Proceedings of CONF-MSS 2026 Symposium: Mechanical Control and Automation
DOI: 10.54254/2755-2721/2026.BJ32383

Robust Indoor Localization for Autonomous Mobile Robots: A Review of Bayesian Sensor Fusion and SLAM Algorithms

Yilin Sun

Chang'an Dublin International College of Transportation, Chang'an University, Xi'an, China
english_is_popular@163.com

Abstract. Indoor autonomous mobile robots (AMRs) must estimate pose without GNSS while facing occlusions, perceptual aliasing, wheel slip, and dynamic obstacles. These conditions make single-sensor localization brittle and motivate probabilistic fusion of motion cues with environment-referenced and infrastructure observations. This algorithm-focused review synthesizes robust indoor localization for AMRs by linking sensing models to estimation theory and engineering trade-offs. This paper is organized around measurement models of common sensors, Bayesian filtering (Kalman and particle) and smoothing via factor-graph optimization, and SLAM pipelines that combine front-end association with back-end optimization and loop closure. Learning-augmented components are discussed as targeted add-ons that improve perception, uncertainty modeling, and outlier rejection. Finally, practical evaluation guidelines are summarized, including recovery behavior and computational metrics that matter in deployment. Open directions include life-long mapping in changing buildings and reliability-aware autonomy regulation. The emphasis is on selection criteria that transfer to real deployments.

Keywords: Indoor Localization, Autonomous Mobile Robots, Bayesian Estimation, Sensor Fusion, SLAM

1. Introduction

Robust indoor localization is a prerequisite for safe navigation, task execution, and human-aware operation of autonomous mobile robots (AMRs). Warehouses, hospitals, and retail/service venues increasingly rely on AMRs to transport goods and interact with people, which makes localization failures costly in both safety risk and downtime. Indoor platforms cannot rely on stable GNSS and must handle perceptual degradation: repetitive corridors and shelving cause aliasing, lighting variation challenges cameras, and dense pedestrian traffic introduces non-stationary occlusions. At the same time, wheel slip and payload changes can rapidly accumulate odometry drift. Modern systems therefore combine complementary sensors and integrate them with probabilistic estimators that reconcile heterogeneous measurements while quantifying uncertainty [1,2].

Localization failures still limit deployment because algorithm choice depends on environment structure, sensor cost, compute budget, and the required recovery behavior after disruptions. A useful review should compare estimation families by their assumptions (linearity, Gaussianity),

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.BJ32383

robustness to outliers, and how naturally they integrate mapping and loop closure. These criteria matter more than small accuracy differences in controlled benchmarks.

This paper surveys recent research on indoor AMR localization with an emphasis on algorithmic foundations. The methodology is literature synthesis: key surveys and representative application studies are grouped into a unified taxonomy of measurement models, Bayesian filtering, graph-based smoothing, and SLAM, and then compared by trade-offs and deployment implications. The goal is to support principled design decisions for real indoor systems rather than to recommend a single "best" algorithm.

2. Scope and taxonomy

2.1. Problem setting

The objective is online estimation of the robot pose x_t (typically SE(2) for planar motion) given a motion model driven by control/odometry inputs and a stream of sensor observations. In localization-only settings, a map is assumed to be known; in SLAM, the map is estimated jointly. The focus here is on indoor ground AMRs in logistics and service scenarios, where recovery from temporary sensor degradation is critical [1,2].

2.2. Algorithmic taxonomy

Indoor localization algorithms can be classified by the representation of the posterior belief over x_t. Filtering methods maintain p(x_t | z_1:t) recursively (Kalman-family or particle filters), while smoothing methods optimize a window of states on a sparse factor graph. SLAM extends smoothing by introducing map variables and loop-closure constraints. Learning is most often used to improve perception or uncertainty modeling inside these estimators rather than to replace them entirely [1].

2.3. Review structure

This review is organized from sensing foundations to system-level deployment. Section 3 examines the main sensor modalities used in indoor AMR localization and their corresponding measurement models. Section 4 then reviews Bayesian sensor-fusion approaches, including Kalman-family filters, particle filters, and graph-based smoothing, with emphasis on their assumptions, advantages, and limitations in indoor environments. Building on these estimation methods, Section 5 surveys representative SLAM pipelines, covering front-end data association, back-end optimization, loop closure, and multi-sensor integration. Sections 6 and 7 shift from algorithmic design to practical use by discussing implementation issues and evaluation criteria for real-world deployment. Finally, Section 8 outlines the major open challenges and emerging research directions in robust indoor localization.

3. Sensor and measurement models

3.1. Proprioceptive motion sensing

Wheel encoders and inertial measurement units (IMUs) provide high-rate motion cues that support smooth short-term tracking. Encoder integration accumulates bias under wheel slip and unmodeled kinematics, while IMU estimates drift due to bias and scale factors. In Bayesian estimators, these

DOI: 10.54254/2755-2721/2026.BJ32383

sensors mainly define the prediction step, so credible process-noise modeling is essential for long-horizon stability.

3.2. Exteroceptive perception sensing

LiDAR and cameras provide environment-referenced observations that correct drift by matching scans or visual features against a map or recent measurements. LiDAR is comparatively invariant to lighting and yields stable geometric constraints, but scan matching can degenerate in corridors or open areas. Cameras provide rich semantics and association cues yet are sensitive to illumination, blur, and texture scarcity. Treating LiDAR and vision as complementary is therefore common in robust indoor designs [1,3,4].

3.3. Infrastructure-based absolute aids

Infrastructure signals can provide absolute constraints when global consistency or rapid recovery is needed. UWB ranging can anchor position and is often fused with motion estimates to improve robustness [5]. BLE beacons can support coarse localization or floor identification in multi-story buildings [6], while RFID tags offer sparse but high-confidence landmarks [7]. These aids trade installation and maintenance costs for improved observability.

4. Bayesian sensor fusion algorithms

Most estimators follow the Bayes filter recursion: predict with the motion model and update with a measurement likelihood. Algorithm choice determines how the posterior is represented (Gaussian, samples, or optimization) and how nonlinearity, non-Gaussian noise, and outliers are handled. Table 1 shows the comparison of sensor fusion method families for indoor AMR localization.

Table 1. Comparison of sensor fusion method families for indoor AMR localization

Method family	Typical inputs	Core assumptions	Strengths	Limitations	Computational cost	Best-fit use cases
Kalman filter (KF)	Wheel odometry; IMU; simple linear observations	Linear models; Gaussian noise	Very fast; closed-form update; easy to implement	Poor fit for strong nonlinearity; sensitive to wrong noise tuning	Low	Short-horizon tracking with well-behaved models
Extended Kalman filter (EKF)	Odometry; IMU; range/bearing (UWB, LiDAR features)	Local linearization valid; near-Gaussian errors; good initialization	Real-time; handles moderate nonlinearity; widely used	Linearization error; inconsistency under high nonlinearity/outliers	Low–Medium	Tracking with occasional absolute constraints
Unscented Kalman filter (UKF)	Odometry; IMU; nonlinear observations	Sigma-point approximation; near-Gaussian noise	Better nonlinear propagation than EKF; no Jacobians	Still Gaussian assumption; higher compute than EKF	Medium	Nonlinear systems with unstable Jacobians

DOI: 10.54254/2755-2721/2026.BJ32383

Table 1. (continued)

Particle filter (PF)	Odometry/IMU; LiDAR scans; beacons	Posterior approximated by weighted samples	Multi-modal beliefs; robust to non-Gaussian noise; recovery	Compute grows with particles; impoverishment risk	Medium–High	Global localization and kidnapped-robot recovery
Factor graph / pose graph optimization	Odometry/IMU; LiDAR/vision constraints; loop closures	MAP on sparse graph; iterative nonlinear optimization	Global consistency; natural loop closure; extensible constraints	Needs robust association; heavier without incremental solvers	Medium–High	SLAM and long-term mapping
Learning-aided fusion	Classical sensors plus learned models	Training data represent domain	Improves perception and uncertainty/outlier modeling	Generalization risk; data/training cost	Medium–High (train), Low–Medium (run)	Hybrid systems with low-cost sensors

4.1. Kalman-family filters

Kalman filters represent the posterior as a Gaussian (mean and covariance), making them attractive for real-time embedded AMRs. EKF and UKF extend the approach to nonlinear models, but heavy-tailed errors and wrong noise tuning can still cause inconsistency. In practice, robust performance depends on credible noise models and simple outlier handling that does not violate Gaussian assumptions [8].

4.2. Particle filtering

Particle filters approximate the posterior with weighted samples, enabling multi-hypothesis tracking in ambiguous indoor spaces and during kidnapped-robot events. They are also convenient when the likelihood is strongly non-Gaussian (e.g., scan matching). Recent indoor systems fuse motion estimates with UWB ranging inside PF frameworks to maintain competing hypotheses under anchor uncertainty [5].

4.3. Graph optimization and smoothing

Smoothing formulates localization/SLAM as MAP inference on a sparse factor graph, unifying constraints from odometry, IMU, and perception. Incremental solvers can provide near-real-time updates while retaining global consistency. Loop-closure-enabled LiDAR odometry and mapping illustrate how global constraints reduce drift but also increase reliance on robust data association [4].

4.4. Learning-augmented estimation

Learning is most reliable as a component that strengthens specific links in the pipeline: semantic perception for better association, learned descriptors for loop closure, or models that predict uncertainty and reject outliers. These modules can improve robustness and enable cheaper sensors, but must be paired with an explicit estimation structure to avoid overconfident failures in unseen indoor domains [1].

DOI: 10.54254/2755-2721/2026.BJ32383

5. SLAM pipelines for indoor AMRs

5.1. Front-end: association and odometry

While the methods reviewed in Section 4 differ in how they represent and update the posterior belief, practical indoor localization systems are not built from estimators alone. In real AMRs, these estimation methods are embedded within complete SLAM pipelines that must generate reliable associations, manage accumulated constraints, and maintain consistency over time. For this reason, the discussion now shifts from algorithmic estimation principles to system-level SLAM design. The SLAM front-end generates relative motion constraints by associating current measurements with past observations. LiDAR front-ends rely on scan registration and geometric features, while visual front-ends track keypoints or direct photometric residuals. Degeneracy detection and robust matching are essential in repetitive corridors and texture-poor scenes [3,9].

5.2. Back-end: optimization and loop closure

Back-end optimization refines poses to satisfy accumulated constraints. Loop closure adds long-range constraints when a place is revisited, greatly reducing drift but risking false positives. Practical designs combine conservative place recognition with robust cost functions that down-weight inconsistent loops [4,9].

5.3. Multi-sensor SLAM

Multi-sensor SLAM exploits complementary observability: IMU stabilizes short-term motion, LiDAR provides metric structure, and cameras add semantics for association. Tightly coupled LiDAR–inertial methods show benefits when one sensor temporarily degrades [3].

5.4. Long-term mapping in dynamic indoor spaces

Indoor maps age as furniture and traffic patterns change, so long-term autonomy needs map maintenance and relocalization strategies. Service deployments also reveal integration needs beyond estimation: mapping must support navigation, task management, and recovery procedures. Restaurant-robot studies highlight the coupling between SLAM, obstacle avoidance, and operational workflows in crowded spaces [10].

6. Implementation considerations

6.1. Calibration and time synchronization

The SLAM pipelines discussed above describe how indoor AMRs build and maintain localization constraints at the system level. However, good pipeline design does not automatically translate into robust field performance. In practice, deployment quality is often determined by implementation details such as calibration accuracy, time synchronization, computational scheduling, and software integration. This section therefore turns from algorithmic structure to the practical considerations that shape real-world robustness. Estimator accuracy is often bounded by integration details. Extrinsic calibration is required to express measurements in a common frame, and small temporal offsets between LiDAR, camera, and IMU streams can bias motion estimates. Time synchronization,

DOI: 10.54254/2755-2721/2026.BJ32383

consistent frame conventions, and monitoring for parameter drift are therefore central to stable indoor fusion pipelines.

6.2. Real-time deployment and middleware

Real-time constraints favor predictable compute and graceful degradation. A common architecture uses a high-rate filter for local tracking plus a slower back-end for loop closure and map optimization. Middleware such as ROS supports modular implementation and debugging, and application reports show how ROS integrates SLAM, navigation, and task management in service environments [10].

7. Evaluation metrics

7.1. Accuracy and robustness

Once implementation issues are taken into account, the next question is how localization systems should be evaluated in a way that reflects deployment reality. Performance cannot be judged by trajectory accuracy alone, because practical AMRs must also recover from failures, operate within limited compute budgets, and remain maintainable over long-term use. This section therefore reviews evaluation metrics that better capture both algorithmic performance and operational usefulness. Accuracy is commonly reported using absolute trajectory error (ATE) and relative pose error (RPE), preferably over repeated runs to capture stochastic failures. Robustness requires additional measures such as relocalization success rate, time-to-recover after occlusion, and stress tests (crowds, lighting changes, corridor degeneracy). Infrastructure aids can be evaluated by how quickly they resolve global ambiguity, for example, via BLE-based floor identification [6].

7.2. Efficiency and operability

Deployable AMRs must balance accuracy with compute, latency, and power. Reporting should include update rate and CPU/GPU load, as well as operational metrics such as map build time, required anchor/tag density, and maintenance overhead. These measures reveal trade-offs between laboratory accuracy and stable long-term indoor operation.

8. Challenges and emerging trends

The evaluation criteria reviewed above make clear that many indoor localization systems still fall short of long-term deployment requirements, even when benchmark accuracy appears strong. In particular, changing environments, recovery behavior, generalization, and reliability monitoring remain insufficiently addressed in current research. Against this background, this section summarizes the major open challenges and emerging trends that are likely to shape the next stage of indoor AMR localization. Long-term autonomy remains difficult because indoor layouts change and dynamic objects can dominate sensor returns. Progress is needed in life-long mapping and principled handling of map updates without accumulating artifacts. Learning-based components also need stronger generalization guarantees; uncertainty-aware training and hybrid designs that preserve Bayesian consistency are likely to matter more than end-to-end replacements. Finally, indoor deployment at scale requires reliability monitoring: estimator confidence should trigger fallback behaviors, human supervision, or autonomy regulation. Work on regulating the degree of automation

DOI: 10.54254/2755-2721/2026.BJ32383

using robot "vitals" and "health" metrics is a step toward localization-aware autonomy management for safety-critical operation [11].

9. Conclusion

This review summarizes robust indoor localization for autonomous mobile robots from an algorithmic viewpoint. Indoor environments amplify failure modes such as aliasing, occlusion, and odometry drift, so reliable navigation typically depends on probabilistic fusion that combines motion sensing with environment-referenced observations and, when available, absolute infrastructure constraints. Kalman-family filters offer efficient tracking under near-Gaussian assumptions, particle filters provide multi-hypothesis recovery under ambiguity, and graph-based smoothing enables global consistency through loop closure. Learning-based modules are most effective as targeted enhancements for association, perception robustness, and uncertainty modeling.

Several gaps remain. Calibration and time synchronization errors are still under-reported despite dominating performance in practice, and many evaluations emphasize average accuracy rather than failure rates, recovery behavior, and maintenance cost. In addition, deployment decisions often depend on operational constraints such as compute budget, infrastructure availability, and the need to recover gracefully from temporary sensing failures.

Future research should connect estimator confidence to operational decisions, including map maintenance policies, safety-aware fallback behaviors, and autonomy management strategies. Benchmarking efforts would benefit from more realistic indoor scenarios that include crowds, rearranged layouts, and long-term operation, as well as reproducible reporting of calibration and synchronization procedures. Such directions can narrow the gap between strong laboratory results and robust AMR localization in everyday indoor deployments.

References

[1] Ušinskis, V.; Nowicki, M.; Dzedzickis, A.; Bučinskas, V. Sensor-Fusion Based Navigation for Autonomous Mobile Robot. Sensors 2025, 25, 1248. https: //doi.org/10.3390/s25041248.
[2] Alatise, M.B.; Hancke, G.P. A Review on Challenges of Autonomous Mobile Robot and Sensor Fusion Methods. IEEE Access 2020, 8, 39830–39846. https: //doi.org/10.1109/ACCESS.2020.2975643.
[3] Chen, J.; Wang, H.; Yang, S. Tightly Coupled LiDAR-Inertial Odometry and Mapping for Underground Environments. Sensors 2023, 23, 6834.
[4] Liu, Y.; Zhang, W.; Li, F.; Zuo, Z.; Huang, Q. Real-Time Lidar Odometry and Mapping with Loop Closure. Sensors 2022, 22, 4373.
[5] Durodié, Y.; Decoster, T.; Van Herbruggen, B.; Vanhie-Van Gerwen, J.; De Poorter, E.; Munteanu, A.; Vanderborght, B. A UWB-Ego-Motion Particle Filter for Indoor Pose Estimation of a Ground Robot Using a Moving Horizon Hypothesis. Sensors 2024, 24, 2164.
[6] Wu, H.; Liu, H.; Roddelkopf, T.; Thurow, K. BLE Beacon-Based Floor Detection for Mobile Robots in a Multi-Floor Automation Laboratory. Transportation Safety and Environment 2023, 6, tdad024.
[7] Tripicchio, P.; D'Avella, S.; Unetti, M.; Motroni, A.; Nepa, P. A UHF Passive RFID Tag Position Estimation Approach Exploiting Mobile Robots: Phase-Only 3D Multilateration Particle Filters With No Unwrapping. IEEE Access 2024, 12, 58778–58788.
[8] Khodarahmi, M.; Maihami, V. A Review on Kalman Filter Models. Archives of Computational Methods in Engineering 2023, 30, 727–747.
[9] Tourani, A.; Bavle, H.; Sanchez-Lopez, J.L.; Voos, H. Visual SLAM: What Are the Current Trends and What to Expect? Sensors 2022, 22, 9297.
[10] Patil, L.N.; Jadhav, S.S.; Waghulde, K.B.; Gaikwad, P.; Duchal, S.; Patil, S.; Fulshete, N.; Chatur-Deokar, M. Precision Mapping and Navigation: A Robotic Restaurant Management System Using SLAM and ROS. Sigma Journal of Engineering and Natural Sciences 2025, 43(4), 1233–1247. https: //doi.org/10.14744/sigma.2025.00117.

DOI: 10.54254/2755-2721/2026.BJ32383

[11] Braun, C.A.; Ramesh, A.; Rothfuß, S.; Chiou, M.; Stolkin, R.; Hohmann, S. Model Predictive Degree of Automation Regulation for Mobile Robots Using Robot Vitals and Robot Health. IFAC-PapersOnLine 2023, 56(2), 8345–8350. https: //doi.org/10.1016/j.ifacol.2023.10.1025.