Tier-Weighted Admissibility

by Nick Clark | Published April 25, 2026

What Tier-Weighted Admissibility Specifies

The architecture's admissibility evaluator applies different default weights to contributions from different fidelity tiers. Tier 1 contributions (full cognitive-state broadcast) carry the highest weight because the broadcasting agent is reporting its own intent with full fidelity. Tier 2 contributions (structured partial signals) carry moderate weight because the signals are constrained but structurally interpretable. Tier 3 contributions (behavior-inferred attribution) carry the lowest weight because the inference may be wrong.

The weights are governance-credentialed and configurable per consumer. A defense-context consumer may weight tiers differently than a civilian-context consumer; the weighting is the consumer's policy choice within the architectural framework.

Why Equal-Weight Fusion Produces Wrong Estimates

Equal-weight fusion across heterogeneous fidelity sources produces estimates dominated by whichever tier has the most contributions. In environments where most entities are non-cooperative (urban traffic with mostly human-driven vehicles), equal-weight fusion would have Tier 3 inferences dominate, drowning out the higher-fidelity Tier 1 contributions from the few cooperative vehicles.

Tier-weighted admissibility produces estimates that respect the structural fidelity differences. The Tier 1 broadcasts contribute with their full informational value; Tier 3 inferences contribute appropriately discounted; the resulting estimate is more accurate than equal-weight fusion.

How Cross-Tier Agreement Modulates Weighting

When Tier 1, Tier 2, and Tier 3 contributions about the same entity agree, the composite confidence increases — multiple independent observation streams confirming the same intent. When they disagree, the architecture surfaces the disagreement for downstream evaluation.

Adversarial-aware operation consumes the disagreement as a structural signal. A vehicle whose Tier 1 broadcast disagrees with its Tier 2 partial signals or Tier 3 behavior inference is structurally suspect — possibly compromised, possibly malfunctioning, possibly adversarial. The architecture flags the disagreement through credentialed observations.

What This Enables for Robust Mixed-Fleet Operation

Robust mixed-fleet operation gains structural support that simpler fusion patterns cannot provide. The architecture handles the realistic mixed-population operating reality (a few cooperative agents broadcasting Tier 1, more entities visible at Tier 2, most entities only inferable at Tier 3) with appropriate weighting.

Defense and security operations gain the disagreement-detection capability that adversarial-aware operation requires. The patent positions the primitive at the layer where mixed-fidelity intent fusion has been operating without architectural support.