Operator Intent: Graduated Fidelity Tiers for Mixed-Fleet Coordination

by Nick Clark | Published April 25, 2026

Mixed-Fleet Coordination Lacks an Intent Substrate

Coordination between autonomous units depends on knowing what other operators intend. Cooperative agents can broadcast intent directly. Non-cooperative agents — human drivers, legacy vehicles, foreign drones, adversarial entities — cannot or will not. Current architectures treat these populations as fundamentally different: cooperative agents get one protocol, everything else gets sensor-based inference, and the two are not unified.

This produces brittleness at the boundary. An autonomous vehicle can coordinate tightly with another autonomous vehicle but treats human drivers as opaque hazards. A drone with cooperative airspace data is paralyzed when a non-cooperative drone enters its window. A defense system distinguishing combatants from noncombatants has no architectural framework that scales from broadcast intent (allied units) to inferred intent (unknown units) to hostile-intent classification (adversarial units).

The result is that mixed-fleet operation gets relegated to constrained environments where everyone is cooperative, or to fully-segregated infrastructure where mixing is prevented architecturally.

1. The Primitive: Three Fidelity Tiers Fused

Operator intent is consumed at three fidelity tiers: full cognitive-state broadcast (cooperative agents publish their planning graphs, intent fields, and capability envelopes), structured partial-fidelity bus extraction (vehicles or devices publish a limited but specified subset such as turn signals, brake lights, route plans, formation orders), and behavior-inferred attribution (operating units infer intent from sensor cues including trajectory, gaze, gesture, formation, and historical pattern).

All three tiers produce intent observations into a unified composite admissibility evaluator. The evaluator weights tier 1 broadcasts most heavily, tier 2 moderately, tier 3 least, and combines them with environmental observations, governance policy, and dispositional state to produce a single coherent intent estimate.

Tier inference is itself an architectural primitive: an operating unit can be configured to consume any subset of tiers based on its policy, and the same physical neighbor can be observed across all three tiers simultaneously, with the evaluator handling agreement and disagreement between tiers structurally.

2. Behavior-Inferred Intent as Governed Observation

Behavior-inferred intent — the lowest-fidelity tier — is structurally distinct from generic 'sensor-based situational awareness.' The inference produces a governed observation that can be contributed back to the mesh, evaluated by other consumers, and challenged by the inferred operator.

When a unit infers another unit's intent (e.g., 'vehicle B is preparing to merge'), the inference is published as an observation with the inferring unit's credential, the inferred operator's identifier (or pseudonym), the inference function reference, and the supporting cues. Other observers can integrate or contradict the inference; the inferred operator can challenge it through governance-credentialed retraction.

This creates a feedback channel that current architectures lack: a unit accused of an intent it does not have can structurally retract; an inference that consistently misclassifies populations can be flagged and refined. The inference function evolves under verification feedback rather than being frozen at training time.

3. Verification-Feedback Inference Evolution

Each behavior-inferred intent observation is paired with a temporal commitment: the inference predicts that a specific behavior will follow within a specified window. After the window expires, the actual behavior is observed and compared to the prediction. The agreement is recorded as a verification observation against the originating inference.

Aggregated verification observations modulate the inference function's parameters. Inference functions that consistently match observed behavior gain weight; functions that consistently miss are demoted; new functions can be proposed and tested under sandboxed admission before being promoted into general use.

This produces a closed loop: every inference is structurally falsifiable, the falsifications accumulate, and the inference function evolves toward higher accuracy under continuous evaluation. The mechanism eliminates the brittleness of frozen-at-training-time classifiers in a domain where adversarial behavior co-evolves.

4. Risk-vs-Hostility Bifurcation Under Due Process

Current insurance-based usage telematics conflate competence-based risk (an operator who is unsafe due to inattention, fatigue, intoxication, or skill limits) with intent-based hostility (an operator who is deliberately attempting harm). The conflation produces actuarial premiums that punish low-skill drivers using the same architecture that should be reserved for adversarial classification.

The intent primitive separates the two. A risk profile is constructed from observed behavior under normal-operation assumptions and is used for actuarial purposes. A hostility profile is constructed from behaviors structurally indicative of adversarial intent (deliberate counter-flow, targeting trajectory, weapon-deployment cues) and is used for protective response.

Critically, hostility profile construction requires due-process credentialing: a regulatory or judicial authority must have credentialed the criteria under which hostility classification operates, the classification event must be governed by audit-grade lineage, and the classified operator has structural standing to challenge. This puts hostility classification on the same footing as protective orders, restraining orders, and other due-process-bound classifications.

5. Counter-Action Selection Under Hostility Classification

When an entity is classified hostile, the operating unit's response set is enlarged: defensive maneuvers, hardened postures, evasive routing, broadcast alerts to allied units, and (in defense contexts) escalation to authorized counter-measures.

Counter-action selection is itself governed by composite admissibility: the hostile classification produces a permissive admissibility envelope that the unit may but need not use, with the actual selection determined by environmental observations, mission policy, and confidence-governed actuation.

Counter-action is therefore not 'hostility = automatic counter-attack.' It is 'hostility = expanded admissibility envelope under which counter-actions become available, subject to the same gating that any other actuation receives.' The classification widens the option space; the gating decides what to do within it.

6. Cross-Domain Adversarial Inference

The same intent primitive scales from civilian mixed-fleet coordination through commercial fleet management, emergency response, defense ISR, and homeland security. The fidelity tiers, the inference function, the risk-vs-hostility separation, and the due-process credentialing remain identical; the configurations change.

A civilian autonomous vehicle uses tier 2 (turn signals, brake lights) and tier 3 (trajectory inference) to coordinate with non-cooperative human drivers, with hostility classification disabled by default. A defense unit operating in contested airspace uses tier 1 (allied broadcasts), tier 2 (transponder data), and tier 3 (behavior inference), with hostility classification enabled under credentialed mission policy.

Cross-domain consistency means the same architecture serves civilian, commercial, public-safety, and defense use cases with configuration changes rather than re-implementation.

7. What This Is Not

This is not Cambridge Mobile Telematics, Nauto, or Lytx. Those products produce behavior-based usage scores from sensor data without an architectural distinction between competence-based risk and intent-based hostility, and without a due-process framework for adverse classifications.

This is not Mobileye REM or HERE high-definition map ingestion. Those systems aggregate observation data into shared maps; the intent primitive aggregates intent observations across heterogeneous fidelity tiers with structural bidirectional retraction.

This is not LAWS (lethal autonomous weapons systems). Those systems' hostility classification is mission-policy-internal. The governed primitive externalizes the classification criteria and audit lineage, supporting accountability across civilian, commercial, and defense deployments under their respective credentialing frameworks.

Conclusion

The operator intent primitive unifies cooperative broadcast, structured partial extraction, and behavior-inferred attribution under a single composite admissibility framework, with verification-feedback evolution and explicit risk-vs-hostility separation.

Mixed-fleet operation — between autonomous and human-driven vehicles, between cooperative and non-cooperative drones, between allied and adversarial entities — moves from special-case workarounds to a unified architectural primitive. This is disclosed under USPTO provisional 64/049,409.