

Fleet-Scale Active Perception

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

What Fleet-Scale Active Perception Specifies

The active-perception research literature (descending from Bajcsy 1988) has explored how a single robot decides where to look next to reduce its own uncertainty. The literature is mature for single-robot and small-team contexts. Fleet-scale active perception extends the same principle: the fleet decides where additional observations should focus to reduce collective uncertainty.

The architecture treats the fleet as a coordinated observation system. The forecasting engine's solicitation propagates as a credentialed observation; participating units consume the solicitation through their own admissibility framework; the response is a credentialed action that the originating engine consumes as updated observation. The cycle operates at fleet timescale rather than per-unit timescale.

Why Cooperative Perception Has Been Hard at Scale

Cooperative-perception literature has produced point solutions for specific contexts (V2X CPM for vehicular cooperative perception, multi-robot SLAM for warehouse robotics, sensor-network coordination for environmental monitoring). Each point solution has its own coordination mechanism, its own credentialing pattern, its own failure modes.

The architectural challenge is that cooperative perception spans many domains with the same fundamental coordination need. Smart-grid utilities cooperatively forecasting load. Weather services cooperatively observing emerging weather patterns. Autonomy fleets cooperatively perceiving traffic patterns. Defense ISR units cooperatively observing operating environments. Each currently uses its own integration; the architectural primitive provides what they share structurally.

How the Architectural Primitive Operates Across Domains

The forecasting engine, the credentialed observation framework, and the composite admissibility evaluation operate identically across domains. The configurations differ — what counts as 'forecast uncertainty' for grid load is different from what counts as 'forecast uncertainty' for traffic patterns — but the architectural mechanism is invariant.

Cross-domain cooperative perception becomes possible. A weather service's forecast-uncertainty solicitation can reach contributors across drone-fleet operators, airline operations, smart-grid utilities, and even individual smartphone-equipped citizens, each contributing under credentialed cross-recognition. The architectural primitive supports cooperative perception that current per-domain integration cannot.

What This Enables for Multi-Domain Cooperative Operations

The Mobileye REM model (fleet-contributed observation aggregation) gains a complementary primitive: forecast-driven solicitation that retasks observation capacity rather than just aggregating what fleets happen to observe. The combination — REM's aggregation plus solicitation-driven retasking — produces fleet-scale active perception that current REM cannot.

Smart-grid forecasting gains cross-utility solicitation, weather services gain cross-fleet observation contribution, and emerging multi-domain coordination scenarios (climate observation, urban-air-mobility coordination, multi-modal transport) gain the same architectural foundation. The patent positions the primitive at the layer where cooperative active perception scales beyond the single-robot research that produced the underlying principle.