Discrepancy Classification Taxonomy

by Nick Clark | Published April 25, 2026

What Each Classification Specifies

Nominal: sensed effect matches predicted effect within expected envelope. The action proceeded as designed; no further response needed.

Expected-noise: discrepancy within sensor and environmental noise floors. Recorded for trending but no immediate operational concern.

Anomaly: discrepancy exceeds noise floor but doesn't match a known fault or adversarial signature. The system flags for further evaluation; subsequent admissibility weights reduce until the cause is identified.

Fault: discrepancy matches a known fault signature (sensor failure, actuator failure, control loop malfunction). The system enters fault-handling mode; the unit's operational envelope adjusts accordingly.

Adversarial-interference: discrepancy matches an adversarial-attribution pattern from the disruption-modeling library. The system triggers expanded admissibility evaluation including cross-system alert broadcast.

Why Five Classes and Not Three

Smaller class sets collapse meaningfully different operational situations. 'Anomaly' alone treats all discrepancies as uniform candidates for evaluation, missing the structural distinction between random noise, cumulative anomaly, attributed fault, and adversarial action.

The five classes match the empirical structure of how autonomous systems actually fail. Random sensor noise is statistically recoverable; cumulative anomaly may indicate emerging failure; identified fault triggers different response than unknown anomaly; adversarial interference triggers governance-flagged cross-system response.

How Classifications Feed Back Through the Architecture

Each classification is itself a credentialed observation. Nominal classifications accumulate as evidence of correct operation. Expected-noise classifications support trending analysis. Anomaly classifications trigger investigation workflows. Fault classifications trigger maintenance and operational restriction. Adversarial-interference classifications trigger cross-system alerts.

The system's confidence-governed actuation evaluator consumes the classification stream as one of its inputs. Recent anomalies modulate admissibility weight; recent faults trigger operational mode adjustments; recent adversarial classifications trigger expanded admissibility envelopes.

What This Enables for Operational Awareness

The architecture produces operational awareness that uniform-discrepancy treatment cannot. The unit knows what kind of discrepancy it is experiencing, with credentialed evidence; downstream systems consume the credentialed classification through their own admissibility framework.

The patent positions the primitive at the layer where autonomous systems have been operating with reconstructed-rather-than-architectural discrepancy awareness.