

Need for an EU Dual-Use Centre for Cognitive Intelligence: Why the EU Needs Vanguard Leaders

Damir Stručić

Croatian Defence Academy "Dr. Franjo Tuđman"

Dražen Kapusta

COTRUGLI Business School

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

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 governance frameworks—but a structural gap persists: no pan-European institution systematically trains, measures, or improves decision superiority under cognitive attack across civil–military boundaries. This article proposes the EU Dual-Use Centre for Cognitive Intelligence (EU-DUCCI) to close that gap. Drawing on NATO scientific findings on cognitive warfare, the Cotruglian merchant ethics tradition (1458), the Vanguard Leadership Framework, and the HAI5 and VIS operational instruments, the article argues that cognitive resilience is a trainable and measurable institutional capability—not a cultural disposition. The article offers a conceptual framework grounded in a mixed methodology combining Vanguard Intelligence Summary analysis, comparative institutional analysis, and Cotruglian historical analysis, outlines the Decision Theatre architecture, proposes evaluation

metrics and governance requirements, and maps a pragmatic eighteen-month implementation pathway.

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 systems. In contemporary strategic competition, adversaries increasingly seek advantage not through kinetic force alone 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 believe is achievable, and how quickly they can coordinate collective response. When cognitive operations succeed, they do not merely confuse individual decision-makers; they fragment the shared epistemic foundation that institutional coordination requires.

Recent NATO scientific work formalizes this challenge, arguing that cognitive warfare targets cognition itself—seeking to influence or disrupt human decision-making and behaviour via military and non-military means, amplified by advanced technologies and hybrid influence operations (Du Cluzel, 2025). The dual-use character of this threat is embedded in that definition: cognitive attacks cut across civilian and military populations simultaneously, and therefore require whole-of-government responses spanning research, policy, training, and resilience planning.

The EU recognizes the significance of this challenge in policy. 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 dual-use technology governance have been updated to manage civil-military research synergies (European Commission, 2024a, 2024b). These are genuine achievements.

Yet a specific institutional capability gap persists that existing instruments do not address. Europe is generating analysis, policy, and legal architecture, but it lacks a pan-European, dual-

use platform for repeatedly training decision superiority at scale—under realistic cognitive threat conditions, with measurable performance outcomes and research-grade after-action learning. 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 defence require.

This article advances a specific proposal: the EU Dual-Use Centre for Cognitive Intelligence (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 human–AI integration framework. The intellectual lineage grounding the proposal extends to Benedetto Cotrugli's 1458 merchant philosophy: the institution that cannot sustain clear epistemic accounts of what it knows, what it assumes, and what it has decided will fail at precisely the moment strategic performance is most required.

The article is structured as follows. Section 2 examines the mechanism of cognitive warfare as an industrialized social proof weapon. Section 3 surveys the EU policy baseline and characterizes the remaining institutional gap. Section 4 argues that mission command logic now extends into the civil domain. Section 5 presents Vanguard Leadership as a dual-use doctrine for cognitive conflict. Section 6 describes the mixed methodology. Section 7 presents the EU-DUCCI proposal, including the Decision Theatre architecture. Section 8 proposes the evaluation framework. Section 9 addresses governance and democratic oversight. Section 10 examines strategic urgency and funding pathways. Sections 11 and 12 provide 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, it is necessary to ground the analysis 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 contexts reduce individual critical thinking, heighten impulsivity, and increase susceptibility to emotional contagion. Under conditions of ambiguity or threat, individuals tend to look to others for epistemic cues—a behaviour formalized as social proof (Cialdini, 1984). This is not a defect of weak minds; it is an adaptive cognitive heuristic operating under information scarcity and time pressure—the precise conditions that adversarial cognitive operations are designed to manufacture.

Contemporary research on misinformation deepens this analysis. Lewandowsky et al. (2017) demonstrate that corrective information frequently 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 expose their artificiality. The problem is not only that false beliefs circulate; it is that repeated exposure to manufactured uncertainty progressively degrades 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 this article's Cotruglian methodological strand, examined in Section 6.3—is the claque system of eighteenth and nineteenth century European theatre. Professional claques were organized groups of paid applauders systematically deployed to manufacture emotional responses in audiences and, critically, to shape other audience members' interpretations of quality. The mechanism was not deception in the conventional 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 behaviour on digital platforms: bot networks, like storms, coordinated

amplification campaigns, and algorithmically distributed narrative pushing. The scale is different; the cognitive mechanism is identical. Modern cognitive operations industrialize social proof at network speed, targeting not just individual minds but the epistemic infrastructure of institutions. What Cotrugli (1458/2017) would have recognized as the corruption of the merchant's reliable market signal—the 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 introduce false beliefs into circulation; they fragment the shared epistemic foundation that collective decision-making requires. When publics and leaders operate from incompatible factual frameworks, consensus becomes unavailable and institutional delay becomes structurally embedded in the response process. This is precisely the outcome adversaries seek: not to win an argument, but to make the argument unresolvable.

Research on Russia's approach to influence operations in Europe documents a consistent pattern of platformized, deniable operations designed to normalize interference below thresholds that trigger decisive collective response (Wardle & Derakhshan, 2017). The objective is not to convince European publics of any particular position, but to degrade confidence in all positions—manufacturing sufficient ambient uncertainty that coordinated action becomes politically unachievable. 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 processes are structurally resistant to cognitive degradation.

3. THE EU POLICY BASELINE AND THE REMAINING INSTITUTIONAL GAP

The EU has not been passive in responding to cognitive dimensions of hybrid threats. A sustained policy effort has produced a substantial institutional and legal architecture whose achievements and limitations must both be characterized accurately.

The Strategic Compass for Security and Defence (Council of the European Union, 2022) represents 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, updated dual-use technology governance frameworks reflect 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 across civilian and military domains. The Digital Services Act creates platform accountability requirements that reduce structural advantages currently enjoyed by coordinated inauthentic behaviour campaigns. Horizon Europe and the European Defence Fund both include provisions relevant to cognitive security research.

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

This gap is exacerbated by the human–AI dimension. As AI systems are increasingly embedded in intelligence fusion, crisis management, and strategic analysis, new cognitive vulnerabilities emerge alongside new capabilities. Automation bias—the documented tendency of human operators to defer to algorithmic recommendations even when those recommendations are flawed or adversarially manipulated—represents a cognitive vulnerability that did not exist at institutional scale 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 training and exercise gap, a measurement gap, and a human–AI governance gap.

4. MISSION COMMAND LOGIC EXTENDED TO THE CIVIL DOMAIN

The doctrinal logic of mission command was developed to address a persistent problem in military operations: how to maintain effective coordinated action under high uncertainty, degraded communication, and rapid tempo change. Hierarchical approval processes fail under these conditions because they are too slow, too dependent on accurate upward information flows, and too brittle when commanders are unavailable or overwhelmed. Mission command responds with disciplined autonomy: subordinates are empowered to act within understood intent without requiring approval for each decision, because they have been trained to understand purpose and exercise judgment about means while remaining aligned on ends.

In the age 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 leader development for human–AI teaming, advanced systems become vulnerabilities rather than capabilities. The organizations that will act most effectively in this environment are those that have already developed the cultural infrastructure—shared mental models, clear intent frameworks, mutual trust—enabling distributed action under uncertainty.

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

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

This adaptation requires careful handling. Mission command in a democratic governance context cannot replicate the command authority relationships of military organizations. Democratic oversight, rule of law, and civil liberties impose constraints that are constitutive requirements of legitimate governance. Adapting mission command logic to civil 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 as a 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 AI as both force multiplier and 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—detecting signals before they become crises, acting within narrowing windows, and restructuring capabilities accordingly. Second, Boyd's (1986) OODA (Observe, Orient, Decide, Act) cycle

provides the operational architecture for decision-making under uncertainty and tempo pressure. Organizations that cycle through OODA faster than adversaries gain decisive advantage regardless of relative resource endowment; in cognitive warfare, the critical resource is not material but the quality and speed of sensemaking. Third, and distinctively, the NEO Cotruglian philosophical tradition provides the ethical and motivational foundation—a foundation examined more fully in the methodology section.

The VL framework operationalizes these foundations through four pillars. Resilience refers not to psychological toughness alone but to the organizational capacity to recover decision quality after cognitive pressure events—a trainable, measurable outcome. Critical Strategic Thinking encompasses disciplined hypothesis testing, structured reasoning, and rapid assumption kill cycles designed to maintain epistemic integrity faster than adversaries can manufacture uncertainty. AI Augmentation addresses deploying artificial intelligence as a genuine force multiplier while managing automation bias, explainability requirements, and adversarial manipulation risks. Tribal Collaboration describes cultivating 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-Use Character of Vanguard Leadership

The proposition 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, epistemic fragmentation under information overload—manifest in boardrooms, government ministries, and operational headquarters simultaneously. The same doctrinal requirements—intent clarity, distributed authority, disciplined hypothesis testing—apply across these contexts. The difference is one of scale, stakes, and the specific character of adversarial operations, not of fundamental cognitive architecture.

For the proposed centre, this dual-use character has a practical institutional implication: the centre can serve both civil and defence stakeholders without requiring separate doctrines or training methodologies. The core training content is the same; scenarios are calibrated to context. This improves the economics of the centre and, more importantly, enables cross-fertilization

between civil and military decision-making practice that generates genuine interoperability—shared cognitive culture rather than merely 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 an explicit discipline on decision processes: hypotheses and options must be articulated separately from evidence; evidence items must be assigned to specific hypotheses; assumptions must be tracked explicitly with kill indicators that would falsify them; and decision briefs must distinguish between what is known, what is assessed, and what is assumed. This structure does not slow decision-making; in trained practitioners, it accelerates decisions by eliminating cognitive overhead from unstructured deliberation and reducing the surface area available for adversarial narrative insertion.

HAI5 is a staged methodology for human–AI integration framed as an infrastructure of trust (Kapusta & Stručić, 2026). The HAI5 architecture connects three domains: technology (AI agents, personalization systems, verifiability mechanisms), governance (traceability requirements, accountability structures, compliance frameworks), and people (the academy, Decision Theatre, and community of practice). Together, VIS and HAI5 address the central academic critique of practitioner frameworks: the lack of operational specificity enabling replication, measurement, and scholarly evaluation. VL is not a leadership philosophy alone; 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 interlocking analytical approaches: Vanguard Intelligence Summary (VIS) structured 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 shared claim that no single approach could sustain alone. This section describes each approach, its application in the article, and the epistemic discipline imposed on the overall argument.

6.1 Vanguard Intelligence Summary (VIS) Structured Analysis

VIS, developed as a practitioner methodology adapted from military intelligence analytical tradecraft (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 where empirical support is unavailable or contested; and kill indicators are identified that would substantially weaken the article's central argument.

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

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

6.2 Comparative Institutional Analysis

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

The second is EU dual-use technology governance, examined through Commission White Papers and regulatory documents (European Commission, 2024a, 2024b). This case provides the civil governance context within which the proposed centre must operate—including the funding landscape, the ethics and oversight requirements, and the institutional culture of EU 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 delivery, HAI5 implementation, and CO-lab innovation platform development (Kapusta, 2025a). This case provides practitioner evidence about what structured leadership training for NEO-era conditions can and cannot accomplish—including documented performance improvements in structured decision-making across 45 cohorts spanning pharmaceutical, technology, energy, and financial services sectors. This case is not presented as proof of effectiveness in cognitive warfare contexts specifically, which would represent 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 the limits of comparability are noted where the cases differ in ways that constrain inference.

6.3 Cotruglian Historical Analysis

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

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 being targeted by contemporary influence operations—susceptibility to manufactured reputation, degradation of epistemic discipline, confusion of manufactured consensus with market signal—are not novel pathologies of the digital age. They are recurrent features of any information environment in which the cost of generating false signals is lower than the cost of verifying them. Cotrugli's merchant was operating in exactly such an environment, and the ethical and practical disciplines he codified—verified sources, explicit separation of known from assumed, trust as commercial infrastructure rather than moral luxury—map directly onto the VIS analytical framework and the VL four pillars.

The claque analogy introduced in Section 2.2 is an instance of Cotruglian historical method: using a historical operating model to make visible a causal mechanism—the engineering of social proof—whose contemporary digital manifestations are more difficult to perceive precisely because of their scale and speed. The World Business Museum in Zagreb, which curates the largest institutional archive of pre-industrial business history globally (covering commercial practice from 4th-millennium Mesopotamia through the 18th century Industrial Revolution, with Cotrugli's legacy as its scholarly centrepiece), provides the primary source foundation for this

analytical approach. No comparable archive exists, making the Cotruglian historical strand a methodological contribution that cannot be replicated by other research groups.

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 the VL doctrine and a longer empirical baseline than any contemporary study can supply.

6.4 Methodological Limitations

This mixed methodology is appropriate to the article's purpose—advancing a conceptual and institutional proposal before empirical testing is possible at the required scale—but it carries limitations that must be acknowledged. The article does not provide controlled empirical testing of VL effectiveness in cognitive warfare contexts; such testing requires scenario exercises with measurement instruments not yet available at EU institutional scale. It does not provide cost-benefit analysis of the proposed centre against alternative responses; the requisite data does not exist in the public domain. And the historical analogies, while illuminating causal mechanisms, cannot substitute for prospective evaluation of training effectiveness. These limitations are genuine; the kill indicators in Section 12 specify the empirical conditions that would substantially weaken the article's central argument.

7. THE EU DUAL-USE CENTRE FOR COGNITIVE INTELLIGENCE: PROPOSAL

7.1 Core Concept and Rationale

The EU Dual-Use Centre for Cognitive Intelligence (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 EU institutional capacity for decision superiority under cognitive attack.

The centre is dual-use in two senses. First, it serves both civil and defence stakeholders: government ministries, military commands, critical infrastructure operators, judicial institutions, regulatory bodies, and business school faculties simultaneously. Second, it combines education and research functions in a configuration designed for the generation and application of knowledge, not merely its dissemination. The institutional rationale is grounded in the capability gap analysis of Section 3: existing EU instruments generate analysis and policy but do not systematically generate trained cognitive resilience at the level of institutional decision-making.

7.2 Mission and Functions

The centre is designed as a three-function institution. The first function is Cognitive Education, Training and Exercises (Cog-ET&E): structured, accredited programmes and immersive scenario exercises conducted in the Decision Theatre environment. Target populations span four groups: senior civil servants and EU institutional officials; military commanders and intelligence professionals; critical infrastructure executives; and faculty at EU business and public administration schools who will subsequently train the next tier of leaders. Training is designed to develop three capabilities: decision superiority under cognitive pressure, human-AI calibration, and cognitive resilience recovery speed.

The second function is Research and Innovation (R&I): applied research into cognitive performance under adversarial conditions, AI-assisted analytics for decision support, training methodology development, and metrics of decision superiority and resilience. Research outputs include peer-reviewed publications, measurement instruments, training scenario repositories, and doctrine notes. The R&I function serves as the validation engine for the training function: Decision Theatre exercises generate data on decision quality under pressure; that data feeds research; research findings refine training methodology. This virtuous cycle distinguishes a serious capability development institution from an educational programme.

The third function is Community of Practice: a pan-European network of instructors, analysts, and leaders who continuously evolve standards, share scenario experiences, and maintain the institutional knowledge base. The community of practice serves an additional strategic function: it creates a European epistemic infrastructure—a network of individuals with shared analytical frameworks, mutual trust, and common standards for evidence evaluation—that is itself a form of cognitive resilience at civilizational scale.

7.3 The 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-based exercises under realistic cognitive threat conditions, 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 simulating the information conditions of actual cognitive attack events—high volume, contested reliability, adversarially curated, algorithmically shaped. Time compression: exercises are conducted under temporal pressure calibrated to shift participants from System 2 to System 1 processing, because it is in this shift that cognitive vulnerabilities most clearly manifest. AI integration: AI decision support tools are present in 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 review and research analysis.

Scenario design draws on two sources: the existing repository of documented cognitive operations against EU member states (adapted and anonymized), and forward-looking scenarios developed by the R&I function to anticipate adversarial modalities not yet operationally deployed but technically feasible given current AI capabilities.

After-action review (AAR) is treated as a core learning intervention rather than an administrative procedure. VIS discipline is applied to AAR: assumptions made during the exercise are evaluated against pre-stated kill indicators; decision quality is assessed against explicit rubrics; human–AI

calibration is measured; and cognitive resilience recovery is tracked from peak pressure events to restored performance. AAR data feeds both individual development plans and institutional research.

7.4 Institutional Location and Implementation Pathway

The centre's institutional anchor should satisfy three criteria: geographical 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 candidate anchor, given its defence-sector research mandate and Croatia's position as both EU member state and NATO ally. Partnership with COTRUGLI Business School—which brings ten years of VL training implementation, the World Business Museum archive, and the CO-lab innovation platform—would provide the civilian education and commercial sector dimensions that a purely military-academic anchor would lack.

A pragmatic three-step implementation pathway is proposed to avoid the declarative partnership failures that have characterized previous EU capability 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 application, using pilot results to support applications to European Defence Fund, Horizon Europe, IPCEI-AI, and Digital Europe instruments. The target timeline from institutional commitment to first operational cohort is eighteen months.

8. METRICS AND EVALUATION FRAMEWORK

Academic evaluators and policy sponsors rightly challenge 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 (compliance with mandate and democratic oversight requirements).

Core operational metrics are as follows. Decision cycle time measures 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 VIS methodology. Cognitive resilience metrics track recovery of decision quality following high-pressure events in scenarios: speed of recovery, completeness of recovery, and residual effects on subsequent decision quality. Human–AI calibration metrics assess the accuracy of participants' confidence in AI recommendations relative to those recommendations' actual accuracy—a direct measure of automation bias and its inverse.

Adversarial robustness metrics represent the most technically demanding element of the framework. These assess participant performance under injected cognitive attack conditions: adversarially manufactured consensus signals, data poisoning cues designed to distort 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 interdisciplinary AI governance directions 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 rationale—provides both a technical requirement for the Decision Theatre's AI systems and a research domain in which the centre can generate internationally significant outputs. His software architecture work on data mediation (Brčić et al., 2024) also connects to the auditable data flows that research-grade Decision Theatre instrumentation requires.

Table 1. *EU-DUCCI Core 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	VIS-derived rubric scoring	Improvement
Cognitive Resilience	Recovery speed after pressure events	Scenario AAR measurement	Faster recovery
Human–AI Calibration	Confidence 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 bibliometric + uptake tracking	Growth over time
Governance Compliance	Ethics board review pass rate, oversight audit outcomes	Independent review	Full compliance

Baseline establishment is a prerequisite for meaningful evaluation. The pilot exercise proposed in Section 7.4 must include full baseline measurement for all core operational metrics, enabling the subsequent training intervention to be assessed against a documented starting point. Longitudinal tracking of cohort performance at six, twelve, and twenty-four months post-training will address the critical question of whether cognitive resilience training produces durable capability change or merely temporary performance improvement.

9. GOVERNANCE, ETHICS, AND DEMOCRATIC OVERSIGHT

A dual-use centre operating in the cognitive domain must be designed from inception to defend cognition rather than manipulate it. This is not merely a reputational requirement; it is an operational one. A centre perceived as a tool for political manipulation would be institutionally destroyed before fulfilling its mission—and would deserve to be. NATO's scientific framing of cognitive warfare explicitly flags ethics, legal concerns, and interdisciplinary foundations as essential to cognitive-domain capability planning (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—using the same techniques to influence beliefs for political purposes—must be embedded in the centre's design, not merely asserted as intent. Independent ethics boards with genuine oversight authority (not merely advisory functions) are a minimum requirement. Scenario content must be reviewed by an interdisciplinary panel including legal scholars, ethicists, and civil society representatives before use in training exercises.

The second principle is human sovereignty over AI. AI systems play substantial roles in both training delivery and decision support in the proposed centre. The risks of AI-mediated influence—both inadvertent (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 audit of AI systems are required components of the technical governance architecture. Brčić's (2025) framework for human–AI decision design provides the technical specification for these requirements.

The third principle is civil–military coordination with democratic safeguards. The dual-use character of the centre creates genuine tension between the speed and operational security requirements of defence participants and the transparency and accountability requirements of civil governance. This tension cannot be dissolved; it must be managed through clear role separation between education, research, and operational intelligence functions; through exclusion of real operational intelligence from training scenarios; and through distinct governance structures for civil and defence programmes that nevertheless share infrastructure and knowledge base. The EU's AI Act tiered risk framework and dual-use technology governance structure 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 generation systems 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 pervasive at the level of routine institutional decision-making—is dramatically cheaper than retrofitting it after damage is done.

The second is transatlantic strategic uncertainty. The 2025 U.S. National Security Strategy has been interpreted by multiple European and U.S. analysts as signalling a sharper reorientation of U.S. priorities and a more transactional posture toward European allies, with significant implications for European strategic agency (Hamilton, 2025; European Policy Centre, 2025). Peternel (2016) argues that asymmetric adversaries specifically target seams between allied institutions; cognitive operations designed to widen transatlantic disagreements represent precisely this form of seam exploitation. Whether or not one accepts the strongest interpretations of current U.S. signalling, the planning lesson holds: European institutions must be able to decide and act under contested narratives, externally driven and internally amplified. An EU-DUCCI that takes five years to establish does not address a threat maturing on a twelve-month timeline; the proposed eighteen-month implementation pathway 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 R&D with dual-use potential; the cognitive domain falls clearly within its scope. Horizon Europe's security research cluster 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 on Artificial Intelligence—is particularly relevant given HashNET Technologies' participation in the IPCEI-CIS consortium (Kapusta, 2025c), which documents an existing institutional relationship between COTRUGLI's partner network and the EU's major AI infrastructure investment vehicle. Erasmus+ supports the educational components, and industrial partnerships with AI decision support technology companies can provide both equipment and applied research co-funding.

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

11. DISCUSSION

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

On the first claim: the existing institutional landscape for cognitive resilience includes elements of genuine value—NATO's cognitive warfare research programme, the EU's StratCom function, 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 feedback loop between practice and research that continuously improves doctrine and training methodology. Existing institutions are either nationally focused, defence-specific, or research-oriented to the exclusion of the operational training function. The gap is specific, not general.

On the second claim: VL is proposed as doctrine, not as established scientific consensus. It has a growing empirical base—the ILJ publication (Kapusta, 2025a), ten years of COTRUGLI cohort data across 45 cohorts, HAI5 implementation evidence across commercial and institutional contexts—but has not been tested in cognitive warfare contexts at EU institutional scale. The appropriate academic posture is to treat VL as the best available framework for institutional training design while 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 been developed: operational experience generates frameworks; institutional research tests and refines them.

On the third claim: dual-use design raises a legitimate concern about mission dilution. Institutions serving multiple constituencies frequently serve none well. The counter-argument 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. The scenarios differ; the cognitive architecture targeted does not. This makes cross-domain training genuinely value-generating: civil participants gain exposure to military decision discipline; defence participants gain exposure to democratic governance requirements and public legitimacy considerations. The interaction generates capability that neither domain would develop in isolation.

Peternel's (2016) analysis of asymmetric warfare logics reinforces the doctrinal bridge between modern strategic competition and institutional leadership requirements, and supports the claim that high-status European institutions must adapt faster than legacy strategic cultures. Brčić's research on explainability and robust software architectures strengthens the technical governance

foundations required for trustworthy human–AI 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. This 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 EU institutional scale. 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 to implement it genuinely are not specified here.

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

Future research should pursue 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 civil cognitive resilience—to identify transferable design elements and genuine innovations in the proposed

centre. Third, technical development of the Decision Theatre's AI integration architecture, building on Brčić's work on explainability and human–AI decision design to produce a validated instrumentation framework for cognitive resilience measurement.

13. CONCLUSIONS

Europe's strategic challenge is increasingly cognitive: adversaries exploit human vulnerabilities at scale, amplified by AI-enabled information environments and by the structural properties of democratic deliberation that make manufactured uncertainty so effective as a paralysis weapon. The EU has made important policy commitments to resilience and to countering hybrid threats, but a critical 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 EU Dual-Use Centre for Cognitive Intelligence as a practical solution. 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 the NEO Cotruglian philosophical tradition that links contemporary institutional requirements to a 567-year-old observation: the institution that cannot maintain clear epistemic accounts of what it knows, what it assumes, and what it has decided will fail at exactly the moment strategic performance is most required.

The mixed methodology employed—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 AI, 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 sustain epistemic discipline under pressure, coordinate distributed action across uncertain information environments, govern AI augmentation without surrendering human judgment, and recover 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 for democratic governance, EU strategic autonomy, and credible defence in a world where the decisive contest is increasingly the contest for shared reality.

Author Note

Dražen Kapusta, DBA, is Principal and Founder of COTRUGLI Business School, and HashNET Technologies, and the architect of the Vanguard Leadership Framework. He serves as an advisor to UNIDO and EU bodies on AI and blockchain strategy and is a participant in the 8ra IPCEI-CIS consortium advisory structure. The World Business Museum, referenced as a primary source archive in this article, is housed at COTRUGLI Business School in Zagreb. The author declares no conflicts of interest beyond the institutional affiliations noted above. Correspondence: drazen.kapusta@cotrugli.eu

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)
- 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 [Unpublished conference paper]. COTRUGLI Business School.
- 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>