Open Horizon Scientific Review

Vol. 04 No.01 2026

Sensor Fusion Algorithms for Improving Autonomous Navigation Accuracy

Zain Shahmeer Abbasi

Independent Researcher in Robotics and Autonomous Systems, Rawalpindi, Pakistan

ABSTRACT

Accurate autonomous navigation requires reliable estimates of position, velocity, and orientation despite noisy measurements, changing environments, and intermittent sensor availability. Sensor fusion combines complementary observations within a common estimation framework to reduce uncertainty and improve continuity. This article presents a narrative review and conceptual framework for navigation-focused sensor fusion, examining nine dimensions: sensor complementarity, calibration and synchronization, Kalman-family estimation, particle filtering, factor-graph optimization, multimodal odometry, robustness, real-time implementation, and performance evaluation. The discussion distinguishes local motion estimation from globally referenced localization and emphasizes that uncertainty reduction depends on valid models, sufficient observability, and correct treatment of correlated information. Representative visual–inertial and lidar–inertial systems illustrate the relationship between sensing and algorithm architecture. An analytical graph shows how correlation limits the uncertainty reduction obtained by averaging two equally precise measurements. It is a mathematical illustration rather than an experimental comparison of navigation algorithms. A proposed implementation workflow combines calibrated sensing, validated probabilistic models, residual monitoring, bounded computational latency, and application-specific benchmarking. The synthesis concludes that adding sensors does not automatically improve accuracy: timing errors, calibration bias, poor data association, and overconfident covariance estimates can undermine fusion. Credible performance claims therefore require consistent reference frames, representative failure scenarios, independent ground truth, and transparent reporting of accuracy, consistency, availability, and execution time.

Key words: *Sensor fusion, Autonomous navigation, State estimation, Kalman filtering, Localization accuracy*

Introduction:

Autonomous vehicles and mobile robots must estimate their motion while interacting with uncertain physical environments. Navigation requires more than recognizing obstacles: the platform needs a sufficiently accurate and timely estimate of where it is, how it is oriented, and how its state is changing. Errors in these quantities can affect mapping, route following, obstacle avoidance, and control.

Individual sensors provide incomplete and imperfect information. Inertial sensors offer rapid motion measurements but accumulate drift when integrated without correction. Cameras

observe rich environmental structure but can be affected by illumination, motion blur, and limited texture. Lidar provides geometric observations but may become weakly constrained in repetitive or geometrically degenerate scenes. Satellite navigation offers a global reference under suitable reception conditions but can suffer from obstruction and multipath.

Sensor fusion combines these observations through explicit models of state evolution and measurement generation. Its purpose is to use complementary information and represent uncertainty coherently. The probabilistic foundations described by Thrun et al. (2005) and Durrant-Whyte and Henderson (2008) explain why this process must account for both measurement quality and dependence between information sources.

Navigation accuracy also depends on the definition of the task. A trajectory can be locally smooth while drifting relative to a global frame. Conversely, a globally corrected trajectory may contain discontinuities that complicate control. Evaluation must distinguish these behaviors rather than treating every position estimate as an interchangeable output.

This article provides a narrative review and conceptual implementation framework, not a systematic review or newly benchmarked estimator. It examines established algorithm families and representative research systems. No dataset is processed here, and no experimental accuracy improvement is claimed. The accompanying graph illustrates a specific statistical relationship that helps explain when fusion benefits can diminish.

1. Sensor Complementarity and Navigation State Representation

The state vector should reflect the quantities required for the application and the measurement models available. A common inertial-navigation state includes position, velocity, orientation, accelerometer bias, and gyroscope bias. Additional states may describe sensor offsets, clock parameters, wheel calibration, or map features. Adding states increases modeling flexibility but also introduces estimation and observability requirements.

Sensor complementarity arises when different measurements constrain weaknesses in other sources. Inertial propagation can bridge short gaps between camera or lidar observations, while environmental measurements can constrain accumulated motion drift. Satellite measurements can provide global information that local odometry lacks. Wheel odometry can support ground-vehicle motion estimation but becomes unreliable under slip or inappropriate motion assumptions.

Observability concerns whether the available measurements and motion allow particular state components to be inferred. Monocular imagery alone does not generally determine metric scale without additional information. Visual–inertial estimation can recover metric quantities under suitable excitation and initialization conditions, but weak or degenerate motion can limit what is identifiable. Global position and heading may remain unobservable in a purely relative system without external references.

Architecture selection should therefore begin with the operating environment and required state, rather than an assumption that more sensors always provide more useful information. Sensor placement, overlapping fields of view, failure mechanisms, mass, power, and cost should be considered together. Redundant sensors can improve continuity, but a shared source of bias or environmental degradation can limit their independence and their contribution to accuracy.

2. Spatial Calibration, Timing, and Data Association

Measurements must be expressed in compatible reference frames before they can constrain a common state. Intrinsic calibration describes a sensor's internal measurement geometry or response, while extrinsic calibration relates the sensor frame to the platform or another sensor. Incorrect camera geometry, lidar alignment, or inertial-axis conventions can introduce systematic errors that a fusion algorithm misinterprets as motion.

Temporal alignment is equally important. A measurement corresponds to a physical acquisition time, not necessarily its arrival time at the estimator. Camera exposure, sensor buffering, communication delay, and clock drift can create offsets between streams. During rapid motion, even a modest timing error can cause substantial disagreement between otherwise accurate measurements.

Rolling-shutter images and scanning lidar measurements are acquired over intervals rather than at one instantaneous pose. Their interpretation may require motion compensation and an appropriate continuous-time or discretized motion model. Treating every point in a moving scan as simultaneous can distort geometric constraints and degrade the estimated trajectory.

Data association establishes which observations refer to the same physical feature or place. Incorrect feature matches or loop closures can introduce strong but false constraints. Geometric verification, consistency checks, and uncertainty-aware association help reduce this risk. Calibration and association quality should be validated independently where possible, because

tuning a filter to hide their errors can produce a misleadingly smooth yet biased estimate.

3. Kalman-Family Filtering for Recursive State Estimation

The Kalman filter recursively combines a predicted state with new observations. Under linear models, appropriate noise assumptions, and correctly specified statistics, it provides a minimum-mean-square-error estimate within its standard formulation (Kalman, 1960). Navigation problems are usually nonlinear, motivating extensions and alternative representations.

The extended Kalman filter linearizes process and measurement models around an estimated state. In inertial navigation, an error-state representation can express small corrections while retaining a suitable nominal orientation representation. Its performance depends on initialization, linearization validity, model quality, and numerical treatment of uncertainty.

The unscented Kalman filter propagates selected sigma points through nonlinear functions to approximate transformed moments (Julier & Uhlmann, 2004). It can avoid explicit Jacobians in parts of the estimator, but it does not guarantee superior performance for every navigation problem. Scaling choices, state dimension, non-Euclidean orientation, and computational cost still require careful handling.

The update balances predicted uncertainty against measurement uncertainty. Underestimating measurement noise can give excessive influence to unreliable observations, while overestimating it can prevent useful corrections. Cross-correlation must also be respected when measurements or upstream estimates share information. Mourikis and Roumeliotis (2007) demonstrate a visual–inertial filtering formulation that exploits constraints across multiple camera states without maintaining every observed feature as a persistent state variable.

4. Particle Filtering and Multimodal Localization

Particle filters represent a probability distribution through weighted samples. They are useful when uncertainty is substantially non-Gaussian or contains multiple plausible location hypotheses, such as global localization in a map with repeated structures. Unlike a single Gaussian approximation, the particle population can represent separated hypotheses before sufficient evidence resolves ambiguity.

Each particle is propagated using a motion model and assigned a weight according to the likelihood of observations. Resampling concentrates computation on promising regions of the state space. However, repeated resampling can reduce diversity, and a

poorly chosen likelihood can cause the filter to discard the correct hypothesis.

The computational burden grows with state complexity and the shape of uncertainty. A particle method that is effective for planar map localization may be expensive for a full high-dimensional inertial state. Proposal design, adaptive particle counts, and analytical treatment of selected subproblems can reduce this burden, but their applicability depends on the model structure.

Particle filtering should not be described as automatically more accurate than Kalman-family methods. Its value lies in representing uncertainty that a simpler approximation cannot adequately capture. Evaluation should examine relocalization success, sensitivity to initialization, particle depletion, and execution time. Thrun et al. (2005) provide the theoretical and practical context for matching probabilistic estimation methods to localization requirements.

5. Factor Graphs, Smoothing, and Inertial Preintegration

Factor graphs express estimation problems through variables connected by probabilistic constraints. Variables can represent poses, velocities, biases, and landmarks, while factors encode inertial motion, visual observations, lidar alignment, satellite fixes, or loop closures. This structure makes information sources and their relationships explicit (Dellaert & Kaess, 2017).

Smoothing jointly refines multiple states instead of updating only a current estimate. A fixed-lag window limits the number of active states, while larger graph formulations can support long-term map consistency. Sparse numerical methods make many such problems tractable, but graph connectivity and variable ordering influence computational cost.

Inertial preintegration summarizes high-rate inertial measurements between selected states. Forster et al. (2017) develop an on-manifold formulation suitable for visual–inertial estimation. Proper uncertainty propagation and bias handling allow the summarized measurements to be incorporated into repeated nonlinear optimization without simply reintegrating every sample at every iteration.

Smoothing introduces its own risks. Marginalization can preserve information using a local linear approximation, and later changes in the estimate can interact with that approximation. Incorrect loop closures may distort many states at once. Robust losses and verification can reduce the impact of outliers, but they do not make incorrect constraints harmless. Solver convergence, initialization, gauge freedoms, and numerical conditioning

should therefore be part of the implementation review.

6. Visual–Inertial and Lidar–Inertial Fusion Architectures

Visual–inertial systems combine image-based motion constraints with inertial propagation. Images provide information about environmental structure and relative movement, while inertial measurements support rapid state prediction and can help constrain metric motion under appropriate conditions. VINS-Mono is a representative tightly coupled monocular visual–inertial estimator with initialization and estimation mechanisms designed for practical operation (Qin et al., 2018).

Lidar–inertial systems combine geometric scan constraints with inertial motion information. Inertial estimates can support scan motion compensation and provide an initial estimate for alignment. LIO-SAM uses a smoothing-and-mapping architecture to integrate these sources (Shan et al., 2020). Such systems remain sensitive to calibration, timing, scene structure, and the quality of lidar correspondences.

Loosely coupled fusion combines outputs from separate estimators, whereas tightly coupled fusion incorporates lower-level measurements or residuals within a shared estimation problem. Tight coupling can preserve information that an upstream summary discards, but it usually increases implementation complexity. Loose coupling can support modular development, although the correlation between upstream estimates must not be ignored.

The preferred architecture depends on the operating conditions and failure tolerance required. Additional sensing should be evaluated through ablation studies that remove selected modalities under controlled conditions. A multimodal system should demonstrate its behavior when one source degrades, rather than only showing nominal results when every sensor performs well. Published research systems illustrate architectural choices; they do not establish a universal ranking across platforms or environments.

7. Robustness, Correlated Errors, and Uncertainty Consistency

Fusion improves estimation only when the information entering the algorithm is sufficiently reliable and correctly modeled. Outliers, multipath satellite signals, moving visual features, wheel slip, and incorrect map matches can create errors that violate assumed distributions. Innovation monitoring, geometric checks, robust estimation, and sensor-health indicators can help identify inconsistent observations.

Covariance describes modeled uncertainty, not a guarantee of actual error. An estimator can report small covariance while remaining biased

because calibration error or shared information was omitted. Consistency assessment compares uncertainty predictions with observed estimation errors under an appropriate statistical model and test design. Small average error alone does not establish consistent confidence reporting.

Correlation provides a simple illustration. Consider two unbiased scalar measurements of the same fixed quantity, each with standard deviation σ and correlation coefficient ρ. Their equal-weight average has variance $\sigma^2(1 + \rho)/2$. With independent errors, averaging reduces variance by half; with perfectly correlated errors, averaging provides no variance reduction. The result follows from the variance-of-a-sum identity and does not require a navigation experiment.

8. Real-Time Implementation and Degraded-Mode Operation

A navigation estimate must arrive in time to support control. High accuracy on a delayed trajectory can be inadequate for a moving platform. Evaluation should therefore measure end-to-end latency from acquisition to usable state output, including synchronization, buffering, feature extraction, optimization, and communication.

High-rate inertial propagation can provide timely predictions between slower optimization updates. The implementation must manage corrections without confusing the controller about frame changes or introducing uncontrolled state jumps. Local odometry and globally corrected map frames can serve different purposes, provided their relationship is explicit and maintained consistently.

Memory limits, processor contention, thermal throttling, and worst-case scene complexity can affect execution time. Average frame rate is not enough to establish reliable timing. Queue growth, dropped measurements, stale outputs, and recovery after computational overload should be examined. Preprocessing and measurement selection should preserve the information needed for observability rather than reducing workload indiscriminately.

Degraded operation requires explicit behavior when a sensor becomes unreliable or unavailable. The estimator may continue with increased uncertainty, switch to another measurement model, or report insufficient localization quality. The navigation system should connect that condition to an application-appropriate response. Continuing to output a precise-looking trajectory without acknowledging lost constraints can be more misleading than a clearly declared reduction in capability.

9. Benchmarking, Accuracy Metrics, and Deployment Validation

A defensible evaluation compares estimates with independent ground truth and states the alignment convention. Absolute trajectory error measures disagreement after a specified alignment, while relative pose error examines local motion differences over chosen intervals. Translation, rotation, scale handling, timestamp association, and evaluation intervals should be reported explicitly.

Reference frames and alignment choices can change the apparent result. Allowing scale alignment can conceal a scale error that matters for metric navigation, while unrestricted trajectory alignment can remove global offsets that the application requires the estimator to control. Evaluations should therefore match the intended localization task rather than selecting the most favorable convention.

The KITTI benchmark supports comparative research in outdoor driving, while the TUM RGB-D benchmark provides established trajectory-evaluation practices for indoor visual systems (Geiger et al., 2012; Sturm et al., 2012). The Oxford RobotCar dataset adds repeated traversals and environmental variation relevant to long-term operation (Maddern et al., 2017). None of these datasets alone establishes suitability for every deployment domain.

Validation should report error distributions, failures, availability, initialization time, recovery behavior, and computational latency alongside mean or root-mean-square error. Controlled sensor outages, calibration perturbations, difficult motion, changing weather, and domain shifts should be considered where relevant. A final deployment decision should connect these results to application requirements and independently verified sensing conditions. The framework presented here is a conceptual synthesis and requires experimental implementation before any platform-specific accuracy claim can be made.

Summary

Sensor fusion can improve autonomous navigation by combining complementary measurements within a coherent estimation framework. Kalman-family filters support recursive state estimation, particle filters represent multimodal uncertainty, and factor-graph methods jointly refine multiple states and constraints. Their effectiveness depends on suitable models, observable motion, calibration, synchronization, and reliable association.

Representative visual–inertial and lidar–inertial architectures show how different sensing modalities can support one another. However, correlated errors, shared bias, invalid constraints, and computational delay can undermine

expected benefits. The analytical graph demonstrates that averaging equally precise measurements yields less uncertainty reduction as their errors become more correlated; it establishes no ranking of navigation algorithms.

Credible performance evaluation should use independent ground truth, transparent alignment rules, representative operating conditions, and consistent reporting of accuracy and uncertainty. Real-time continuity and failure recovery are also necessary to understand practical navigation quality. Sensor quantity is therefore an incomplete measure of capability: the value of fusion lies in how reliably the system interprets, combines, and validates its information.

References:

- Dellaert, F., & Kaess, M. (2017). Factor graphs for robot perception. Foundations and Trends in Robotics, 6(1–2), 1–139. https://doi.org/10.1561/2300000043
- Durrant-Whyte, H., & Henderson, T. C. (2008). Multisensor data fusion. In B. Siciliano & O. Khatib (Eds.), Springer handbook of robotics (pp. 585–610). Springer. https://doi.org/10.1007/978-3-540-30301-5_26
- Forster, C., Carlone, L., Dellaert, F., & Scaramuzza, D. (2017). On-manifold preintegration for real-time visual–inertial odometry. IEEE Transactions on Robotics, 33(1), 1–21. https://doi.org/10.1109/TRO.2016.2597321
- Geiger, A., Lenz, P., & Urtasun, R. (2012). Are we ready for autonomous driving? The KITTI vision benchmark suite. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 3354–3361). https://doi.org/10.1109/CVPR.2012.6248074
- Julier, S. J., & Uhlmann, J. K. (2004). Unscented filtering and nonlinear estimation. Proceedings of the IEEE, 92(3), 401–422. https://doi.org/10.1109/JPROC.2003.823141
- Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. Journal of Basic Engineering, 82(1), 35–45. https://doi.org/10.1115/1.3662552
- Maddern, W., Pascoe, G., Linegar, C., & Newman, P. (2017). 1 year, 1000 km: The Oxford RobotCar dataset. The International Journal of Robotics Research, 36(1), 3–15. https://doi.org/10.1177/0278364916679498
- Mourikis, A. I., & Roumeliotis, S. I. (2007). A multi-state constraint Kalman filter for vision-aided inertial navigation. In Proceedings of the IEEE International Conference on Robotics and Automation (pp. 3565–3572).

https://doi.org/10.1109/ROBOT.2007.364024

- Qin, T., Li, P., & Shen, S. (2018). VINS-Mono: A robust and versatile monocular visual-inertial state estimator. IEEE Transactions on Robotics, 34(4), 1004–1020. https://doi.org/10.1109/TRO.2018.2853729
- Shan, T., Englot, B., Meyers, D., Wang, W., Ratti, C., & Rus, D. (2020). LIO-SAM: Tightly-coupled lidar inertial odometry via smoothing and mapping. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (pp. 5135–5142). https://doi.org/10.1109/IROS45743.2020.9341176
- Sturm, J., Engelhard, N., Endres, F., Burgard, W., & Cremers, D. (2012). A benchmark for the evaluation of RGB-D SLAM systems. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (pp. 573–580). https://doi.org/10.1109/IROS.2012.6385773
- Thrun, S., Burgard, W., & Fox, D. (2005). Probabilistic robotics. MIT Press.