reputation, is now executed at civilizational scale by state and non-state actors alike.

2.3 Strategic Consequences for EU Institutions

The strategic consequence for democratic governance is structural. Cognitive operations do not merely inject false beliefs into circulation; they fragment the shared epistemic foundation that collective decision-making requires. When publics and leaders operate from incompatible factual frames, consensus becomes unavailable, and institutional delay becomes structurally embedded in the response process. This is precisely the outcome adversaries seek: not to win the argument, but to render the argument unresolvable.

Research on the Russian approach to influence operations in Europe documents a consistent pattern of platformized, deniable operations designed to normalize interference below thresholds that would trigger a decisive collective response (Wardle & Derakhshan, 2017). The goal is not to convince European publics of any particular position, but to inject enough ambient uncertainty into their views that coordinated action becomes politically infeasible. For EU institutions, the practical

implication is that the challenge is not primarily informational. The question is not how to win information contests, but how to build organizations whose decision-making processes are structurally resilient to cognitive degradation.

3 The EU Policy Baseline and the Remaining Institutional Gap

The EU has not been passive in responding to the cognitive dimensions of hybrid threats. A sustained policy effort has produced a substantial institutional and legal architecture whose achievements and limits must be described precisely.

The Strategic Compass for Security and Defence (Council of the European Union, 2022) is the most comprehensive strategic document the EU has produced in the security domain. It explicitly commits to strengthening resilience against hybrid threats, including FIMI, improving situational awareness through intelligence fusion, and developing response capabilities across the full spectrum of hybrid instruments, acknowledging their civil-military dual character. Separately, the updated frameworks for managing dual-use technologies reflect a sophisticated understanding of the tension between open research environments and security requirements (European Commission, 2024a, 2024b).

Instrument-level responses have also been significant. The East StratCom Task Force monitors and counters disinformation. The Hybrid Fusion Cell integrates intelligence from civilian and military domains. The Digital Services Act creates platform accountability requirements that reduce the structural advantages currently enjoyed by coordinated inauthentic behavior campaigns. Horizon Europe and the European Defence Fund both contain provisions relevant to cognitive security research.

Despite this breadth, a specific capability gap remains, better described as a human-institutional gap than a hardware gap. The EU's policy and legal architecture produces analyses and reports; it does not systematically produce a population of leaders and institutions trained to sustain high-quality sensemaking and aligned action under cognitive attack. Training in this domain exists, but it is inconsistent, without standardized measurement, and without research-grade feedback loops that would enable systematic capability improvement over time.

This gap is aggravated by the human-AI dimension. As AI systems become increasingly embedded in intelligence fusion, crisis management, and strategic analysis, new capabilities arrive together with new cognitive vulnerabilities. Automation bias, the documented tendency of human operators to defer to algorithmic recommendations even when they are wrong or adversarially manipulated, represents a cognitive vulnerability that did not exist at the institutional level a decade ago (Brčić, 2025; Eftimov et al., 2024). Managing this vulnerability requires not only technical safeguards but trained human judgment. The institutional gap is therefore threefold: a gap in training and exercises, a gap in measurement, and a gap in the governance of the human-AI relationship.

4 Mission Command Logic Extended to the Civilian Domain

The doctrinal logic of mission command was developed to solve a persistent problem in military operations: how to sustain effective coordinated action under high uncertainty, degraded communication, and rapid changes of tempo. Hierarchical approval processes fail under these conditions: they are too slow, too dependent on accurate upward information flows, and brittle when

commanders become unavailable or overloaded. Mission command answers with disciplined autonomy: subordinates are empowered to act within an understood intent without requiring approval for every decision, because they are trained to understand purpose and to exercise judgment about means while remaining aligned on ends.

In the era of AI-driven operations, this logic has acquired new urgency. Algorithmic systems generate recommendations faster than traditional approval processes can evaluate them. As Kapusta (2026b) argues, the paradox is educational and institutional: AI promises decision advantage, but without developing leaders for human-AI collaboration, advanced systems become vulnerabilities rather than capabilities. The organizations that will operate most effectively in this environment are those that have already developed the cultural infrastructure (shared mental models, clear intent frameworks, mutual trust) that enables distributed action under uncertainty.

The thesis of this section is that mission command logic now extends beyond the defense domain. The conditions that make it necessary, high uncertainty, degraded information environments, rapid tempo, and distributed actors, increasingly characterize major EU institutional challenges. The coordination demands of the COVID-19 pandemic, the response to the 2022 energy crisis following Russia's invasion of Ukraine, the management of large-scale migration, and cyber incidents affecting critical infrastructure all involved distributed actors across civil-military boundaries, contested information environments, adversarially fueled narratives, and decision timelines that exceed the capacity of hierarchical coordination.

In these contexts, the effective institutions were those able to act with incomplete information, maintain alignment across distributed stakeholders, resist manufactured narratives, and preserve decision quality under pressure. These are not qualities that arise naturally from democratic culture; they are cultivated capabilities that require deliberate development. The adaptation of mission command's doctrinal logic, clarity of intent, distributed authority, mutual trust, and disciplined initiative, offers civilian governance contexts a structural foundation for cognitive resilience in EU institutions.

That adaptation demands careful handling. Mission command in the context of democratic governance cannot replicate the command-authority relationships of military organizations. Democratic oversight, the rule of law, and civil liberties impose constraints that are constitutive requirements of legitimate governance. Adapting mission command logic to civilian contexts therefore requires explicit attention to accountability and transparency, even when decision speed is at a premium. Section 9 addresses these governance requirements.

5 Vanguard Leadership as a Dual-Use Doctrine for Cognitive Conflict

5.1 Theoretical Foundations

Vanguard Leadership (VL) was developed at COTRUGLI Business School in response to three converging structural shifts: the exponential acceleration of change in the NEO era (Networked, Exponential, Orchestrated); the militarization of commercial and institutional competition through asymmetric tactics and cognitive operations; and the integration of artificial intelligence as simultaneously a force multiplier and a source of organizational risk (Kapusta, 2025a).

The framework's theoretical foundations are threefold. First, dynamic capabilities theory (Teece et al., 1997) provides the economic and organizational grounding: in rapidly changing environments, competitive advantage derives from the capacity to sense, seize, and transform, that is, to detect signals before they become crises, to act within ever narrower windows, and to restructure capabilities coherently. Second, Boyd's (1986) OODA loop (Observe, Orient, Decide, Act) provides the operational architecture for decision-making under uncertainty and tempo pressure. Organizations that cycle through OODA faster than their adversaries gain a decisive advantage independent of relative resources; in cognitive warfare, the key resource is not materiel but the quality and speed of sensemaking. Third, and distinctively, the NEO Cotruglian philosophical tradition provides the ethical and motivational foundation, a foundation discussed more fully in the methodology section.

The VL framework operationalizes these foundations through four pillars. Resilience refers not merely to psychological toughness but to the organizational capacity to recover decision quality after cognitive-pressure events, an outcome that can be trained and measured. Critical strategic thinking encompasses disciplined hypothesis testing, structured reasoning, and rapid assumption-falsification cycles designed to preserve epistemic integrity faster than adversaries can manufacture uncertainty. AI augmentation refers to deploying artificial intelligence as a genuine force multiplier while managing automation bias, explainability requirements, and the risks of adversarial manipulation. Tribal Collaboration describes the cultivation of communities of practice and cross-organizational trust networks that reduce institutional susceptibility to manufactured consensus by providing alternative epistemic anchors during information-disorder events.

5.2 The Dual Character of Vanguard Leadership

The claim that VL constitutes a dual-use doctrine requires elaboration. Applied to leadership doctrine, dual-use means that the same cognitive vulnerabilities, susceptibility to manufactured consensus, automation bias, and epistemic fragmentation under information overload, manifest simultaneously in corporate boards, government ministries, and operational headquarters. The same doctrinal requirements, clarity of intent, distributed authority, and disciplined hypothesis testing, hold in all of these contexts. The difference lies in scale, stakes, and the specific character of adversarial operations, not in the underlying cognitive architecture.

For the proposed centre, this dual character has a practical institutional implication: the centre can serve both civilian and defense constituencies without requiring separate doctrines or training methodologies. The core training content is the same; scenarios are calibrated to context. This improves the centre's economics and, more importantly, enables the cross-fertilization of civilian and military decision practices that produces genuine interoperability, a shared cognitive culture rather than mere procedural compatibility.

5.3 VIS and HAI5 as Operational Instruments

Two operational instruments translate VL doctrine into repeatable practice. The Vanguard Intelligence Summary (VIS) is a structured analytic method adapted from military-intelligence practice for executive and institutional decision-making (Kapusta, 2025b). VIS imposes explicit decision-process discipline: hypotheses and options must be articulated separately from evidence; evidence items must be assigned to specific hypotheses; assumptions must be explicitly tracked, with kill indicators that would falsify them; and decision summaries must distinguish what is known from what is assessed and what is assumed. This structure does not slow decision-making; among

trained practitioners it accelerates decisions by removing the cognitive overhead of unstructured debate and by reducing the surface available for adversarial narrative insertion.

HAI5 is a staged methodology for human-AI integration framed as trust infrastructure (Kapusta & Stručić, 2026). The HAI5 architecture links three domains: technology (AI agents, personalization systems, verifiability mechanisms), governance (auditability requirements, accountability structures, alignment frameworks), and people (academia, the Decision Theatre, and communities of practice). Together, VIS and HAI5 answer the central academic critique of practitioner frameworks: the lack of operational specificity that would enable replication, measurement, and scientific evaluation. VL is not merely a leadership philosophy; it is a training and governance architecture with explicit methods, measurable outcomes, and institutional requirements.

6 Methodology

This article employs a mixed methodology combining three interrelated analytic approaches: structured Vanguard Intelligence Summary (VIS) analysis, comparative institutional analysis, and Cotruglian historical analysis. The combination is not an aggregation of convenient methods; each approach addresses a distinct dimension of the research problem, and the three streams converge on a joint claim that no single approach could sustain on its own. This section describes each approach, its application in the article, and the epistemic discipline it imposes on the overall argument.

6.1 Structured Vanguard Intelligence Summary (VIS) Analysis

VIS, developed as a practitioner methodology adapted from the military-intelligence analytic craft (Kapusta, 2025b), is applied as a methodological discipline throughout the analysis. Its contribution to the article's methodology is the imposition of explicit epistemic accountability on every significant claim: hypotheses are stated separately from evidence; evidence items are assigned to specific claims; assumptions are flagged wherever empirical support is unavailable or contested; and kill indicators are identified that would materially weaken the article's central argument.

The article's central hypothesis, that a dedicated EU-DUCCI would close the identified cognitive-resilience gap more effectively than increased funding of existing institutions, is treated as a falsifiable claim rather than a normative statement. Evidence assigned to that hypothesis includes: NATO's documented capability-gap analysis (Du Cluzel, 2025), the absence of pan-European civil-military capacity for cognitive training in the current EU institutional landscape, and operational evidence from the application of VL across 45 COTRUGLI MBA cohorts and the deployment of HAI5. The assumptions underlying the hypothesis, that VL transfers effectively to EU institutional contexts, that Decision Theatre scenarios can be validated for research-grade measurement, and that the dual-use design is institutionally sustainable, are stated explicitly, and kill indicators for them are given in the limitations section.

This approach addresses a structural risk in security-policy proposals: the tendency of conceptual work to accumulate supporting evidence while suppressing disconfirming considerations. VIS discipline requires that disconfirming evidence and unfalsified assumptions be visible to readers and evaluators. The article's limitations section (Section 12) is therefore not a conventional academic hedge; it is an integral component of the VIS analytic framework applied to the article's own central claim.

6.2 Comparative Institutional Analysis

The comparative institutional analysis draws on three cases of partial but instructive relevance to the proposed centre. The first is NATO capability development in the emerging cognitive domain, examined through publicly available reports of NATO's Science and Technology Organization. This case provides a defense-sector reference point for what an institutionalized cognitive-resilience program looks like at scale, including its doctrinal, governance, and measurement dimensions. It also illustrates the limitation of an exclusively defense-sector approach: cognitive attacks directed at civilian populations and civilian governance institutions require a training architecture that NATO, by mandate and culture, cannot fully provide.

The second is dual-use technology governance in the EU, examined through the Commission's White Papers and regulatory documents (European Commission, 2024a, 2024b). This case provides the civilian-governance context within which the proposed centre must operate, including the funding landscape, ethics and oversight requirements, and the EU's institutional culture of capability development. It also illustrates the gap between technology governance (which the EU does well) and human capability development (which remains inconsistent and unmeasured).

The third is COTRUGLI Business School's operational experience with the VL framework across ten years of executive MBA programs, the deployment of HAI5, and the development of the CO-lab innovation platform (Kapusta, 2025a). This case provides practitioner evidence of what structured leadership training for NEO-era conditions can and cannot achieve, including documented performance improvements in structured decision-making across 45 cohorts spanning the pharmaceutical, technology, energy, and financial sectors. This case is not presented as evidence of effectiveness in cognitive-warfare contexts specifically, which would constitute methodological overreach; it is presented as evidence of institutional feasibility and doctrinal coherence.

The comparative analysis is conducted explicitly through the VIS lens: evidence from each case is assigned to specific claims in the article's argument, and comparability limits are recorded wherever the cases differ in ways that constrain inference.

6.3 Cotruglian Historical Analysis

The Cotruglian historical thread is this article's most distinctive methodological contribution and requires the most careful explanation. Cotrugli's *Libro del arte dela mercatura* (written in 1458, published posthumously in 1573; modern critical edition 2017) is the earliest known systematic treatment of business practice and commercial ethics in Western literature. Written in Ragusa (modern Dubrovnik) by a merchant-diplomat who operated across Mediterranean trading networks, it codifies the conditions of trustworthy commerce: how merchants maintain trust at a distance, manage uncertainty, evaluate counterparties, and sustain epistemic discipline about the difference between verified information and rumor.

The methodological argument for applying Cotruglian historical analysis to a twenty-first-century cognitive-security proposal rests on a specific claim: that the cognitive vulnerabilities targeted by contemporary influence operations, susceptibility to manufactured reputation, degradation of epistemic discipline, and the confusion of manufactured consensus with market signal, are not novel pathologies of the digital age. They are recurring features of any information environment in which the cost of producing false signals is lower than the cost of verifying them. Cotrugli's merchant operated in exactly such an environment, and the ethical and practical disciplines he codified,

verified sources, the explicit separation of the known from the assumed, and trust as commercial infrastructure rather than moral luxury, map directly onto the VIS analytic framework and the four pillars of VL.

The claque analogy introduced in Section 2.2 is an example of the Cotruglian historical method: the use of a historical operational model to make visible a causal mechanism, the engineering of social proof, whose contemporary digital manifestations are harder to perceive precisely because of their scale and speed. The World Business Museum in Zagreb, which holds the world's largest institutional archive of pre-industrial business history (covering commercial practice from Mesopotamia in the fourth millennium BCE to the industrial revolution of the eighteenth century, with the Cotruglian legacy as its scholarly centerpiece), provides the primary-source foundation for this analytic approach. No comparable archive exists, which makes the Cotruglian historical thread a methodological contribution that other research groups cannot replicate.

Cotruglian historical analysis does not claim that fifteenth-century merchant philosophy provides direct operational guidance for twenty-first-century cognitive warfare. The claim is more specific: that the recurring structural problem of maintaining epistemic integrity in adversarial information environments generates persistent practical solutions, and that documenting those solutions across five and a half centuries provides both intellectual legitimacy for VL doctrine and a longer empirical base than any contemporary study can offer.

6.4 Methodological Limitations

This mixed methodology is appropriate to the article's purpose, advancing a conceptual and institutional proposal before empirical testing at the required scale is possible, but it carries limitations that must be acknowledged. The article provides no controlled empirical test of VL's effectiveness in cognitive-warfare contexts; such testing requires scenario exercises with measurement instruments that do not yet exist at the EU institutional level. It provides no cost-benefit analysis of the proposed centre against alternative responses; the necessary data do not exist in the public domain. And historical analogies, however illuminating of causal mechanisms, cannot substitute for prospective evaluation of training effectiveness. These limitations are real; the kill indicators in Section 12 state the empirical conditions that would materially weaken the article's central argument.

7 The European Dual-Use Cognitive Intelligence Centre: The Proposal

7.1 Core Concept and Rationale

The European Dual-Use Cognitive Intelligence Centre (EU-DUCCI) is proposed as a pan-European platform combining accredited education, immersive training and exercises, applied research and innovation, and a self-sustaining community of practice. Its primary mission is to develop, institutionalize, and continuously improve the EU's institutional capacity for decision superiority under cognitive attack.

The centre is dual-use in two senses. First, it serves both civilian and defense constituencies: government ministries, military commands, critical infrastructure operators, judicial institutions, regulatory bodies, and business school faculties alike. Second, it combines educational and research functions in a configuration designed for the creation and application of knowledge, not merely its

dissemination. The institutional rationale is grounded in the capability-gap analysis of Section 3: existing EU instruments produce analyses and policies, but they do not systematically produce trained cognitive resilience at the level of institutional decision-making.

7.2 Mission and Functions

The centre is designed as an institution with three functions. The first function is cognitive education, training, and exercises (Cog-ET&E): structured, accredited programs and immersive scenario exercises conducted in a Decision Theatre environment. Target populations span four groups: senior civil servants and officials of EU institutions; military commanders and intelligence professionals; critical infrastructure managers; and faculty of EU business and public administration schools, who will in turn train the next layer of leaders. The training is designed to develop three capabilities: decision superiority under cognitive pressure, calibration of the human-AI relationship, and speed of cognitive-resilience recovery.

The second function is research and innovation (R&I): applied research on cognitive performance under adversarial conditions, AI-supported decision-support analytics, training methodology development, and metrics of decision superiority and resilience. Research outputs include peer-reviewed publications, measurement instruments, training scenario repositories, and doctrinal notes. The R&I function serves as the validation mechanism for the training function: Decision Theatre exercises generate data on decision quality under pressure; those data feed research; research findings refine training methodology. This virtuous loop is what distinguishes a serious capability-development institution from an educational program.

The third function is a community of practice: a pan-European network of instructors, analysts, and leaders who continuously develop standards, share scenario experience, and maintain an institutional knowledge base. The community of practice has an additional strategic function: it creates European epistemic infrastructure, a network of individuals with shared analytic frameworks, mutual trust, and common standards for evidence evaluation, which is itself a form of cognitive resilience at the civilizational level.

7.3 Decision Theatre: Architecture and Function

The Decision Theatre is the operational engine of the proposed centre: an immersive, instrumented environment in which teams conduct scenario exercises under realistic conditions of cognitive threat, generating research-grade data on decision quality for analysis and continuous improvement.

The architecture is designed around four requirements. Information richness: participants operate in environments that simulate the information conditions of real cognitive-attack events, large in scale, contested in reliability, adversarially curated, and algorithmically shaped. Time compression: exercises are conducted under time pressure calibrated to the shift from System 2 to System 1 processing, because it is precisely in that shift that cognitive vulnerabilities manifest most strongly. AI integration: AI-based decision-support tools are present in the exercises, enabling direct training and measurement of human-AI calibration and automation bias. Measurability: all decisions, decision processes, and information interactions are recorded for after-action analysis and research-level analysis.

Scenario design draws on two sources: an existing repository of documented cognitive operations against EU member states (adapted and anonymized) and future-oriented scenarios developed by the

R&I function to anticipate adversarial modalities not yet deployed operationally but technically feasible given current AI capabilities.

The after-action review (AAR) is treated as the core learning intervention rather than an administrative procedure. VIS discipline applies to the AAR: assumptions made during the exercise are evaluated against pre-specified kill indicators; decision quality is assessed against explicit rubrics; human-AI calibration is measured; and cognitive-resilience recovery is tracked from peak pressure events through restored performance. AAR data feed both individual development plans and institutional research.

7.4 Institutional Location and Implementation Path

The centre's institutional anchor should satisfy three criteria: geographic accessibility to EU and NATO institutions, academic and research credibility, and existing civil-military collaboration infrastructure. The University of Defence and Security “Dr. Franjo Tuđman” (SOIS) in Zagreb represents a possible anchor, given its research mandate in the defense sector and Croatia's position as simultaneously an EU member state and a NATO ally. A partnership with COTRUGLI Business School, which brings ten years of VL training practice, the World Business Museum archive, and the CO-lab innovation platform, would supply the civilian-education and commercial-sector dimensions that a purely military-academic anchor would lack.

A pragmatic three-step implementation path is proposed to avoid the declarative failures of partnerships that afflict EU capability-development initiatives. Step one: a joint SOIS-COTRUGLI working group designs the Cognitive Intelligence Academy, learning outcomes, target populations, certification frameworks, and measurement instruments. Step two: a pilot Decision Theatre exercise, one scenario, one cohort, conducted under full measurement conditions with published results and public after-action documentation. Step three: curricular integration and funding applications, using the pilot results to support applications to the European Defence Fund, Horizon Europe, IPCEI-AI, and Digital Europe instruments. The target timeline from institutional commitment to the first operational cohort is eighteen months.

8 Metrics and Evaluation Framework

Academic evaluators and policy sponsors rightly interrogate capability proposals that cannot specify how success will be measured. The proposed evaluation framework distinguishes three categories of metrics: operational (training outcomes), research (R&I productivity and impact), and governance (alignment with mandate and democratic-oversight requirements).

The core operational metrics are as follows. Decision cycle time measures the elapsed time from event detection to coordinated institutional response, assessed before and after training cohorts; the expected direction of effect is compression. Decision quality rubrics assess the internal coherence, evidence traceability, option robustness, and risk reasoning of decisions made during scenario exercises, using structured scoring instruments derived from the VIS methodology. Cognitive resilience metrics track the recovery of decision quality after high-pressure events within scenarios: speed of recovery, completeness of recovery, and residual effects on subsequent decision quality. Human-AI calibration metrics assess the accuracy of participants' trust in AI recommendations

relative to the actual accuracy of those recommendations, a direct measure of automation bias and its inverse.

Adversarial robustness metrics represent the technically most demanding element of the framework. They assess participant performance under injected cognitive-attack conditions: adversarially manufactured consensus signals, data-poisoning cues designed to skew AI decision-support recommendations, false urgency triggers, and coordinated narrative insertion during exercises. Performance under these conditions is the most direct measure of the training's operational relevance.

Brčić's research on explainability in AI systems and on interdisciplinary directions in AI governance is directly relevant to the measurement framework (Eftimov et al., 2024). Explainable-AI governance, the requirement that AI recommendations in high-stakes environments be accompanied by interpretable justifications, provides both a technical requirement for the Decision Theatre's AI systems and a research domain in which the centre can generate internationally significant results. His work on data-mediation software architectures (Brčić et al., 2024) also connects to the verifiable data flows required by research-grade Decision Theatre instrumentation.

Table 1. Core EU-DUCCI evaluation metrics by domain

Metric domain	Specific indicator	Assessment method	Expected direction
Decision cycle time	Time to coordinated response	Pre/post cohort comparison	Compression
Decision quality	Coherence, evidence traceability, option robustness	Scoring against a VIS-derived rubric	Improvement
Cognitive resilience	Speed of recovery after pressure events	Measured through scenario AAR	Faster recovery
Human–AI calibration	Trust accuracy relative to AI recommendation quality	Calibration scoring	Reduced automation bias
Adversarial robustness	Performance under injected cognitive-attack conditions	Controlled scenario injection	Sustained performance
Research impact	Publications, doctrine adoption, dataset sharing	Standard bibliometrics + adoption tracking	Growth over time
Governance alignment	Ethics-board review pass rate, oversight audit outcomes	Independent review	Full alignment

Baseline establishment is a precondition of meaningful evaluation. The pilot exercise proposed in Section 7.4 must include full baseline measurement across all core operational metrics, enabling later training interventions to be assessed against a documented starting point. Longitudinal tracking of cohort performance at six, twelve, and twenty-four months after training will address the

key question of whether cognitive-resilience training produces durable capability change or only a temporary performance improvement.

9 Governance, Ethics, and Democratic Oversight

A dual-use centre operating in the cognitive domain must be designed from the outset to defend cognition, not to manipulate it. This is not merely a reputational requirement; it is an operational one. A centre perceived as an instrument of political manipulation would be institutionally destroyed before fulfilling its mission, and deservedly so. NATO's scientific framing of cognitive warfare explicitly highlights ethics, legal concerns, and interdisciplinary foundations as critical to capability planning in the cognitive domain (Du Cluzel, 2025). Three governance principles are proposed as non-negotiable design requirements.

The first principle is transparency of purpose. The centre's training objectives, scenario content, and research protocols must be publicly documented and subject to external review. The structural distinction between cognitive resilience, building institutional capacity to sustain decision quality under attack, and cognitive manipulation, the use of the same techniques to shape beliefs for political purposes, must be engineered into the centre's design rather than merely asserted as intention. Independent ethics boards with real oversight authority (not merely advisory functions) are a minimum requirement. Scenario content must be reviewed before use in training exercises by an interdisciplinary panel including legal experts, ethicists, and civil society representatives.

The second principle is human sovereignty over artificial intelligence. AI systems play significant roles both in training delivery and in decision support within the proposed centre. The risks of AI-mediated influence, both unintentional (through algorithmic bias in scenario design) and adversarial (through potential manipulation of the centre's own AI systems), must be explicitly managed. Clear autonomy boundaries, explainability requirements for all AI recommendations in Decision Theatre exercises, traceability of AI decision-support outputs, and regular adversarial audits of AI systems are mandatory components of the technical governance architecture. Brčić's (2025) framework for the design of human-AI collaborative decision-making provides a technical specification for these requirements.

The third principle is civil-military coordination with democratic safeguards. The centre's dual character creates genuine tension between the speed and operational-security requirements of defense constituencies and the transparency and accountability requirements of civilian governance. This tension cannot be resolved; it must be managed: through clear separation of the educational, research, and operational-intelligence roles; through the exclusion of live operational intelligence from training scenarios; and through separate governance tracks for civilian and defense programs that nonetheless share infrastructure and a knowledge base. The graduated risk framework of the EU AI Act and the structure of dual-use technology governance provide the regulatory context within which the centre's governance framework should be situated (European Commission, 2024a).

10 Strategic Urgency and Funding Pathways

10.1 The Implementation Window

The strategic urgency of the proposed centre rests on two converging developments that have created a narrow implementation window.

The first is the maturation of AI-enabled cognitive operations. The same large language models and multimodal generative systems that are transforming commercial software development are available to state and non-state actors for generating synthetic media, personalized influence operations, and adversarial data poisoning at industrial scale. The asymmetry between the cost of generating and the cost of countering AI-enabled cognitive operations is currently extreme: generation is cheap; detection and response are expensive and slow. Building institutional cognitive resilience now, before AI-enabled operations become ubiquitous at the level of routine institutional decision-making, is dramatically cheaper than retrofitting it after the damage has occurred.

The second is transatlantic strategic uncertainty. The 2025 US National Security Strategy has been read by multiple European and American analysts as signaling a sharper reprioritization of American commitments and a more transactional stance toward European allies, with significant implications for European strategic agency (Hamilton, 2025; European Policy Centre, 2025). Petener (2025) argues that asymmetric adversaries specifically target the seams between allied institutions; cognitive operations designed to widen transatlantic disagreements represent precisely this form of seam exploitation. Whether or not one accepts the sharpest interpretations of current American signaling, the planning lesson stands: European institutions must be capable of deciding and acting under contested narratives, externally seeded and internally amplified. An EU-DUCCI that requires five years to establish does not answer a threat that matures within twelve months; the proposed eighteen-month implementation path is a direct response to this urgency.

10.2 Funding Architecture

The funding architecture maps naturally onto existing EU instruments, and a diversified strategy across instruments reduces both political and budgetary risk. The European Defence Fund supports research and development with dual-use potential; the cognitive domain falls clearly within its scope. The security-research cluster of Horizon Europe covers resilience, critical infrastructure protection, and hybrid threats. Digital Europe supports the AI and data-infrastructure components of the Decision Theatre. IPCEI-AI, the Important Project of Common European Interest for artificial intelligence, is especially relevant given the participation of HashNET Technologies in the IPCEI-CIS consortium (Kapusta, 2025c), which documents an existing institutional relationship between the COTRUGLI partner network and the EU's principal investment vehicle for AI infrastructure. Erasmus+ supports the educational components, and industrial partnerships with firms in AI-based decision-support technology can provide equipment and co-funding for applied research.

A diversified funding strategy across these instruments also creates institutional resilience within the centre itself: multiple funding relationships distribute dependency and reduce the leverage that any single funder could exert over the research agenda and governance structure. This is not merely a financial consideration; it is a first-order governance consideration for an institution whose mission is the defense of epistemic independence.

11 Discussion

The argument of this article rests on three claims that merit critical examination. The first is that a dedicated new centre, rather than increased funding of existing institutions, is the appropriate response to the cognitive-resilience gap. The second is that Vanguard Leadership provides a doctrine adequate to the challenge of cognitive warfare. The third is that the dual-use design is operationally sustainable rather than merely rhetorically attractive.

On the first claim: the existing institutional landscape for cognitive resilience includes elements of genuine value, NATO's cognitive-warfare research program, the EU's StratCom function, and national resilience centres in Finland, Sweden, and the Baltic states. The argument for a new centre is not that these institutions are failing, but that none of them occupies the specific institutional position of the proposed EU-DUCCI: pan-European scope, civil-military integration, accredited training with standardized measurement, immersive scenario exercises with research-grade instrumentation, and a practice-research feedback loop that continuously improves doctrine and training methodology. Existing institutions are either nationally focused, defense-specific, or research-oriented to the exclusion of operational training functions. The gap is specific, not general.

On the second claim: VL is proposed as a doctrine, not as settled scientific consensus. It has a growing empirical base, publications in the ILJ (Kapusta, 2025a), ten years of data on COTRUGLI cohorts across 45 generations, and evidence of HAI5 application in commercial and institutional contexts, but it has not been tested in cognitive-warfare contexts at the EU institutional level. The appropriate academic posture is to treat VL as the best available framework for the design of institutional training while simultaneously building the research infrastructure, through the centre's R&I function, that would enable genuine empirical testing. This is consistent with how military doctrine has historically developed: operational experience generates frameworks; institutional research tests and refines them.

On the third claim: the dual-use design raises a legitimate concern about mission dilution. Institutions that serve multiple constituencies often serve none of them well. The counterargument rests on a specific structural claim: the cognitive vulnerabilities the centre addresses derive from universal properties of human cognition under pressure, not from domain-specific operational requirements. Scenarios differ; the target cognitive architecture does not. This is what makes cross-domain training genuinely valuable: civilian participants gain exposure to military decision discipline; defense participants gain exposure to the demands of democratic governance and considerations of public legitimacy. That interaction generates a capability that neither domain would develop in isolation.

Petener's (2025) analysis of the logic of asymmetric warfare reinforces the doctrinal bridge between contemporary strategic competition and the requirements of institutional leadership, and supports the claim that high-status European institutions must adapt faster than their inherited strategic cultures. Brčić's research on explainability and robust software architectures strengthens the technical-governance foundations required for trustworthy human-AI collaborative decision environments in the proposed centre (Brčić et al., 2024; Eftimov et al., 2024).

12 Limitations and Future Research Directions

This article advances a conceptual and institutional proposal; it does not provide empirical validation of the proposed centre's effectiveness. That is the primary limitation and must be stated without qualification.

The effectiveness of VIS and HAI5 in cognitive-warfare contexts has not been tested in controlled scenario exercises at the EU institutional level. The Decision Theatre architecture is conceptual; its technical specification, instrumentation, and scenario design require substantial additional development. The governance framework is principled but not yet operationalized; the institutional procedures and oversight mechanisms required for its genuine implementation are not specified here.

The kill indicators for the article's central argument, the empirical conditions that would materially weaken the case for EU-DUCCI, are as follows. First: empirical evidence that existing institutions (NATO's cognitive-warfare program, national resilience centres) already provide adequate training for decision superiority across civil-military boundaries at the EU level. Second: evidence that the dual-use design produces institutional dysfunction, conflict between the requirements of civilian and defense programs, governance failures, or mandate drift, that outweighs its benefits. Third: evidence that the four VL pillars do not transfer effectively to EU institutional contexts from the commercial and defense contexts in which they were primarily developed. Any of these findings would require substantial revision of the proposal.

Future research should proceed along three directions. First, empirical testing of VIS and HAI5 training protocols in Decision Theatre scenarios using the measurement instruments proposed in Section 8, with pre-registered protocols and independent evaluation. Second, comparative analysis of the Finnish resilience model, internationally recognized as best practice for civilian cognitive resilience, to identify transferable design elements and genuine innovations in the proposed centre. Third, technical development of the Decision Theatre AI-integration architecture, building on Brčić's work on explainability and the design of human-AI collaborative decision-making, to produce a validated instrumentation framework for measuring cognitive resilience.

13 Conclusions

Europe's strategic challenge is increasingly cognitive: adversaries exploit human vulnerabilities at scale, amplified by AI-enabled information environments and by structural properties of democratic deliberation that make manufactured uncertainty so effective as a weapon of paralysis. The EU has undertaken important policy commitments on resilience and countering hybrid threats, but a key institutional gap remains: no pan-European institution systematically trains, measures, and improves decision quality under cognitive attack across civil-military institutional boundaries.

This article has proposed the European Dual-Use Cognitive Intelligence Centre as a practical remedy. The centre combines accredited education, immersive training and exercises through a Decision Theatre, applied research and innovation, and a self-sustaining community of practice. Its doctrine is Vanguard Leadership, operationalized through VIS and HAI5 and grounded in a NEO Cotruglian philosophical tradition that connects contemporary institutional requirements to an observation 567 years old: an institution that cannot maintain clear epistemic accounts of what it knows, what it

assumes, and what it has decided will fail precisely at the moment when strategic performance is most needed.

The applied mixed methodology, VIS structured analysis, comparative institutional analysis, and Cotruglian historical analysis, provides three distinct evidentiary foundations for the central claim, while the explicit identification of kill indicators in the limitations section ensures that the proposal is falsifiable rather than merely aspirational. The governance framework is grounded in three non-negotiable principles: transparency of purpose, human sovereignty over artificial intelligence, and civil-military coordination with democratic safeguards.

The Vanguard leader this centre is designed to develop is not a heroic archetype. It is a capability profile: a leader who can maintain epistemic discipline under pressure, coordinate distributed action across uncertain information environments, manage AI augmentation without surrendering human judgment, and restore institutional decision quality after cognitive-attack events. These are trainable capabilities. In the NEO era, networked, exponential, orchestrated, cognitive intelligence is not optional. It is the enabling condition of democratic governance, of EU strategic autonomy, and of credible defense in a world where the decisive contest is increasingly the contest over shared reality.

Author Notes

Dražen Kapusta, DBA, is Principal and Founder of COTRUGLI Business School and of HashNET Technologies, and the architect of the Vanguard Leadership framework. He serves as an advisor to UNIDO and to EU bodies on artificial intelligence and blockchain strategy, and participates in the advisory structure of the 8ra IPCEI-CIS consortium. The World Business Museum, cited as a primary-source archive in this article, is located at COTRUGLI Business School in Zagreb. Correspondence: drazen.kapusta@cotrugli.eu

Brigadier Damir Stručić, PhD, is Deputy Commandant of the Croatian Defence Academy “Dr. Franjo Tuđman” in Zagreb and lectures at the University of Defence and Security “Dr. Franjo Tuđman”. His professional and academic work centers on national security, defense, and intelligence and counterintelligence. During the Croatian Homeland War he served in the intelligence system of the Republic of Croatia, gaining direct operational experience in a range of intelligence and counterintelligence missions; after the war he held a series of senior leadership positions within Croatia's national security system, participating in drafting the legislation governing the security and intelligence services, the National Security Strategy of the Republic of Croatia, and numerous other strategic, planning, and implementation documents in the fields of security and intelligence and counterintelligence. His teaching and research focus on national security, intelligence and counterintelligence, strategic management of security systems, military education, and contemporary security challenges.

The authors declare that there are no conflicts of interest beyond the institutional affiliations noted above.

References

Boyd, J. R. (1986). *Patterns of conflict* [Unpublished briefing]. United States Air Force.

- Brčić, M. (2025). Leading with AI: How to blend human judgment with machine intelligence for superior decision-making. *International Leadership Journal*, 17(1), 34–49.
- Brčić, M., Bošnić, I., & Vrdoljak, B. (2024). Mask–Mediator–Wrapper architecture as a data mesh driver. *IEEE Transactions on Software Engineering*. <https://doi.org/10.1109/TSE.2024>.
- Cialdini, R. B. (1984). *Influence: The psychology of persuasion*. Harper Collins.
- Cotrugli, B. (2017). *The book of the art of trade* (J. Kirshner, Ed.; T. Smail, Trans.). Palgrave Macmillan. (Original work written 1458, published 1573)
- COTRUGLI CoLab. (2025a). *VIS Framework and AI simulation* [Internal resource]. COTRUGLI Business School.
- COTRUGLI CoLab. (2025b). *Vanguard Leadership* [Internal resource]. COTRUGLI Business School.
- Council of the European Union. (2022). *A strategic compass for security and defence: For a European Union that protects its citizens, values and interests and contributes to international peace and security*. Publications Office of the European Union. <https://www.consilium.europa.eu/media/54412/20220325-strategic-compass-en.pdf>
- Du Cluzel, F. (2025). *Cognitive warfare: The advent of the concept of 'cognitics' in the field of conflict*. NATO Science and Technology Organization Research Report. <https://www.sto.nato.int>
- Eftimov, T., Kop, M., Brčić, M., et al. (2024). Explainable artificial intelligence (XAI) 2.0: A manifesto of open challenges and interdisciplinary research directions. *Information Fusion*, 106, 102301. <https://doi.org/10.1016/j.inffus.2024.102301>
- European Commission. (2024a). *Dual-use technologies: Enhancing synergies between civilian and defence research and innovation*. Publications Office of the European Union. <https://eur-lex.europa.eu>
- European Commission. (2024b). *White paper on options for enhancing support for research and development involving technologies with dual-use potential*. Publications Office of the European Union. <https://eur-lex.europa.eu>
- European Policy Centre. (2025, December 5). *Trump's new national security strategy: An existential threat to Europe?* <https://www.epc.eu>
- Hamilton, D. S. (2025, January). *Breaking down Trump's 2025 national security strategy*. Brookings Institution. <https://www.brookings.edu>
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kapusta, D. (2025a). Vanguard leadership: Integrating dynamic warfare strategies and innovation tactics in an AI-driven world. *International Leadership Journal*, 17(1), 4–28.
- Kapusta, D. (2025b). *Vanguard intelligence summary: A practitioner's framework for executive decision-making under uncertainty* [Unpublished working paper]. COTRUGLI Business School.
- Kapusta, D. (2025c). *NEO Cotruglian triple entry (NCTE): Trust infrastructure for the machine economy* [White paper]. HashNET Technologies. <https://hashnet.tech>
- Kapusta, D. (2026a). *Elaborat – NOVI: Framework for SOIS–COTRUGLI cooperation on AI-augmented leadership and cognitive intelligence* [Unpublished project elaboration]. COTRUGLI Business School.
- Kapusta, D. (2026b). *Mission command in the age of algorithmic warfare: From sequential operations to NEO leadership in multidomain environments*. COTRUGLI CoLab internal archive.
- Kapusta, D. (2026c). *NSS 2025 intelligence assessment: Implications for European sovereignty and transatlantic relations* [Unclassified internal assessment]. COTRUGLI Business School.

- Kapusta, D. (2026d). *A strategic report on countering state-sponsored cognitive warfare in Europe* [Unpublished report]. COTRUGLI Business School.
- Kapusta, D., & Stručić, D. (2026). *Vanguard (VIS): Operational manual for disciplined decision-making* [Unpublished manuscript]. COTRUGLI Business School.
- Le Bon, G. (2002). *The crowd: A study of the popular mind*. Dover Publications. (Original work published 1895)
- Lewandowsky, S., Ecker, U. K. H., & Cook, J. (2017). Beyond misinformation: Understanding and coping with the 'post-truth' era. *Journal of Applied Research in Memory and Cognition*, 6(4), 353–369. <https://doi.org/10.1016/j.jarmac.2017.07.008>
- Petener, Z. (2025). Business as warfare: A tactical playbook for modern leadership. *International Leadership Journal*, 17(1), 29–33.
- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533. [https://doi.org/10.1002/\(SICI\)1097-0266\(199708\)18:7<509::AID-SMJ882>3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1097-0266(199708)18:7<509::AID-SMJ882>3.0.CO;2-Z)
- The White House. (2025). *2025 national security strategy*. <https://www.whitehouse.gov>
- U.S. Naval Institute Staff. (2025, December 5). 2025 U.S. national security strategy. *USNI News*. <https://news.usni.org>
- Wardle, C., & Derakhshan, H. (2017). *Information disorder: Toward an interdisciplinary framework for research and policy making*. Council of Europe. <https://edoc.coe.int>