

# Autonomous + Human-Driven Mixed-Fleet Coordination

by [Nick Clark](#) | Published April 25, 2026

## What Mixed-Fleet Coordination Actually Requires

Real autonomous deployment shares the road with human-driven vehicles. The autonomous vehicle must coordinate with the human-driven vehicle without relying on cooperative protocols the human driver does not have access to. Current architectures handle this through computer vision and motion modeling — the human driver becomes a computer-vision target rather than a coordination partner.

Three-tier intent fusion treats the human driver as a Tier 2 + Tier 3 source: Tier 2 from observable signals (turn signals, brake lights, hazard lights, horn) and Tier 3 from inferred attribution (trajectory, gaze if observable, formation, historical pattern). The autonomous vehicle is a Tier 1 + Tier 2 + Tier 3 source. Both populations contribute to the same composite admissibility evaluation.

## Why Treating Humans as Computer-Vision Targets Limits Deployment

The computer-vision-target pattern produces structurally limited coordination. The autonomous vehicle observes the human-driven vehicle's motion and infers intent from the motion alone. Tier 2 signals (turn signal, brake light) are consumed informally if at all; the architecture has no first-class concept for them. Tier 3

inference is a single category with no corroboration mechanism across heterogeneous observation sources.

Three-tier fusion is structural. Each tier feeds the same admissibility evaluator with declared fidelity. Cross-tier corroboration handles agreement and disagreement explicitly. Mixed-fleet operation moves from special-case workaround to architectural primitive.

## **How Three-Tier Fusion Operates Across the Fleet**

The composite admissibility evaluator runs continuously per neighboring entity. For each entity in observation range, the evaluator collects available Tier 1 broadcasts (V2X cooperative messages where present), Tier 2 signals (visual indicators, transponder data, structured partial broadcasts), and Tier 3 inferences (motion-based attribution).

The output is a per-entity intent estimate with declared confidence and fidelity. The autonomous vehicle's planning consumes the estimate as a credentialed observation, weighting it appropriately for the planning horizon. Cross-entity coordination — formation maintenance, lane changes, intersection negotiation — operates through the same mechanism for all observed entities regardless of which tier they're contributing through.

## **What This Enables for L4/L5 Commercial Expansion**

The narrow-geography constraint that has held L4 to specific urban deployments largely reflects the difficulty of mixed-traffic operation. Geographies where human-driver-population density is high (most of the United States outside specific neighborhoods) defeat the computer-vision-target pattern at operational scale.

Three-tier fusion expands the operational domain. Autonomous vehicles coordinate with human drivers structurally rather than as obstacles. The architecture also supports mixed-cooperation transitions: as V2X deployment progresses, human-driven vehicles may broadcast Tier 2 partial-fidelity signals through aftermarket transponders, gradually upgrading their intent fidelity without requiring full L4 capability. The patent positions the primitive at the layer mixed-fleet expansion has been waiting for.