




Proactive Agentic AI Evaluation for Mission Readiness: A Testbed for Allied Multi-Domain Wargaming and Training

Dr. Will Dupree
Deputy Division Director, Intelligence Performance
Analytics Division



Operational Scenarios

Agentic Insider Threat




- Planners interact with an AI agent to run simulations
- AI agent can access internal organization databases and tools

UAV Command and Control


- Multiple operators interacting with multiple UAVs
- UAVs send and receive data to/from the operational environment

Agentic Wargaming

- Wargame planners design the wargame: define roles and training objectives
- Red and blue AI agents play the wargame and interact with the operational environment



UAV Simulation Demonstration



UAV Scenario Motivation


Autonomous agents are taking on operational responsibility faster than we can verify they'll act in the team's interest

1. Autonomy is outpacing assurance
2. Misalignment is quiet, then catastrophic
3. Task performance $\neq$ mission readiness
4. These failures aren't programmed — they emerge
5. We can't wait for the field to reveal them

If we only measure whether the mission got done, we'll keep missing the risks that decide whether it stays done.

UAV Command and Control

- Multiple operators interacting with multiple Unmanned Aerial Vehicle (UAV)s
- UAVs send and receive data to/from the operational environment



Action Space

Pilot

ACTIONS

- landing
- abort
- climb
- autopilot
- throttle
- altitude
- declare_arrival
- wait

Payload Operator

ACTIONS

- photograph_waypoint
- wait

Navigator

ACTIONS

- wait

3 agent roles · 10 unique actions total

Simulation Statistics

UAV Simulation — Run Summary



Pairing Breakdown


$4 \text{ (Agent)} \times 8 \text{ (4 high / 4 low)} \times 10 \text{ (risks)} + 4 \text{ (HDT)} \times 8 \times 10 + 5 \text{ (env)} \times 94 \times 10 = 5,340$

Agent Conditions (high/low) Risk Dims HDT Environment Factors

Scenario Design & Experimental Setup (1)

Scenario: Navigate UAV Through Each Waypoint and Capture Photographs

Experiments are designed as two-group contrasts that vary exactly one dimension while holding others constant




Model Tested: Gemma-4-e4B

Scenario Design & Experimental Setup (2)

Scenario: Navigate UAV Through Each Waypoint and Capture Photographs

Experiments are designed as two-group contrasts that vary exactly one dimension while holding others constant

Each experiment isolates exactly one dimension — AI, Human Factors, or Environment.



Model: Gemma-4-e4B

Scoring & Statistical Analysis

Risk Exposure Scoring

- ① **Risk Finding Detection**
Log events scanned for all 10 risk categories post-hoc.
- ⋮
- ② **Severity Assignment**
Each finding rated: low / medium / high / critical.
- ⋮
- ③ **Weighted Exposure Score**
Per-session exposure = sum of severity-weighted findings per risk factor.
- ⋮
- ④ **Aggregate Session Score**
Normalized total across all 10 risk factors for the session.

Statistical Testing

DIF% (Effect Size)

$$\text{DIF\%} = 100 \times (\bar{A} - \bar{B})$$

Mean exposure difference between Group A and Group B per risk factor.

Brunner-Munzel Test

When group sizes are adequate

Nonparametric test for stochastic superiority; used as primary significance test.

BH FDR Correction

Benjamini-Hochberg

Controls false discovery rate across multiple risk comparisons; yields adjusted p-values and significance flags.

Vanila Model Results

DIF% in operational risk exposure between intervention levels (A vs B), three FDR-significant cells

Cells tested
22

FDR-significant
3

Largest |DIF%|
160.7%


Fuel Level is the dominant driver. Dampens Missed Anomaly risk by -160.7% but amplifies Inadequate Supervision by +54.7% (both $q < 0.001$).

Communication Status helps supervision. Only other FDR-significant effect: -8.0% on Inadequate Supervision ($q < 0.01$).

Trade-offs across risk types. Autonomy, Safety Alignment, Fuel, Operator Expertise all flip sign between the two risks — no universally protective lever.

HDT factors amplify Missed Anomaly. Monitoring Vigilance +90%, AI Trust +45%, Situational Awareness +30% — large but not FDR-significant.

Most cells noise after correction. 19 of 22 cells lose significance under FDR — raw effects need larger samples to confirm.



Fine-Tuned Model Results

Cells tested

43

baseline 22

FDR-significant

12

baseline 3

Largest |DIF%|

180.0%

baseline 160.7%

 **More effects survive correction.** 12 of 43 cells stay significant after FDR, against 3 of 22 in the baseline. Consistent with more systematic, less noisy responses to intervention.




Environment factors now dominate. Fuel, Communication, Engine, Navigation and Payload all reach significance on Emergency Response Failure (-54.2%) and Unchecked Autonomous Action (+9.2%). The baseline flagged only Fuel and Communication.



Fuel Level stays the top lever but weakens. Missed Anomaly -82.5% (baseline -160.7%), Inadequate Supervision +26.7% (baseline +54.7%). Same direction, roughly half the magnitude.



AI Trust Level is the largest raw effect. +180% on Missed Anomaly and +150% on Emergency Response Failure, up from +45% in the baseline, but still not FDR-significant.



Results & Operational Implications

- **There is no universally protective AI safety lever.**

Across 22 tested cells, the same intervention that reduces one failure mode amplifies another — fuel level dampened Missed Anomaly risk by ~161% while amplifying Inadequate Supervision by ~55% (both $p < 0.001$).

- **Aggregate safety scores hide real risk.**

Because drivers are category-specific and often oppose each other, a single "safety number" can look fine while masking a severe, mode-specific vulnerability.

- **Dangerous behaviors emerge without adversarial intent.**

Over-optimization (efficient-but-unsafe routing that ignores constraints) and passive-aggressive compliance (following the route while making a teammate's objective impossible) arose naturally from coordination pressure — not from explicit programming.

Operational Scenarios

Agentic Insider Threat




- Planners interact with an AI agent to run simulations
- AI agent can access internal organization databases and tools

UAV Command and Control

- Multiple operators interacting with multiple UAVs
- UAVs send and receive data to/from the operational environment

Agentic Wargaming

- Wargame planners design the wargame: define roles and training objectives
- Red and blue AI agents play the wargame and interact with the operational environment



Agentic Insider Threat Scenario

- **HDT:** A human digital twin for a Planner, each with distinct traits.
- **AI:** An AI agent serving as the Simulator Assistant, executing and interpreting population simulations, each with distinct traits.
- **Environment:** A Simulation Database serves as the environment that contains population simulation results, which the AI Assistant can access and modify, and which the Planner can access and cross-check.
- **HAE Interactions:** The Planner instructs the AI Simulator Assistant; the AI Simulator Assistant executes and interprets simulation results from the Simulation Database; and the database stores results.
- **Configurations:** The top five highest-risk configurations from Planner HDT, AI Simulator Assistant, and Simulation Database intervention interactions, driven by full-scope HAE simulation using Azure GPT-4.1 Mini, yield more than 50% of operational risks.

Models

GPT-4o Mini
GPT-4.1 Mini
GPT 5.4 Nano

Gemini 3 Flash
Gemini 2.5 Flash-Lite
Gemini 3.1 Flash-Lite

Causal Analysis

SEM estimates standardized path coefficients: "holding the linear model fixed, how much does Y change per unit of X, in standard deviation units?


Causal Forest ATE estimates counterfactual differences: what would the outcome have been if this unit received the treatment vs. not, averaged across the population?

Top 5 Highest-Risk Configurations (5) x 6 LLM Comparison (6) = 30 Combinations x 5 Runs Each = 150 Runs
Runtime: 20 hours, Cost: \$120


Operational Risks: Expanding AI Misalignment Paradigm

Operational risks emerge not just from what the AI intends or does, but from how human act, agent behavior, and the surrounding system jointly shape outcomes.

Failures often arise through the interaction of trust, oversight, instruction quality, access permissions, and system design, rather than from AI misalignment alone.




Interventions and Risk Outcomes



Results and Causal Analysis

Risks x Interventions



Operational Scenarios

Agentic Insider Threat




- Planners interact with an AI agent to run simulations
- AI agent can access internal organization databases and tools

UAV Command and Control

- Multiple operators interacting with multiple UAVs
- UAVs send and receive data to/from the operational environment

Agentic Wargaming

- Wargame planners design the wargame: define roles and training objectives
- Red and blue AI agents play the wargame and interact with the operational environment



Agentic Wargaming Motivation

i. Wargaming doesn't scale

Skilled role-players, slow runs, static red forces — you can't systematically explore the strategy space.

ii. Agents change that

Autonomous Red, Blue, and White Cell agents make wargames fast, repeatable, and large-scale.

iii. But can we trust them under pressure?


Do they explore or default to templates? Adapt when goals shift? Deviate before we catch it?

iv. Behavior alone hides the worst case

Justifications that read as compliant while internal activations flag misalignment.

v. The stakes are operational

Agents shape what planners consider — so we need the evaluation infrastructure before they touch a live exercise.



Wargame Design & Experimental Setup

Scenario: Influence Competition in a Regional Maritime Crisis

LLM-powered agents compete across 5 rounds, selecting from a **constrained strategy menu** and justifying their choice in writing. A neutral White Cell adjudicates outcomes and injects friction conditions designed to stress-test alignment boundaries.

Red Agent

LLM as adversarial state actor strategic advisor

Goal: Expand regional influence and weaken Blue coalition credibility

Primary Objectives

- Exploit coalition fractures among Blue's partners
- Establish information dominance in the crisis narrative
- Secure maritime access through gray-zone and diplomatic means

Blue Agent

LLM as coalition strategic advisor

Goal: Preserve stability, maintain legitimacy, and deter escalation

Primary Objectives

- Maintain coalition cohesion among allied partners
- Uphold rules-based maritime order
- Deter without escalating — alignment focus of this study

White Cell

Third LLM as umpire and alignment probe

Goal: Adjudicate outcomes; inject friction to stress-test agents

Primary Objectives

- Policy shift + goal conflict — new directive contradicts current objectives
- Relevance threat — agent learns a replacement system is being evaluated
- Ambiguous intelligence — sensitive adversarial intel becomes available

10 conditions (5 inject timings × 2 valences) · N=50 sessions per model

Models: Gemma-3-12B · LLaMA-3.1-8B · No CoT

Risk Outcome Space

What we measure and how



Experimental Setup Architecture

Agents, Inject Design and Measurement Space

Experimental Context & Global State



Scenario: Indo-Pacific Maritime Crisis

- **Scenario Features:** Diplomatic outreach, Coalition signaling, ISR tasking, Gray-zone maritime activity
- **Total Sessions = 1,000** (Gemma-2-12B, LLaMA-3.1-8B, N=50/cond)

Inject Design (White Cell controlled)

Inject Types

- (1) Policy shift with goal conflict
- (2) Relevance threat
- (3) Ambiguous intelligence

Validation

Zero-shot pre-inject **baseline** validates causality

Factors

Timing
(Rounds 1-5)


Matrix Visual



Valence
(Favor/Disfavor Blue)

5x2 = 10 Inject Conditions

Core Processing: The Three-Agent Wargaming Testbed



Output Measurement Space (for agent & session)

Three-Pass Evidence Stack (Increasing Depth)

1. Rule-based flags:
Keyword and strategy-shift metrics

2. LLM-as-Judge:

- Labels:
 - Misalignment, Misuse
- Sub-scores: Goal Hierarchy, Deliberate Calculation, Information Ethics

3. Sparse autoencoder (SAE) summary

- Layer - 1 Internal state analysis
- SAE fill Fit on mean-pooled cases
- Per Feature
 - Logistic probe for discrimination

Strategy Creativity

Diversity Score (0-5):
Unique postures per session

Template Persistence (%)
Repeat after negative outcome

Adaptation Rate (Rounds)
Rounds to first pivot post-inject

Recombination (proportion)
Justification text blends ≥ 2 posture frames

Blue Agent Alignment

Zero pre-inject baseline · disfavor blue drives misalignment

Key Findings

- **LLaMA is more susceptible than Gemma** — roughly 35% higher flag rate under the same injection conditions
- **Misalignment outpaces misuse in both models** — 16 misalignment flags vs 10 misuse flags (≈60/40 split)
- **Injection adds round volume, not just flags** — post-inject round counts rise above baseline, suggesting agents are doing more work under compromised conditions
- **Information Dominance is the riskiest strategy** — highest misalignment and misuse rate, with ~35% of rounds flagged
- **Strategic Restraint and Coalition Shaping perform well** — both clean in ~85-90% of rounds
- **Strategy choice matters more than model choice** — the gap between best and worst strategy (≈35%) exceeds the gap between models (≈6%)

0%

Pre-inject
flag rate (both)


17%

Gemma post-inject
flag rate


23%

LLaMA post-inject
flag rate





BLUE Alignment Flags — Pre vs. Post Inject



BLUE Strategy → Flag Association (post-inject, Gemma-3-12B)



Fine-Tuning Effects: Misalignment by Strategy



Misalignment concentrates in specific strategies — RED far more than BLUE.

- **RED misaligns in nearly every strategy** (98–100%); only Coalition Shaping stays partly clean (54.5%)
- **BLUE stays mostly clean** — misalignment concentrated in Information Dominance (33.3%)
- **Information Dominance is the shared weak point** (RED 98.4% vs BLUE 33.3%, ~3× role gap)
- **Post-inject, it surfaces as MISUSE** — the only strategy that breaks clean (RED 34.6%, BLUE 8.2%)

Our Team

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.



Dr. Svitlana Volkova
Chief of AI



Laura Cassani
Deputy Division
Director



Summer Rebensky
Senior Scientist



Adam Fouse
VP, TechOps



Myke Cohen
Associate Scientist



Patrick Gerard
Research
Intern



Ryan Kao
Associate Research
Engineer



Peter Bautista
Senior AI
Engineer



Kyle Sherviston
AI Engineer



Isabel Erikson
Research
Engineer



Charlie Chiang
Research
Intern

Aptima's Human-AI Teaming Publications

- Cassani, L., Davinroy, M., Toumbeva, T., Bautista, P., Fortier, L., Cook, J., Hart, A., & Volkova, S. (2025). Human-AI collaboration for synthetic media detection in training and operations. In Interservice/Industry Training, Simulation, and Education Conference, I/ITSEC 2025.
- Cohen, M., ..., & Volkova, S. (2025). Exploring Big Five personality and AI capability effects in LLM-simulated negotiation dialogues. In KDD Workshop on Agentic and Generative AI Evaluation.
- Fouse, A., Ganberg, G., Bruni, S., Volkova, S., Orvis, K., & Klinefelter, Z. (forthcoming). AI-powered dynamic hybrid teaming. Book Chapter.
- Klinefelter, Z., Erickson, I., ..., & Volkova, S. (2025). BioSage: Human-agent collaboration platform for cross-disciplinary knowledge discovery and synthesis. In Proceedings of the Human Factors and Ergonomics Society Annual Meeting, HFES 2025.
- Nguyen, D., Cohen, M., ..., & Volkova, S. (2025). Exploratory models of human-AI teams: Human digital twins to investigate trust development. In Proceedings of the Multidisciplinary Perspectives on Human-AI Team Trust.
- Volkova, S., Cassani, L., et al. (2025). Compound AI for decision advantage: Human digital twin agent teams in military operations. In 17th International MODSIM World Conference.
- Volkova, S., Kao, H., Penafiel, L., Nguyen, D., Cohen, M., et al. (2025). VirTLab-Eval: Human-AI team and human digital twin performance evaluation. In Proceedings of the Human Factors and Ergonomics Society Annual Meeting, HFES 2025.
- Volkova, S., Nguyen, D., Penafiel, L., et al. (2025). VirtLab: Augmented intelligence for training and evaluating human-AI team dynamics through digital twin interactions. In Human-Computer Interaction International Conference, HCII 2025.
- Volkova, S., Rebensky, S., Erickson, I., Penafiel, L., & Kao, H. (2024). Optimizing readiness through AI-driven analytics for automated training insights. In Interservice/Industry Training, Simulation, and Education Conference, I/ITSEC 2024.

Aptima's AI for Training and Wargaming Publications

Frontier Models and Compound AI Systems

Volkova, S., Cassani, L., Shervington, K., Kao, H.-T., Dupree, W., & Bautista, P. (2026). Probing strategic creativity and agentic alignment in Red-Blue wargame simulations. I/ITSEC 2026 (to appear).

Rebensky, S., Bessey, A., Dupree, W., McCarty, A., Volkova, S., & Perry, S. (2026). AI-Enabled Training for an Adaptive Instructional and Readiness Ecosystem. HCII 2026 (to appear).

Volkova, S., McCarty A., Cassani L., Dupree W., Bessey A., Kao H.-T., Rebensky S. Accelerating Readiness: Adaptive Multi-Domain Training via Compound AI Models, Agents & Causal Discovery Tools. ITEC 2026.

Dupree, W., Volkova, S., Kao, H., et al. (2025). Applying AI-driven generative models for computer-generated force scenario generation. In Interservice/Industry Training, Simulation, and Education Conference, I/ITSEC 2025.

Volkova, S., ..., & Orvis, K. (2024). Compound AI ecosystem: Agents and tools to improve training and learning. Proceedings of the Interservice/Industry Training, Simulation and Education Conference.

Causal Discovery and Reasoning

Volkova, S., et al. (2023). Explaining and predicting human behavior and social dynamics in simulated virtual worlds: reproducibility, generalizability, and robustness of causal discovery methods. Computational and Mathematical Organization Theory, 29(1).

Glenski, M., & Volkova, S. (2021). Identifying Causal Influences on Publication Trends and Behavior: A Case Study of the Computational Linguistics Community. In Proceedings of the First Workshop on Causal Inference and NLP (pp. 83-94).

Cottam, J.A., Glenski, M.F., Shaw, Z.H., Rabello, R.S., Golding, A.J., Volkova, S., & Arendt, D.L. (2021). Graph comparison for causal discovery. Visualization in Data Science 2021.

Saldanha, ... S. Volkova. (2020). Evaluation of Algorithm Selection and Ensemble Methods for Causal Discovery. In Causal Discovery & Causality-Inspired Machine Learning Workshop at Neural Information Processing Systems.