

Article

Reliability Assessment and Adaptive Fusion Algorithm for Multi-Sensor Data in Autonomous Driving under Adverse Weather Conditions

Yi Guo ^{1,*}¹ Computer and Information Science, University of Pennsylvania, PA, USA

* Correspondence: Yi Guo, Computer and Information Science, University of Pennsylvania, PA, USA

Abstract: Atmospheric disturbances impose systematic degradation on multi-sensor perception systems in autonomous vehicles, necessitating a fundamental rethinking of sensor fusion strategies. This study presents a comprehensive reliability assessment framework combined with an adaptive fusion algorithm designed to mitigate sensor-specific performance deterioration under adverse meteorological conditions. Using an empirical dataset encompassing 10,000 hours of vehicle operation, we establish quantitative correlations between atmospheric parameters and measurement uncertainty across heterogeneous sensor modalities. Real-time trustworthiness estimation is achieved through a dynamic scoring mechanism that integrates environmental context with temporal performance evolution. The proposed adaptive fusion algorithm performs reliability-weighted integration through probabilistic decision modeling, optimizing sensor data combination while minimizing perception errors. Experimental validation demonstrates that the method improves object detection accuracy by 22.9 percentage points compared with majority-vote fusion (AP@0.5) and by 18.4 percentage points compared with fixed-weight fusion in heavy rain (AP@0.5). Furthermore, the false positive rate is reduced by 66-73% relative to the fixed-weight and majority-vote baselines, respectively. These improvements directly enhance collision avoidance performance, substantially advancing the safety and robustness of autonomous vehicles operating in challenging environmental conditions.

Keywords: multi-sensor fusion; reliability assessment; adverse weather; autonomous driving

Received: 22 October 2025

Revised: 29 October 2025

Accepted: 09 November 2025

Published: 15 November 2025



Copyright: © 2025 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

1.1. Background and Motivation for Multi-Sensor Fusion in Autonomous Driving

Autonomous vehicle perception systems integrate heterogeneous sensor modalities, each providing complementary information to environmental understanding pipelines [1]. Camera sensors operating at 30-60 Hz excel in texture discrimination and semantic classification through high-resolution imaging; however, depth estimation remains limited to stereoscopic configurations, with accuracy dependent on baseline separation. LiDAR systems generate three-dimensional point clouds at 10-20 Hz, achieving millimeter-level spatial precision for geometric reconstruction, though their performance is susceptible to atmospheric scattering during precipitation. Radar sensors maintain detection capabilities in dense precipitation while providing instantaneous velocity measurements via Doppler processing at 20 Hz, albeit with angular resolution constraints of 1-3 degrees.

Mathematical modeling of sensor fusion extends beyond simple aggregation, employing probabilistic state estimation where measurement uncertainties propagate through covariance matrices. Each sensor provides observations z_i with uncertainty characterized by Σ_i , requiring weighted combination strategies to minimize posterior uncertainty. Contemporary architectures employ Kalman filtering for linear systems with Gaussian noise, particle filters for non-linear measurement models using Monte Carlo sampling, and factor graph optimization for batch processing over temporal windows [2].

ISO 26262 defines hardware ASIL targets for functional safety, while ISO/PAS 21448 (SOTIF) does not mandate a specific failure rate. In this work, we target a failure probability of less than or equal to 10^{-8} to 10^{-9} per operational hour for safety-critical functions as an engineering goal rather than a normative requirement. These constraints influence fusion algorithm design, necessitating approaches that maintain operational integrity under sensor degradation or failure. Robust fusion algorithms are thus essential prerequisites for Level 4 and Level 5 autonomous vehicle capabilities.

1.2. Challenges of Sensor Performance Degradation under Adverse Weather Conditions

Atmospheric phenomena produce interference patterns that degrade sensor performance through absorption, scattering, and refraction [3]. Camera sensors experience contrast reduction following the Koschmieder visibility model, where the apparent luminance equals the object luminance multiplied by the exponential of negative extinction coefficient times propagation distance, plus the atmospheric luminance multiplied by one minus the same exponential term. Water droplet accumulation introduces geometric distortions, which can be analyzed via point spread functions; modulation transfer function values decline by 40-60% during rainfall exceeding 2 millimeters per hour.

LiDAR degradation follows Mie scattering, where the scattering cross-section depends on particle radius and wavelength through the size parameter defined as two times π times radius divided by wavelength [4]. Precipitation generates forward and backscatter, producing spurious points and range bias; refraction and multipath effects can distort apparent range, mitigated via outlier rejection and attenuation modeling. The attenuation coefficient can be empirically estimated, for example, k_{rain} equals 0.43 times rainfall rate to the power of 0.61 at 905 nanometers, where rainfall rate is expressed in millimeters per hour.

Electromagnetic propagation through precipitation is subject to frequency-dependent attenuation, where γ equals k times rainfall rate to the power of α , expressed in decibels per kilometer, with empirically calibrated coefficients. Multipath effects during precipitation generate ghost targets through ground reflection, causing bearing uncertainties up to approximately 5 degrees and range ambiguities exceeding 10 meters under severe conditions. Thermal expansion further affects calibration, where the angular error equals the thermal coefficient multiplied by the temperature difference, multiplied by structural span divided by mounting separation, inducing systematic drift.

1.3. Research Objectives and Contributions

This study introduces three key innovations to mitigate weather-induced perception degradation. First, weather-aware sensor performance modeling constructs empirical degradation functions via regression analysis of 10,000 hours of annotated sensor data. Piecewise linear approximations balance computational efficiency with modeling fidelity, maintaining correlation coefficients above 0.85 across weather categories while enabling real-time execution.

Second, dynamic reliability scoring applies Bayesian inference to quantify real-time trustworthiness, computing posterior distributions of reliability given measurements and weather conditions, where reliability represents the sensor state, measurements are z , and environmental conditions are w . Operating at 20 Hz, the algorithm integrates

environmental context using 30-60 second temporal smoothing for stability against transients, while 5-15 second response windows capture genuine transitions.

Third, adaptive fusion formulates weight allocation as a constrained optimization problem: minimize the sum over i of weight_i times the sensor uncertainty squared, subject to the sum of weights equal to one, and each weight greater than or equal to a minimum threshold. Here, sensor uncertainty quantifies the individual sensor's reliability, while the minimum weight prevents complete exclusion. The framework demonstrates agility, adapting within 2-second windows while maintaining $O(n^2)$ computational complexity for an n -sensor configuration.

2. Related Work and Technical Challenges

2.1. Current Multi-Sensor Fusion Approaches in Autonomous Vehicles

Contemporary fusion architectures implement hierarchical processing by combining measurements at multiple abstraction levels [5]. Early fusion concatenates raw sensor data into joint feature spaces and applies convolutional networks with three-dimensional kernels for simultaneous processing of images and LiDAR voxels. These architectures achieve approximately 8.3% higher detection accuracy compared with single-modality baselines but require around 4.2 times more computational resources for operation at 10 Hz.

Late fusion processes sensor streams through independent pipelines before combining outputs using non-maximum suppression and probabilistic association [6]. Modular architectures enable parallel GPU processing, achieving 30 Hz operation with 250 milliseconds latency. Detection-level fusion typically employs the Hungarian algorithm for bipartite matching, optimizing assignments based on spatial proximity and appearance similarity.

Transformer-based networks leverage attention mechanisms to learn sensor importance through self-supervised training on extensive driving datasets. Attention modules compute cross-modal correlations by performing dot-product operations between encoded features, producing importance maps that highlight complementary regions. These approaches demonstrate approximately 15.8% improvement in detection recall compared with manually designed weighting strategies, albeit with a computational cost of around 500 GFLOPS.

Probabilistic graphical models represent measurements as factor graphs, connecting state variables through measurement factors and motion constraints. Optimization minimizes the negative logarithm of the posterior probability of the state given the measurements, typically converging within ten iterations for urban scenarios. Factor graphs can naturally handle asynchronous measurements and missing data through marginalization.

2.2. Weather-Induced Sensor Performance Limitations and Failure Modes

Empirical studies reveal sensor-specific vulnerability profiles under adverse weather [7]. Camera detection range exhibits exponential deterioration, calculated as the clear-weather range multiplied by the exponential of negative attenuation coefficient times precipitation rate. For example, with an attenuation coefficient of 0.15, classification accuracy can drop from a baseline of 92% to 61% during heavy precipitation exceeding 5 millimeters per hour.

LiDAR point cloud degradation can be modeled as the clear-weather number of returns multiplied by one minus the per-meter scattering probability raised to the propagation distance. Empirical campaigns show a 73 percentage-point reduction at 50 meters during 10 millimeters per hour rainfall, with false positive rates increasing from 2.1% to 18.7% due to precipitation backscatter. Missing points concentrate in forward-facing sectors, creating directional blind spots [8].

Radar resolution deteriorates from the nominal 2.5 degrees to 4.1 degrees beam width in dense fog due to refraction and multipath effects. Thermal cycling induces approximately 0.3 degrees angular error per 10-degree Celsius temperature change. Sub-

zero conditions amplify vibration magnitudes by 35%, introducing high-frequency noise exceeding 50 Hz, which can corrupt inertial measurements. Thermomechanical coupling generates differential stress patterns that manifest as systematic biases.

2.3. Existing Reliability Assessment Methods and Their Limitations

Statistical frameworks use chi-squared testing to validate measurement consistency against nominal distributions [9]. The test statistic is calculated as the sum over all observations of the squared difference between measurement and mean, divided by the variance. This statistic follows a chi-squared distribution, enabling hypothesis evaluation with 95% confidence. Adaptive outlier rejection adjusts sensitivity from three standard deviation limits under favorable conditions to five standard deviation limits during disturbances.

Machine learning methods employ random forest models with 100 decision trees, achieving approximately 94.2% classification accuracy on annotated failure datasets [10]. Feature engineering extracts 47 statistical descriptors, including temporal autocorrelation, spectral entropy, and cross-sensor correlation over one-second windows.

Information-theoretic approaches estimate reliability using mutual information normalized by measurement entropy, where sensors with high mutual information are assigned increased weights, while random behavior results in reliability approaching zero [11]. Computation requires approximately 50 milliseconds for urban scenes containing 100 objects.

Existing frameworks generally lack explicit weather modeling, treating sensor degradation uniformly regardless of the underlying cause. This limitation prevents targeted mitigation and often results in conservative rejection, unnecessarily reducing system capability. Incorporating weather-aware reliability assessment allows retention of partially degraded sensors by appropriately modeling their uncertainty.

3. Multi-Sensor Reliability Assessment Framework

3.1. Weather-Aware Sensor Performance Modeling and Degradation Analysis

Sensor performance modeling establishes degradation functions through regression analysis of annotated measurements across atmospheric conditions. Empirical coefficients $\alpha_{\text{sensor}(w)}$ relate the weather state w to sensor performance, where α is in the range [0,1] representing capability retention. Piecewise linear interpolation between discrete categories enables continuous estimation while maintaining computational efficiency.

Camera degradation incorporates visibility V_{met} to compute extinction using $\sigma_{\text{ext}} = 3.912 / V_{\text{met}}$ following Koschmieder's law. Contrast reduction is quantified as $C_{\text{r}} = \exp(-\sigma_{\text{ext}} * d)$ at distance d , establishing the detection probability $P_{\text{detect}} = P_{\text{clear}} * C_{\text{r}} * \alpha_{\text{precipitation}}$, where $\alpha_{\text{precipitation}}$ accounts for lens contamination. Validation across 5,000 images demonstrates a correlation coefficient $r = 0.87$ between predicted and observed detection rates.

LiDAR models implement modified Mie theory for polydisperse distributions. Backscatter is computed as $\beta(r) = \sum (n_i * \sigma_i)$, aggregating contributions from particles with density n_i and cross-section σ_i . Range-dependent attenuation follows $A(r) = \exp(-2 \int \beta(r') dr')$, creating non-linear patterns varying with atmospheric stratification. Validation against 3,000 events confirms accuracy within 15% for rainfall rates below 20 mm/hr.

Radar modeling follows ITU-R recommendations for 77 GHz attenuation. Two-way attenuation is calculated as $L = 2 \int \gamma(r) dr$, where $\gamma = 0.38 * R^{0.72}$ for horizontal polarization. Doppler broadening quantifies velocity degradation as $\sigma_v = 0.13 * R^{0.5}$ [12].

As shown in Table 1, sensor performance degradation coefficients are provided for different weather conditions.

Table 1. Sensor Performance Degradation Coefficients Under Different Weather Conditions.

Weather Condition	Visibility (m)	Camera Coefficient	LiDAR Coefficient	Radar Coefficient	Temperature Factor
Clear	>10000	0.95	0.92	0.88	1.00
Light Rain 0.5-2mm/h	5000 - 8000	0.78	0.75	0.85	0.98
Moderate Rain (2 - 5mm/h)	2000 - 5000	0.52	0.45	0.78	0.96
Heavy Rain (>5mm/h)	500 - 2000	0.31	0.23	0.71	0.94
Light Fog	1000 - 3000	0.65	0.82	0.83	0.97
Dense Fog	50 - 200	0.15	0.68	0.76	0.95
Light Snow	800 - 2000	0.58	0.38	0.74	0.89
Heavy Snow	100 - 500	0.22	0.12	0.62	0.85
Freezing Rain	300 - 800	0.28	0.18	0.68	0.82
Mixed Precipitation	200 - 600	0.35	0.25	0.65	0.88
Night + Light Rain	1000 - 3000	0.42	0.71	0.82	0.96
Night + Fog	50 - 500	0.18	0.64	0.73	0.94

Temperature compensation is modeled as $T_factor = 1 - 0.02 * |T - 20|$, where 20°C is the reference temperature. The composite coefficient $\alpha_total = \alpha_weather * T_factor * \alpha_mechanical$ combines atmospheric, thermal, and mechanical factors. Cross-validation demonstrates a mean absolute error of 0.08 for coefficient prediction.

3.2. Real-Time Reliability Estimation Algorithm for Camera, LiDAR and Radar

Environmental assessment employs computer vision for atmospheric characterization. Precipitation detection isolates rain streaks through temporal differencing, while fog quantification analyzes Fourier spectra to characterize scattering intensity. These methods enable robust classification across meteorological conditions.

Measurement validation uses the Mahalanobis distance:

$$d_M = \sqrt{(z - \mu)^T \Sigma^{-1} (z - \mu)}$$

for outlier identification, where μ and Σ encode Kalman filter predictions. Observations violating $d_M > \chi^2_alpha(n)$ trigger reliability penalties, with $\chi^2_alpha(n)$ representing critical values at significance α for n dimensions. Adaptive thresholds transition from $\alpha = 0.01$ under benign conditions to $\alpha = 0.05-0.10$ under adverse conditions, balancing sensitivity against false alarms.

Temporal consistency is evaluated through sliding windows computing the coefficient of variation:

$$CV = \sigma / \mu$$

Stable sensors maintain $CV < 0.15$, while degraded sensors manifest $CV > 0.35$. Intermediate values trigger proportional adjustments. Spectral analysis identifies periodic disturbances via FFT to detect mechanical resonances [13].

As shown in Figure 1, the real-time sensor reliability estimation architecture integrates these modules.




Figure 1. Real-Time Sensor Reliability Estimation Architecture.

Cross-sensor correlation is quantified as:

$$q_{ij} = \text{cov}(z_i, z_j) / (\sigma_i * \sigma_j)$$

Strong correlation ($q > 0.7$) confirms sensor veracity, while deficiency signals malfunction or occlusion. Matrix computation at 5 Hz enables fault detection within 200 ms.

Historical tracking uses exponentially weighted averaging:

$$R_{\text{avg}} = 0.1 * R_{\text{current}} + 0.9 * R_{\text{previous}}$$

Linear regression over 60-second windows identifies degradation trends, triggering alerts when rates exceed 0.01 per minute. Performance history is maintained for 24 hours at 1 Hz granularity.

As shown in Table 2, reliability scoring parameters and thresholds are defined for each sensor.

Table 2. Real-Time Reliability Scoring Parameters and Thresholds.

Assessment Criteria	Parameter	Camera	LiDAR	Radar	Weight Factor
Environmental Impact	Precipitation Weight	0.65	0.58	0.32	0.40
Environmental Impact	Fog Weight	0.72	0.45	0.28	0.35
Environmental Impact	Temperature Weight	0.25	0.30	0.15	0.15
Environmental Impact	Wind Weight	0.18	0.22	0.12	0.10
Measurement Consistency	Spatial Accuracy ($\pm m$)	0.8	0.3	0.5	0.30
Measurement Consistency	Temporal Stability (CV)	0.15	0.12	0.18	0.25
Measurement Consistency	Inter-sensor Agreement	0.75	0.80	0.70	0.20

Measurement Consistency	Detection Range (m)	150	200	300	0.25
Historical Performance	Short - term Weight (180s)	0.60	0.55	0.65	0.60
Historical Performance	Long - term Weight (3600s)	0.40	0.45	0.35	0.40
Historical Performance	Degradation Threshold	0.20	0.25	0.15	-
Historical Performance	Recovery Time (s)	45	60	30	-
Quality Thresholds	Minimum Reliability	0.15	0.20	0.25	-
Quality Thresholds	Warning Level	0.35	0.40	0.45	-
Quality Thresholds	Optimal Level	0.85	0.80	0.75	-

Computational optimization achieves 8 ms latency via SIMD vectorization and GPU-accelerated inference. Memory bandwidth is optimized through caching and prefetching. The modular architecture allows selective computation based on sensor availability.

3.3. Dynamic Reliability Scoring Mechanism Based on Environmental Conditions

Hierarchical Bayesian inference is formulated as:

$$P(R | z, w, h)$$

where R represents reliability, z the observations, w the atmospheric state, and h the historical context. The prior $P(R | w, h)$ integrates weather degradation with long-term trends, while the likelihood $P(z | R)$ evaluates consistency. Posterior approximation uses variational inference, converging in 5 iterations via coordinate descent.

Short-term reliability is aggregated as:

$$R_{\text{short}} = \sum (w_i * r_i)$$

with r_i in $[0,1]$ and normalized weights w_i summing to 1. Weight determination minimizes entropy:

$$\sum (w_i * \log(w_i))$$

subject to performance constraints, promoting balanced utilization. Interior point methods achieve convergence within 2 ms for 3-sensor configurations.

Long-term evolution is modeled as a state-space system:

$$R_t = f(R_{t-1}, u_t) + \epsilon_t, \epsilon_t \sim N(0, Q)$$

$$z_t = h(R_t) + \eta_t, \eta_t \sim N(0, R)$$

Extended Kalman filtering provides recursive estimation with $O(n^2)$ complexity.

As shown in Figure 2, the evolution of dynamic reliability scores during weather transitions is illustrated.




Figure 2. Dynamic Reliability Score Evolution During Weather Transition Events.

Spatial mapping divides the field-of-view into angular θ_i and range r_j cells, computing $R(\theta_i, r_j)$ for each. Weather effects are non-uniform, with forward sectors experiencing up to 40% greater degradation during frontal precipitation. This enables selective utilization, extracting maximum information while suppressing unreliable regions.

The geometric mean:

$$(\prod R_i)^{1/n}$$

achieves optimal balance, maintaining sensitivity to degradation while preventing catastrophic collapse, outperforming arithmetic averaging and minimum operators.

As shown in Table 3, reliability score integration weights are defined for different operational scenarios.

Table 3. Reliability Score Integration Weights for Different Operational Scenarios.

Operational Scenario	Speed (km/h)	Camera Weight	LiDAR Weight	Radar Weight	Safety Factor
Urban Driving	20 - 50	0.45	0.35	0.20	1.2
Suburban Driving	40 - 70	0.38	0.32	0.30	1.1
Highway Driving	80 - 120	0.25	0.25	0.50	1.0
Construction Zone	10 - 30	0.25	0.55	0.20	1.5
Parking Operations	0 - 10	0.33	0.33	0.34	1.8
School Zone	15 - 25	0.50	0.30	0.20	1.6
Night Urban	20 - 50	0.30	0.45	0.25	1.3
Night Highway	80 - 120	0.20	0.30	0.50	1.2

Rain					
Modifier	-	0.6	0.7	1.2	-
Fog					
Modifier	-	0.4	0.8	1.1	-
Snow					
Modifier	-	0.5	0.6	1.3	-
Clear					
Weather	-	1.0	1.0	1.0	-
Modifier					

Operational context modulates thresholds through risk-aware boundaries. Dense urban environments with pedestrian concentrations exceeding 10 per km² mandate 1.5× safety amplification. Highway scenarios prioritize long-range detection, redistributing weights toward distant bins.

4. Adaptive Fusion Algorithm Design and Implementation

4.1. Weight Allocation Strategy Based on Sensor Reliability Scores

Weight allocation is formulated as a quadratic optimization problem, determining $w = [w_{\text{camera}}, w_{\text{LiDAR}}, w_{\text{radar}}]$ to minimize

$J(w) = w^T \Sigma w$

where Σ encodes reliability-scaled covariance. Constraints enforce $\Sigma w_i = 1$ and $w_i \geq R_i * w_{\text{min}}$ to ensure proportional contribution.

Sequential quadratic programming uses gradient $\nabla J = 2 * \Sigma * w$ and Hessian $H = 2 * \Sigma$, achieving convergence within 3-5 iterations. Warm-starting reduces computation to 2 iterations for frame transitions. Cholesky decomposition exploits matrix structure, requiring 0.8 ms on embedded processors. Active set methods maintain feasibility throughout iterations.

Uncertainty propagation employs the unscented transform for nonlinear models. Sigma points are defined as $X_i = \mu \pm \sqrt{((n + \kappa) * \Sigma)}$, where n is the dimension and $\kappa = 3 - n$ provides accuracy. The transformed covariance is computed as $P_y = \Sigma (W_i * (Y_i - \mu_y) * (Y_i - \mu_y)^T)$, aggregating contributions and generating uncertainty ellipsoids.

Cross-validation across 1,000 scenarios confirms an 18.3% RMSE reduction under adverse weather. Metrics include position RMSE, classification accuracy, and track consistency, with validation emphasizing edge cases where sensor degradation deviates from typical patterns [14].

As shown in Figure 3, the adaptive weight allocation response to environmental changes is illustrated.




Figure 3. Adaptive Weight Allocation Response to Environmental Changes.

Temporal adaptation follows

$$w_t = 0.3 * w_{\text{optimal}} + 0.7 * w_{t-1}$$

balancing responsiveness against stability. Abrupt transitions trigger verification against safety constraints, while monitoring prevents pathological dependencies by enforcing $w_i \leq 0.75$ upper bounds.

As shown in Table 4, sensor weight allocation matrices are provided for weather-specific scenarios.

Table 4. Sensor Weight Allocation Matrix for Weather-Specific Scenarios.

Weather Scenario	Precipitation (mm/h)	Visibility (m)	Camera Weight	LiDAR Weight	Radar Weight	Confidence Level
Clear Day	0	>10000	0.40	0.35	0.25	0.95
Clear Night	0	>5000	0.28	0.42	0.30	0.88
Light Rain Day	0.5 - 2.0	3000 - 8000	0.30	0.28	0.42	0.82
Light Rain Night	0.5 - 2.0	1500 - 4000	0.22	0.25	0.53	0.76
Heavy Rain Day	>5.0	500 - 2000	0.15	0.20	0.65	0.71
Heavy Rain Night	>5.0	200 - 1000	0.08	0.20	0.72	0.65
Light Fog	0	1000 - 3000	0.25	0.35	0.40	0.80
Dense Fog	0	50 - 200	0.10	0.32	0.58	0.68
Light Snow	1.0 - 3.0	800 - 2000	0.20	0.25	0.55	0.73
Heavy Snow	>5.0	100 - 500	0.12	0.18	0.70	0.62
Freezing Rain	0.5 - 2.0	300 - 800	0.18	0.22	0.60	0.58
Mixed Precipitation	2.0 - 5.0	200 - 600	0.16	0.24	0.60	0.64

Implementation employs lookup tables for common scenarios, SIMD vectorization for matrix operations, and pipeline parallelization across channels. Memory footprint remains below 10 MB including precomputed tables.

4.2. Environmental Context-Aware Fusion Decision Making

Context-aware decision making integrates meteorological measurements, traffic estimation, and map priors into strategy selection. Hierarchical state machines with $S_{\text{weather}} \in \{\text{clear, light, moderate, severe}\}$ and $S_{\text{operation}} \in \{\text{parking, urban, highway, emergency}\}$ trigger parameter updates, adapting thresholds, gates, and tracking parameters.

Weather classification processes precipitation rate, visibility, and temperature/humidity through fuzzy membership functions $\mu_i(x)$, enabling gradual transitions. Defuzzification via the centroid provides a continuous severity score $s \in [0,1]$, allowing smooth interpolation.

Traffic density is analyzed using radar and camera detections, computing vehicles per area and categorizing as sparse ($<10/\text{km}^2$), moderate ($10\text{--}50/\text{km}^2$), or dense ($>50/\text{km}^2$). Higher densities bias the system toward conservative strategies, increasing confidence thresholds from 0.6 to 0.85, with initialization requiring 3-5 detections.

Map priors incorporate HD annotations for topology, lanes, and patterns. Prior probabilities $P(\text{object}|\text{location})$ weight detections according to distribution. Intersections increase pedestrian priors 3×, while highways suppress pedestrian hypotheses. Detection integration follows:

$$P(\text{object}|\text{detection}, \text{location}) \propto P(\text{detection}|\text{object}) * P(\text{object}|\text{location})$$

Predictive adaptation anticipates changes using weather radar forecasts for 15 minutes, enabling preemptive adjustments. Route analysis identifies challenging segments, triggering specialized configurations. This framework reduces adaptation latency from 2 s to 0.5 s.

Safety envelope verification ensures risk-aware decisions, using formal verification:

$$\varphi = G(\text{distance_min} > d_{\text{safe}})$$

Reachability analysis at 1 Hz validates configurations against worst-case degradation. Failed verification triggers a conservative fallback prioritizing collision avoidance [15].

4.3. Robust Data Association and Conflict Resolution Mechanisms

Data association establishes correspondences through:

$$L(A) = \prod P(z_i|x_j) * P(A)$$

where A represents the assignment. The Hungarian algorithm solves $\max_A L(A)$ with $O(n^3)$ complexity, while multiple hypothesis tracking maintains K -best solutions for ambiguous cases.

Spatial gates use:

$$d^2_M = (z - Hx)^T S^{-1} (z - Hx) < \gamma$$

where $S = H * P * H^T + R$, and γ denotes thresholds adapting from $\gamma = 9.21$ (99% confidence for 2D) in clear conditions to $\gamma = 16.27$ (99.9%) in adverse weather.

Temporal association enforces consistency:

$$\text{LLR}_t = \text{LLR}_{t-1} + \log(P_D * P(z|x) / P_{FA})$$

where P_D denotes detection probability and P_{FA} is the false alarm rate. Tracks exceeding $T_{\text{confirm}} \in [3, 8]$ are confirmed, while scores below $T_{\text{delete}} \in [-5, -2]$ are terminated.

Conflict resolution aggregates estimates through:

$$x_{\text{consensus}} = \sum (w_i * x_i) / \sum w_i$$

weighted by reliability and precision. RANSAC with iterations $N = \log(1-0.99) / \log(1-w^m)$ maintains robustness against failures and interference.

False positive suppression verifies geometric consistency, including object dimensions, ground alignment, and occlusions. Temporal filters require $N_{\text{min}} \in [2, 5]$ detections within windows. Semantic validation rejects implausible detections based on scene context.

Measurement fusion employs covariance intersection:

$$P^{-1} = \omega * P_1^{-1} + (1-\omega) * P_2^{-1}$$

$$x = P * (\omega * P_1^{-1} * x_1 + (1-\omega) * P_2^{-1} * x_2)$$

where ω minimizes $\text{trace}(P)$, providing conservative estimates without knowledge of correlations.

5. Experimental Evaluation and Performance Analysis

5.1. Simulation Setup and Real-World Dataset Description

Validation encompasses both simulation and real-world datasets totaling 10,000 operational hours. Physics-based simulations achieve correlation coefficients of 0.92 for camera, 0.89 for LiDAR, and 0.94 for radar. Weather simulation utilizes particle systems for precipitation, volumetric ray marching for fog, and physically-based lighting for diurnal variations.

Field data collection deployed five vehicles across diverse regions: Pacific Northwest (3,500 hours, predominantly in precipitation), Upper Midwest (2,800 hours, snow-prone conditions), and California Central Valley (3,700 hours, fog-prone regions). The aggregated weather distribution across all areas includes 5,600 hours of precipitation (56%), 3,700 hours of snow (37%), and 700 hours of fog/mixed (7%). Sensor configurations consist of six cameras (1920×1080 at 30 Hz), two LiDAR units (128 channels at 10 Hz), and five radars (77 GHz at 20 Hz). Ground truth is provided via differential GPS with RTK, achieving 2 cm accuracy and synchronized within 1 ms through PTP.

The weather distribution further details clear skies (4,200 hours, 42%), light precipitation (2,300 hours, 23%), moderate precipitation (1,800 hours, 18%), heavy precipitation (1,000 hours, 10%), and fog/mixed conditions (700 hours, 7%). Precipitation rates range from 0.1 to 50 mm/hr, with visibility spanning 10 m to unlimited. Temperatures vary from -25°C to 45°C.

Dataset annotation includes 8.7 million vehicles, 2.3 million pedestrians, and 1.1 million cyclists, each with 3D bounding boxes and tracking. Weather labels are derived from meteorological stations within 5 km, calibrated against on-vehicle sensors. Semi-automated pipelines achieve 96.3% inter-annotator agreement for objects and 91.7% for weather labels.

5.2. Comparative Analysis with Existing Fusion Methods under Various Weather Scenarios

Evaluation compares the adaptive fusion method against fixed-weight, majority voting, Kalman filter, and learned fusion baselines. Metrics include detection accuracy, false positives, tracking consistency, and computational efficiency across different weather conditions.

Average precision at IoU 0.5 demonstrates notable improvements during adverse weather. In clear weather, performance remains comparable (proposed: 89.3%, fixed-weight: 88.1%, majority-vote: 86.7%, Kalman: 87.9%, learned: 88.5%). Under heavy rain, the proposed method achieves AP@0.5 of 71.2%, compared to 52.8% for fixed-weight and 48.3% for majority-vote, corresponding to absolute improvements of +18.4 and +22.9 percentage points, respectively, enabled by weather-aware adaptation.

False positive rates during precipitation are significantly reduced. The proposed method maintains 8.3 false positives per kilometer during heavy rain, compared to 24.7 for fixed-weight, 31.2 for majority voting, 22.1 for Kalman, and 15.8 for learned fusion. This 73% reduction relative to majority voting translates to fewer phantom braking events. Statistical significance is confirmed with $p < 0.001$ for adverse weather differences.

MOTA scores demonstrate system robustness: the proposed method achieves 76.4% during moderate rain versus 61.2% for fixed-weight, representing a 24.8% improvement. Track fragmentation is reduced by 41% through adaptive selection, preserving continuity during sensor failures. Identity switches decrease from 4.7/minute to 1.8/minute in challenging scenarios.

Computational analysis indicates 18.3 ms latency on an NVIDIA Jetson AGX Xavier. Adaptive fusion adds only 3.7 ms overhead compared to the fixed-weight baseline. Memory utilization remains below 850 MB, while power consumption increases by 2.3 W during adverse weather for reliability assessment tasks.

References

1. C. Xiang, C. Feng, X. Xie, B. Shi, H. Lu, Y. Lv, and Z. Niu, "Multi-sensor fusion and cooperative perception for autonomous driving: A review," *IEEE Intelligent Transportation Systems Magazine*, vol. 15, no. 5, pp. 36-58, 2023. doi: 10.1109/imits.2023.3283864
2. A. Abdulkasoud, and R. Ahmed, "Transformer-based sensor fusion for autonomous vehicles: A comprehensive review," *IEEE Access*, 2025. doi: 10.1109/access.2025.3545032
3. L. Han, J. Wang, C. Li, F. Tao, and Z. Fu, "A novel multi-object tracking framework based on multi-sensor data fusion for autonomous driving in adverse weather environments," *IEEE Sensors Journal*, 2025.
4. Y. Zhang, C. Tu, K. Gao, and L. Wang, "Multisensor information fusion: Future of environmental perception in intelligent vehicles," *Journal of Intelligent and Connected Vehicles*, 2024. doi: 10.26599/jicv.2023.9210049

5. F. A. Butt, J. N. Chattha, J. Ahmad, M. U. Zia, M. Rizwan, and I. H. Naqvi, "On the integration of enabling wireless technologies and sensor fusion for next-generation connected and autonomous vehicles," *IEEE Access*, vol. 10, pp. 14643-14668, 2022. doi: 10.1109/access.2022.3145972
6. A. S. Alam, X. Hei, and Y. Zhang, "In-progress: Enhancing traffic signal perception for connected and autonomous vehicles (CAVs) via multi-sensor fusion of camera, LiDAR, radar, and SPaT data," In *2025 IEEE Security and Privacy Workshops (SPW)*, May, 2025, pp. 361-363.
7. Y. Qiu, Y. Lu, Y. Wang, and C. Yang, "Visual perception challenges in adverse weather for autonomous vehicles: A review of rain and fog impacts," In *2024 IEEE 7th Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*, September, 2024, pp. 1342-1348. doi: 10.1109/itnec60942.2024.10733168
8. I. P. P. A. Sumalatha, P. Chaturvedi, S. Patil, H. P. Thethi, and A. A. Hameed, "Autonomous multi-sensor fusion techniques for environmental perception in self-driving vehicles," In *2024 International Conference on Communication, Computer Sciences and Engineering (IC3SE)*, May, 2024, pp. 1146-1151.
9. K. Huang, B. Shi, X. Li, X. Li, S. Huang, and Y. Li, "Multi-modal sensor fusion for auto driving perception: A survey," *arXiv preprint arXiv:2202.02703*, 2022.
10. D. Khan, S. Aslam, and K. Chang, "Vehicle-to-infrastructure multi-sensor fusion (V2I-MSF) with reinforcement learning framework for enhancing autonomous vehicle perception," *IEEE Access*, 2025. doi: 10.1109/access.2025.3551367
11. Z. Shao, H. Wang, Y. Cai, L. Chen, and Y. Li, "UA-Fusion: Uncertainty-aware multimodal data fusion framework for 3D object detection of autonomous vehicles," *IEEE Transactions on Instrumentation and Measurement*, 2025.
12. H. Su, H. Gao, X. Wang, X. Fang, Q. Liu, G. Huang, and Q. Cao, "Object detection in adverse weather for autonomous vehicles based on sensor fusion and incremental learning," *IEEE Transactions on Instrumentation and Measurement*, 2024. doi: 10.1109/tim.2024.3472860
13. M. Nawaz, S. Khan, M. Daud, M. Asim, G. A. Anwar, A. R. Shahid, and W. Yuan, "Improving autonomous vehicle cognitive robustness in extreme weather with deep learning and thermal camera fusion," *IEEE Open Journal of Vehicular Technology*, 2025. doi: 10.1109/ojvt.2025.3529495
14. D. J. Yeong, G. Velasco-Hernandez, J. Barry, and J. Walsh, "Sensor and sensor fusion technology in autonomous vehicles: A review," *Sensors*, vol. 21, no. 6, p. 2140, 2021. doi: 10.3390/s21062140
15. E. Ghiasi, G. Ghajari, M. Gottipati, P. S. S. Gogineni, R. Galla, and N. Bourbakis, "Association of multi-sensor data for autonomous car driving: A comparative evaluation," In *2024 IEEE 36th International Conference on Tools with Artificial Intelligence (ICTAI)*, October, 2024, pp. 427-436.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of the publisher and/or the editor(s). The publisher and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.