




Adaptive Doctrine Selection under Uncertainty

AN EXPLAINABLE APPROACH FOR IADS SIMULATION

MESAS 2026

Outline

1

Motivation and Domain

2

Doctrine

3

Method

4

Experiment

5

Results


6

Summary and Outlook

The Use Case: An IADS Opponent for Simulation and Training

Training and testing need an opponent that behaves like a real air defense

- **Integrated air defense systems (IADS)** challenge air superiority
- Live testing is expensive and covers few configurations and scenarios
- Simulation provides repeatable opponents for training, testing, and capability analysis
- The simulated opponent must act plausibly, and its behavior must be explainable



Which behavior should the simulated air defense use, and can it adapt while the attack unfolds?

The Current Threat: Mixed and Massed

Example: attacks on the United Arab Emirates, reported on 8 April 2026

- **Mixed:** drones, cruise missiles, and ballistic missiles in one campaign
- **Massed:** thousands of threats since the attacks began, so interceptors become scarce
- **Different flight profiles:** low-level routes and ballistic trajectories

Same weapons, two attack strategies

Force Flow

spread over time and space



Force Packaging

concentrated in time and space



The defense must use limited interceptors against a mixed threat it only partly sees.

A Mixed Air Threat

UAE Ministry of Defence

8 April 2026

Reported on 8 April 2026



UAVs

35



Ballistic missiles

17

Cumulative figures
since the onset of attacks



UAVs

2,256



Cruise missiles

26



Ballistic missiles


537

Source: UAE Ministry of Defence (@modgovae)
8 April 2026. Figures as reported.
English adaptation of the supplied post

x.com/modgovae/status/2041835088053420543

The Model: A Layered Red IADS

Radars and launchers grouped into batteries, units, and one IADS team, in three range layers



Our scenario: 9 units, 17 batteries

1
unit

Long range

Early-warning radar (600 km); launchers up to 100 km; the only layer against ballistic missiles

3
units

Medium range

3 batteries each; launchers up to 40 km against aircraft, cruise missiles, drones

5
units

Short range

1 battery each; launchers up to 12 km against aircraft, cruise missiles, drones

Detection: radar equation, terrain line of sight.

Engagement: kill probability by missile, target, range.

Doctrine: Definition and Implementation

Five policy dimensions turn a doctrine into executable rules for the Red CGF

Doctrine: a high-level behavioral policy for how a force employs its capabilities under mission conditions. It sets the frame; tactical engagement decisions happen inside it.

	Command level	Track ordering	Target assignment	Launcher selection	Firing policy	Magazine policy
Doctrine 1	IADS (centralized)	Highest kill probability first	Defense in depth	Highest kill probability	Shoot-look-shoot	1 missile in reserve
Doctrine 2	Battery	Nearest to a launcher first	Shortest-range battery that can engage	Nearest launcher	Shoot-shoot	No reserve
Doctrine 3	Unit	Nearest to a battery first	Least-engaged target per unit	Highest kill probability in unit	Shoot-look-shoot	2 to 4 missiles in reserve

Shared by all doctrines: radar policy and a self-defense rule. **Not modeled:** organization, command relationships, emission control, deception.
So this is a doctrine configuration, not a complete military doctrine.

Earlier Work: Ranking Fixed Doctrines Offline

Monte Carlo simulation plus game-theoretic analysis: no doctrine wins against both strategies

Method




Red's best response

share of 5,000 sampled games

	Force Flow	Force Packaging
Doctrine 1	17.7 %	3.0 %
Doctrine 2	23.9 %	83.3 %
Doctrine 3	58.4 %	13.8 %



Doctrine 2 under robust criteria

share of games selecting Doctrine 2



A fixed doctrine is a compromise. Doctrine 2 is the robust choice if Blue's strategy is unknown.


If Red could recognize the strategy during the attack, it could switch. That is this paper.



8

Method: Online Doctrine Adaptation

A closed loop that runs once per 30-s decision epoch



Orange: new in this paper. Identical inputs give identical decisions, so every run can be re-derived after the fact.

Method 1: Belief Tracking with a Discriminative HMM

Hidden Markov model (HMM): Blue's strategy is the hidden state, Red's air picture the observation

Observation

18 features from the fused Red air picture: counts of newly detected hostile tracks over 30 to 300 s windows, plus the spatial pattern of recent detections. Standardized with training statistics.

1 Predict

$$b_t^-(s) = \sum_{s'} T_{s's} b_{t-1}(s')$$

Markov persistence $T_{ss} = 0.95$ per epoch

2 Evidence

$$q_t = \sigma(\beta_0 + \beta^T z_t), \quad \psi_t(s) = \left(\frac{q_t(s)}{\pi(s)} \right)^\lambda$$

regularized logistic regression, tempered with $\lambda = 0.2$

3 Update

$$b_t(s) \sim b_t^-(s) \psi_t(s)$$

fewer than 7 observed features: $\psi_t = 1$, the belief only propagates

Why this design

- **Discriminative:** no generative model for correlated, partly missing detection features
- **Full belief:** uncertainty is kept, not reduced to a hard label
- **Tempering:** keeps the belief from becoming overconfident

Training protocol

- Classifier fitted on 36 fixed-strategy runs
- Evidence weight and persistence tuned on separate runs
- Frozen model passed a predefined acceptance gate
- Evaluation only on disjoint, held-out seeds

Method 2: Expected Utility and Guarded Switching

Rank the doctrines under the belief, then switch only if every guard passes

Offline payoff $\bar{U}(d, s)$

	Force Flow	Force Packaging
Doctrine 1	0.768	0.681
Doctrine 2	0.780	0.763
Doctrine 3	0.790	0.720

$$EU_t(d) = \sum_s b_t(s) \bar{U}(d, s)$$

$$\Delta_t(d) = EU_t(d) - EU_t(d_{\text{act}}) - c_{\text{sw}} [d \neq d_{\text{act}}]$$

switching cost $c_{\text{sw}} = 0.002$; ties favor the active doctrine

Leave Doctrine 2 for Doctrine 3 only if $P(\text{Packaging}) < 0.193$



Guards: all must pass, otherwise keep the active doctrine

Net improvement

$\Delta_t \geq 0.001$ after switching cost

Minimum dwell time

300 s (10 epochs) since the last switch

Confirmation


same candidate in 2 consecutive epochs

Switch limit

at most 4 switches per run

Record: each guard's value, threshold, and outcome, the first failed guard as the reason, and a plain-text explanation from a fixed template

Four attack scenarios (6,000 s, 199 decision epochs)



Hypotheses

- H1** The tracker beats an uninformed prediction and the per-epoch classifier without filtering (Brier score, log loss)
- H2** Guards reduce doctrine switches compared with unguarded adaptation
- H3** Exploratory: effect of guarded adaptation on mission utility vs. fixed Doctrine 2

Eight selectors, all starting with Doctrine 2

- Fixed Doctrine 1, 2, or 3
- Guarded, full belief (**ours**)
- Guarded, most likely strategy only (MAP)
- Unguarded, with and without switch cap
- Payoff oracle that knows Blue's true strategy

Runs and statistics

- 1,600 runs: 4 scenarios × 8 selectors × 50 held-out seeds
- Paired by scenario and seed: differences come from doctrine choices only
- 95 % seed-cluster bootstrap, 10,000 resamples; an effect counts only if the interval excludes zero
- Plus 1,800 runs for the guard sensitivity study

Result 1: Red Recognizes Blue's Strategy

Beats both baselines, with well-calibrated confidence

72 %

of decisions name the correct strategy

-22 %

prediction error compared with guessing

90 s

median time to recognize a change to Force Flow (46 of 50 runs)

Known blind spot

The first 20 minutes of a Force Packaging attack look like Force Flow.


Early signal

Before a change to Force Packaging, the quiet gap between attack phases already resembles packaging: the belief moves before the change.

Result 2: The Guards Keep Adaptation Calm

Fewer switches, fewer unnecessary ones, and no measurable loss in mission value


Doctrine switches per run



- Guards remove 0.11 unnecessary switches per run compared with no guards
- Robust: across nine guard settings, mission value moves by less than 0.3 %
- One clear control knob: a 10-min dwell time cuts switches from 2.1 to 1.2 per run

Every Decision Is Traceable

Each decision is recorded with its reasons and can be replayed after the run



- 1 Belief: 92 % Force Flow
- 2 Doctrine values: Doctrine 3 ranks first
- 3 Decision: keep Doctrine 2, the minimum dwell time is not yet reached
- 4 Ground truth, shown for analysis only
- 5 Guard chain: which rule allowed or blocked the switch
- 6 Explanation in plain text, generated from the record

Trainers and analysts can step through a run and see why the IADS kept or changed its doctrine.

Summary and Outlook

An adaptive opponent that can be audited

Recognize

A discriminative HMM infers Blue's strategy from the Red air picture, better than both baselines

Adapt cautiously

Expected utility plus guards: fewer unnecessary switches, no measurable cost

Explain

Every decision is recorded and could be re-derived after the run

Measure the value

Strategy knowledge was worth up to 1 % here; inference captured about 80 % of it

Adaptation becomes a measurable design choice.

The framework shows when it pays off, and every decision comes with its reason.

Next

State- and time-dependent doctrine values

Tracker persistence matched to the scenario

Tactical multi-agent coordination

Evaluation with military experts and trainees


Full paper: F. Stracker, P. Stütz: Adaptive Doctrine Selection under Uncertainty: An Explainable Approach for IADS Simulation. MESAS 2026

Offline doctrine evaluation: A Monte-Carlo Game Theory Framework for Doctrine Selection. Interservice/Industry Training, Simulation, and Education Conference 2026, Paper 26280.

QUESTIONS?



Fabian Stracker
Fabian.Stracker@unibw.de
fabianstracker@bundeswehr.org



Research Gate



Linked In