

From Simulation Outputs to Explainable C4I Insights

Neuro-Symbolic Approaches for MDO Planning and Critical-Sector Resilience



Michael Romei de Socio — PhD Candidate

Alessio Merlo — Full Professor of Computer Engineering

CASD — School of Advanced Defense Studies



NATO CAX FORUM 2026 | Track T3

Rome · 29 September–01 October 2026



We do not lack outputs. We lack operational meaning.


Simulation evidence becomes useful only when its meaning, timing and assumptions remain visible.



The analytical bottleneck is translation — not data volume.

MDO is a dependency problem

The operational meaning often lies in relationships between events—not in any isolated event.




Illustrative pathway

Dependencies expose possible propagation paths; they do not establish deterministic causality.



A bridge, not a replacement

Preserve the full chain from simulated dynamics to human verification.




The value is continuity: evidence, semantics, assumptions and constraints stay connected.



Why neural + symbolic + human?

The value comes from composition—each part resolves a different analytical limitation.




No single component supplies pattern discovery, explicit admissibility and accountable judgement by itself.



ENRICH — Explore alternative futures under constraints

Explore more futures without assuming that every generated future is valid.


PARTIAL SCENARIO HISTORY





WARN — See cascades before they converge

Individual signals may appear weak; their temporal and structural convergence may not.




Reason over distributed evidence—not only isolated alerts.



COMPARE — Make COA consequences inspectable

Do not ask only which option looks best. Ask what supports each assessment.



INSPECTION QUESTIONS

- What assumptions support it?
- Which capabilities become critical?
- Which constraints activate?
- Where can second-order effects emerge?
- What evidence supports the assessment?

Transparent comparison—not automated selection of a ‘best’ COA.



Incremental integration with the CAX ecosystem

Add an analysis layer without rebuilding the simulation, C4I or command stack.



POTENTIAL EXCHANGE PATHS



Federation path



Semantic exchange path




Service-delivery path

CONCEPTUAL COMPATIBILITY ≠ IMPLEMENTED INTEROPERABILITY ≠ OPERATIONAL DEPLOYMENT



Wargaming and CAX as a validation environment


Use WG & CAX experiments to test analytical value, verification burden and integration assumptions.





Explainable means inspectable

An explanation is useful only if evidence and assumptions remain recoverable.



A viable path to verification—not merely a final approval button.



From research components to collaborative capability

WHAT WE HAVE STUDIED

Ontology-centred architectures

Ontology-guided information structuring

Temporal structural reasoning

Learned scenario continuation + validation

Digital human-in-the-loop wargaming

WHAT WE STILL NEED TO TEST TOGETHER

Integrated CAX workflows + interoperability

Larger heterogeneous simulation environments

Degraded and uncertain semantic conditions

Operator interaction + human factors

Operationally meaningful validation criteria

**M&S +
INTEROPERABILITY**

**AI + KNOWLEDGE
ENGINEERING**

**HUMAN-CENTRED
C4I**

From simulated events to trusted operational evidence.

Questions

From Simulation Outputs to
Explainable C4I Insights

Michael Romei de Socio — PhD Candidate
Alessio Merlo — Full Professor of Computer Engineering

CASD — School of Advanced Defense Studies
Piazza della Rovere 83 · 00165 Rome · Italy

michael.romeidesocio@unicasd.it
alessio.merlo@unicasd.it




GOOGLE SCHOLAR



ORCID

Research foundations

Five connected studies support the workflow:



[1] Romei de Socio, Pozzato & Merlo — A Scalable Ontology-Driven Architecture for Situational Awareness and Decision Support in Multi-Domain Operations.

[2] Romei de Socio, Pozzato & Merlo — When Ontologies Guide LLMs: A Controlled Study of Zero-Shot Schema-Aligned Information Extraction.

[3] Romei de Socio, Pozzato & Merlo — TRACTA: Benchmarking Temporal Reasoning over Semantic Trajectories.


[4] Romei de Socio, Pozzato & Merlo — Learning and Structurally Validating Simulation Scenario Continuations in Dynamic Graph Systems.

[5] Romei de Socio, Bernardi & Merlo — From Tabletop to Web Application: A Lightweight and Modular Architecture for Politico–military Digital Wargaming.



Selected experimental evidence

TRACTA — semantic-input recurrent configuration vs strongest raw-event reference [3]



SCENARIO CONTINUATION [4]

Medium-complexity synthetic regime

**Generated
invalid mass
15.59%**

**Hard filter
retained
84.41%**

**Invalid mass after hard filter
0% in retained ensemble**

Both studies use synthetic data. Results support method feasibility—not operational effectiveness.


[3] Romei de Socio, Pozzato & Merlo — TRACTA: Benchmarking Temporal Reasoning over Semantic Trajectories, <https://doi.org/10.48550/arXiv.2607.22365>.

[4] Romei de Socio, Pozzato & Merlo — TRACTA: Benchmarking Temporal Reasoning over Semantic Trajectories, <https://doi.org/10.48550/arXiv.2607.21421>.



BACKUP 3 — Integration architecture and maturity

Every status label applies to its own block; arrows show the proposed evidence flow.



No status applies to the complete diagram as an integrated deployed system.