

What is different about AI?

Valerie Sitterle



Plenary Panel: Humans and AI Systems [Sep 2025]

Artificial Intelligence Strategy for the Department of War [Jan 2026]

Speed to capability and alignment with commander intent are treated as opposite ends of the same axis.

Timing problem vs Capability problem. Different problems.




Combined Air Operations Center at Al Udeid Air Base, Qatar. (Air Force photo)



Wilcox and Metcalf (2026)

AI Command and Staff—Operational Evidence and Insights from Wargaming. *Military Strategy Magazine*, 10(4).



Today's agentic AI are often implemented as:
generative AI + tools/memory/autonomy to act

System possesses functional agency if it can:

- (i) generate actions toward an objective
- (ii) represent a model of how actions relate to outcomes
- (iii) adapt its behavior when that model changes

[Miehling et al., "Agentic AI Needs a Systems Theory" (2025)]

- Pursue persistent but mutable objectives
- Perceive environment
- Plan over multiple steps
- Act to affect other systems
- Adapt internal state

Framing: 2 Pathways in Military Decision Making

Integrated Sensing and Communications (ISAC) & Electromagnetic Spectrum Operations (EMSO)

Agentic in the decision sense.
No human in the primary loop.

Common Operating Picture (COP) Synthesis & Command Staff Augmentation

Agentic in the staff sense.
Human must judge.

Both are moving targets, not fixed baselines to measure against once

Legacy – Closed World Assumption.
State-of-the-art – Open-World Assumption (OWA).

Assumes operational environment is incomplete, out-of-distribution signals will appear, and explicitly engineered to detect that incompleteness.

Legacy – human staff supported by fusion software.
State-of-the-art – agentic AI increasingly prevalent.

Evolved from a classification system with a human reading the output to an agentic synthesis layer that reasons across retrieved intelligence and proposes interpretations a human must judge.

Difference in Kind instead of Scale: Unbounded Open-World Generation

Common misconception: "*agentic AI is open-world, legacy <X> is closed-world.*" **False.**

Difference in kind is what happens when a system hits incompleteness:

Bounded open-world recognition → reject/novelty option is a designed, first-class output

Unbounded open-world generation → no structural equivalent; model always produces something coherent – continues generating a syntactically and semantically well-formed response, filling the gap statistically from patterns learned

ISAC & EMSO

Difference between system whose response to incompleteness is bounded

vs

an unbounded, undifferentiated continuation of output behavior.

COP & Staff

Legacy tools always had analog of reject option, missing, or unidentified data.

vs

Now more likely to return plausible gap filling narrative/output.

Difference in Kind from Interplay Across 4 Interacting Properties

STRUCTURAL ROOT

Unbounded output space with non-observable failure.

Direct consequence of open-world generation; source from which everything else follows.

OPERATIONAL CONSEQUENCES

Non-monotonic reasoning across context

System behavior is sensitive to framing, context, and presentation in ways that cannot be fully characterized by specifying inputs and interfaces

Semantic goal drift

Agentic AI systems can revise the practical meaning of a goal during operation

WHERE CONSEQUENCES
BECOME OPERATIONALLY
HAZARDOUS & COMMANDERS
ARE MOST EXPOSED

Shift from evidence to argument

Agentic AI increasingly provides arguments in the form of synthesized narratives, reasoned recommendations, explanations of adversary intent, and courses of action with embedded justifications. Humans verify evidence, but they must judge arguments.

Mitigations not Interchangeable across Contexts

Integrated Sensing and Communications (ISAC) & Electromagnetic Spectrum Operations (EMSO)

Physics Informed Neural Networks (PINNs)
and CBOs work because the domain is
physics-grounded and actions are
constrainable

Common Operating Picture (COP) Synthesis & Command Staff Augmentation

No/less physics to constrain against; CBOs
can block unauthorized actions but can't
touch the reasoning/narrative itself

Shared across both: RAG, RLHF, red-teaming – improve statistical reliability, reduce
frequency of factual drift or known failure modes.

Physics Informed Neural Networks (PINNs)
Constraint-Based Ontologies (CBOs)
Retrieval-Augmented Generation (RAG)
Reinforcement Learning from Human Feedback (RLHF)

*Current mitigations shift locus of risk from content
of the output to the reliability of the reasoning
process in both contexts. Do not alter the
underlying structural basis of behavior.*

Reframing the Question

What operational confidence threshold is sufficient for agentic AI, in this role, at this echelon, and how quickly can that threshold be established?

Operational confidence does not require perfection.

Different roles and different contexts carry different thresholds and assurance requirements.

T&E reframed from speed penalty to the mechanism by which meaningful operational speed is achieved

Speed in deployed solutions may best be achieved by engineering when and where humans enter the system

Reframe risk according to military context and role:

Locus of control for agentic AI is not inside the model, but at the interfaces between the AI and the humans, systems, and operational contexts it acts within.

Usable operational superiority.



ACQUISITION INNOVATION
RESEARCH CENTER

Thank you!

Email the presenter

✉ vsitterl@stevens.edu

Stay connected with SERC/AIRC Online

