

Towards All-Weather sub-50-cm Radar-Inertial Positioning

Lakshay Narula

Department of Electrical & Computer Engineering

The University of Texas at Austin

Austin, TX 78712

`lakshay.narula@utexas.edu`

Peter A. Iannucci

Department of Aerospace Engineering & Engineering Mechanics

The University of Texas at Austin

Austin, TX 78712

Todd E. Humphreys

Department of Aerospace Engineering & Engineering Mechanics

The University of Texas at Austin

Austin, TX 78712

Abstract

Deployment of automated ground vehicles beyond the confines of sunny and dry climates will require sub-lane-level positioning techniques based on radio waves rather than near-visible-light radiation. Like human sight, lidar and cameras perform poorly in low-visibility conditions. This paper develops and demonstrates a novel technique for robust sub-50-cm-accurate urban ground vehicle positioning based on all-weather sensors. The technique incorporates a computationally-efficient globally-optimal radar scan batch registration algorithm into a larger estimation pipeline that fuses data from commercially-available low-cost automotive radars, low-cost inertial sensors, vehicle motion constraints, and, when available, precise GNSS measurements. Performance is evaluated on an extensive and realistic urban dataset of all-weather sensors. Comparison against ground truth shows that during 60 min of GNSS-denied driving in the urban center of Austin, TX, the technique maintains 95th-percentile errors below 50 cm in horizontal position and 0.5° in heading.

1 Introduction

Development of automated ground vehicles (AGVs) has spurred research in lane-keeping assist systems, automated intersection management [Fajardo et al., 2011], tight-formation platooning, and cooperative sensing [Choi et al., 2016, LaChapelle et al., 2020], all of which demand accurate (e.g., 50-cm at 95%) ground vehicle positioning in an urban environment. But the majority of positioning techniques developed thus far depend on lidar or cameras, which perform poorly in low-visibility conditions such as snowy whiteout, dense fog, or heavy rain. Adoption of AGVs in many parts of the world will require all-weather localization techniques.

Radio-wave-based sensing techniques such as radar and GNSS (global navigation satellite system) remain operable even in extreme weather conditions (see Sec. 2) because their longer-wavelength electromagnetic radiation penetrates snow, fog, and rain. Carrier-phase-differential GNSS (CDGNSS) has been successfully applied for the past two decades as an all-weather decimeter-accurate localization technique in open-sky conditions. Proprioceptive sensors such as inertial measurement units (IMUs) also continue to operate regardless of external conditions. Coupling a CDGNSS receiver with a tactical-grade inertial sensor, as in [Petovello et al., 2004, Scherzinger, 2006, Zhang, 2006, Kennedy et al., 2006] delivers robust high-accuracy positioning even during the extended signal outages common in the urban environment, but such systems are far too expensive for widespread deployment on AGVs. Recent work has shown that 20-cm-accurate (95%) CDGNSS positioning is possible at low cost even in dense urban areas, but solution availability remains below 90%, with occasional long gaps between high-accuracy solutions [Humphreys et al., 2020]. Moreover, the global trend of increasing radio interference in the GNSS bands, whether accidental or deliberate [Humphreys, 2017], underscores the need for GNSS-independent localization: GNSS jamming cannot be allowed to paralyze an area’s automated vehicle networks.

Clearly, there is a need for AGV localization that is low cost, accurate at the sub-50-cm level, robust to low-visibility conditions, and continuously available. This paper is the first to establish that low-cost inertial- and automotive-radar-based localization can meet these criteria.

Mass-market commercialization has brought the cost of automotive radar down enough that virtually all current production vehicles are equipped with at least one radar unit, which serves as the primary sensor for adaptive cruise control and automatic emergency braking. But use of automotive radar for localization faces the significant challenges of data sparsity and noise: an automotive radar scan has vastly lower resolution than a camera image or a dense lidar scan, and is subject to high rates of false detection (clutter) and missed detection. As such, it is nearly impossible to deduce semantic information or to extract distinctive environmental features from an individual radar scan. This is clear from Fig. 1c, which shows a sparse smattering of reflections from a single composite scan using three radar

units. The key to localization is to aggregate sequential scans into a batch, as in Fig. 1d, where environmental structure is clearly evident. Even still, the data remain so sparse that localization based on traditional machine vision feature extraction and matching is not promising. Additionally, stable short-term odometry is a pre-requisite for aggregating radar scans, which in itself is a challenge when dealing with low-cost inertial sensing.

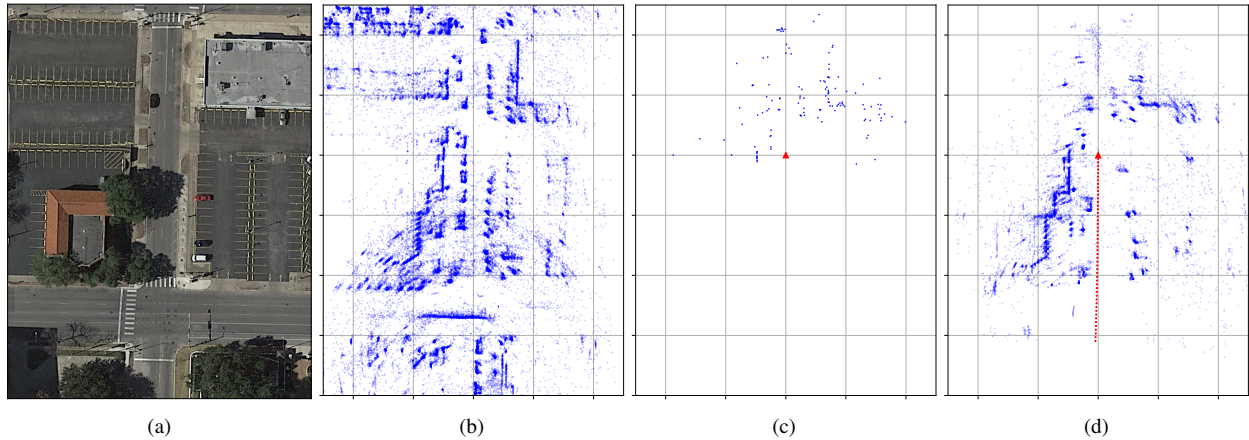


Figure 1: Panel (a) shows a satellite view of the environment being mapped with automotive radar. Panel (b) shows the generated radar map point cloud with vehicle pose obtained from a reference localization system. Note the repeating structure along the road side due to parked vehicles. An individual radar scan obtained during localization is shown in panel (c), along with the red triangle denoting vehicle location and heading. The scan is sparse and contains significant clutter, making it challenging to register to the prior map. Panel (d) shows a batch of radar scans during localization, with the red dots denoting the vehicle trajectory over the past five seconds. The batch captures the underlying structure which can be registered to the prior map.

This paper proposes a two-step process for radar-based localization. The first is the mapping step: creation of a geo-referenced two-dimensional aggregated map of all radar targets across an area of interest. Fig. 1b shows such a map, hereafter referred to as a radar map. The full radar map used throughout this paper, of which Fig. 1b is a part, was constructed with the benefit of a highly stable inertial platform so that a trustworthy ground truth map would be available against which maps generated by other techniques could be compared. But an expensive inertial system or dedicated mobile mapping vehicle is not required to create a radar map. Instead, it can be crowd-sourced from the very user vehicles that will ultimately exploit the map for localization. During periods of favorable lighting conditions and good visibility, user vehicles can exploit a combination of low-cost CDGNSS, as in [Humphreys et al., 2020], and GNSS-aided visual simultaneous localization and mapping, as in [Narula et al., 2018b], to achieve the continuous decimeter-and-sub-degree-accurate geo-referenced position and orientation (pose) required to lay down an accurate radar map. In other words, the radar map can be *created* when visibility is good and then *exploited* at any later time, such as during times of poor visibility.

Despite aggregation over multiple vehicle passes, the sparse and cluttered nature of automotive radar data is evident from the radar map shown in Fig. 1b: the generated point cloud is much less dense and has a substantially higher

fraction of spurious returns than a typical lidar-derived point cloud, making automotive-radar-based localization a significantly more challenging problem.

The second step of this paper’s technique is the localization step. Using a combination of all-weather odometric techniques such as inertial sensing, radar odometry, and ground vehicle dynamics constraints, a sensor fusion filter continually tracks the changes in vehicle pose over time. Over the latest short interval (e.g., 5 s), pose estimates from the filter are used to spatially organize the multiple radar scans taken over the interval and generate what is hereafter referred to as a batch of scans, or batch for short. Fig. 1d shows a five-second batch terminating at the same location as the individual scan in Fig. 1c. In contrast to the individual scan, some environmental structure emerges in the batch of scans, making robust registration to the map feasible. Even so, the localization problem remains challenging due to the dynamic radar environment: note the absence of parked cars on the left side of the street during localization. The batch of scans is matched against the prior map of the surroundings to estimate the pose offset of the batch from the truth. This pose offset is then applied as a measurement to the sensor fusion filter to correct odometric drift.

Contributions. This paper’s overall contribution is a robust pipeline for all-weather sub-50-cm urban ground vehicle positioning. This pipeline incorporates a computationally-efficient correlation-maximization-based globally-optimal radar scan registration algorithm that estimates a two-dimensional translational and a one-dimensional rotational offset between a prior radar map and a batch of current scans. Significantly, the registration algorithm can be applied to the highly sparse and cluttered data produced by commercially-available low-cost automotive radars. Maximization of correlation is shown to be equivalent to minimization of the L^2 distance between the prior map and the batch probability hypothesis densities. The pipeline supports the construction of the radar registration estimate and optimally fuses it with inertial measurements, radar range rate measurements, ground vehicle dynamics constraints, and cm-accurate GNSS measurements, when available. A novel technique for online estimation of the vehicle center of rotation is introduced, and calibration of various other extrinsic parameters necessary for optimal sensor fusion is described.

This paper also presents a thorough evaluation of the positioning pipeline on the large-scale dataset described in [Narula et al., 2020b]. Data from automotive sensors are collected over two 1.5 h driving sessions through the urban center of Austin, TX on two separate days specifically chosen to provide variety in traffic and parking patterns. The dataset is collected in clear weather conditions, but only includes data from sensors that are expected to remain unaffected in adverse weather. Comparison with a post-processed ground truth trajectory shows that proposed pipeline maintains 95th-percentile errors below 35 cm in horizontal position and 0.5° in heading during 60 min of GNSS-denied driving.

A preliminary version of this paper describing the radar scan registration algorithm was published in [Narula et al., 2020a]. The current version develops and tests a full sensor fusion pipeline that includes the radar batch estimation as a sub-component.

Organization of the rest of this paper. Sec. 2 provides a survey of millimeter-wave radar performance in adverse weather. Sec. 3 establishes the significance of this contribution in view of the prior work in related fields. The radar batch-based pose estimation technique for the low-cost automotive radar sensor model is developed in Sec. 4. Sec. 5 describes the overall sensor fusion architecture involving inertial sensing, GNSS, motion constraints, and radar measurements. Implementation details and experimental results from field evaluation are presented in Sec. 6, and Sec. 7 provides concluding remarks.

2 Radar-Based All-Weather Perception

Perception in adverse weather conditions has been recognized as one of the major challenges for automated driving [Kyle Stock, 2018]. Optical cameras, widely featured as the primary sensors in AGVs, are limited by the prevailing visibility range during dense fog or heavy rain, and become inoperative with even slight build up of rain drops, snow, or dust on the camera lens (e.g., see [Hong et al., 2020, Fig. 9]). Moreover, the camera-perceived environment changes drastically during and after adverse weather events such as a snow storm. These impairments make cameras ill-suited for all-weather localization.

Several experimental automated driving projects employ lidars in addition to camera-based sensing. Lidars provide direct depth measurements of the surrounding environment and remain equally effective at all times of the day. Unfortunately, however, lidars do not fare much better than cameras under adverse weather conditions. The suspended water droplets in fog cause significant attenuation and backscatter of the 900 nm or 1500 nm laser radiation emitted by lidars (e.g., see [Bijelic et al., 2018, Figs. 1, 5] or [Li et al., 2020, Fig. 7]). In addition to the reduced sensing range, ghost returns due to backscatter from fog are challenging for mapping and localization systems. Similarly, lidars are inoperative during and after severe snow storms as shown in [Jokela et al., 2019]; lidars from five different manufacturers fail to obtain any reflections from a vehicle 20 m in front of the lidar due to severe backscatter from the snow dust kicked up by the leading vehicle.

Millimeter-wave radar performance in adverse weather conditions has been empirically and theoretically studied for several decades in many different applications [Hong et al., 2020, Yoneda et al., 2018, Reina et al., 2011, Ryde and Hillier, 2009, Yen et al., 2015, Brooker et al., 2007, Foessel et al., 1999, Chen, 1975, Kobayashi, 1980, Wallace, 1988,

[Mohammed et al., 2020](#), [Kutla et al., 2018](#), [Rasshofer and Gresser, 2005](#), [Zang et al., 2019](#)], ranging from missile guidance and mining to collision warning systems and vehicular perception. As the survey below suggests, millimeter-wave radar is remarkably robust in all weather conditions due to its much longer wavelength as compared to visible-light cameras and lidars.

The feasibility of radar-based all-weather mapping and localization has been recently studied in [[Hong et al., 2020](#)] and [[Yoneda et al., 2018](#)]. Reference [[Yoneda et al., 2018](#)] compares the accuracy of 76 GHz automotive-radar-based localization in clear weather and snowy conditions. Reference [[Yoneda et al., 2018](#), Table I] shows that radar-based localization achieves the same accuracy in clear conditions and in partially-covered snow conditions, while the lidar-based baseline method breaks down in partially-covered snow. Longitudinal accuracy of radar-based localization begins to degrade in fully-covered snow conditions, while lateral accuracy remains unaffected. The SLAM system presented in [[Hong et al., 2020](#)] uses a mechanically-rotating 77 GHz radar, which, although different in the working principle from the phased-array automotive radars, is identical in terms of signal propagation. Through experiments in rain, fog, and snow, [[Hong et al., 2020](#)] demonstrates the feasibility of all-weather localization and mapping with millimeter-wave radar, while camera- and lidar-based SLAM failed in the same experiments. Interestingly, [[Hong et al., 2020](#), Fig. 8] reports partial loss of radar returns after a thick layer of ice deposits on the radar, but the proposed system nevertheless performs adequately.

The SLAM algorithm developed in [[Hong et al., 2020](#)] is not directly applicable to the data made available by low-cost automotive radars, since the SURF feature extraction technique used in [[Hong et al., 2020](#)] cannot operate on scans such as the one shown in Fig. 1c. On the other hand, the results from [[Yoneda et al., 2018](#)] are especially interesting for this paper, since the two systems employ similar radar sensors, as well as a correlation-based registration method. In comparison to this paper, [[Yoneda et al., 2018](#)] provides no probabilistic justification for the correlation-based approach, assumes perfect, hypothetical odometry information during generation of radar batches, and only estimates a two-dimensional translational offset with heading assumed to be perfectly known. Meanwhile, this paper develops a complete sensor fusion pipeline with radar, GNSS, and IMU, and estimates the full three degrees-of-freedom state of the vehicle.

Beyond localization and mapping, [[Reina et al., 2011](#)] evaluates the performance of a 95 GHz radar for obstacle perception in low-visibility conditions. Reference [[Reina et al., 2011](#), Figs. 19, 20] provides an excellent example of the advantage offered by millimeter-wave radar in a dust storm. Whereas lidar-based perception totally breaks down in a moderate dust storm, the radar output remains unaffected. Similarly, an empirical study of a 95 GHz radar in [[Ryde and Hillier, 2009](#)] concludes that the radar measurements suffer no perceptible degradation even when tested in severe rain (50–70 mm/h) and dense dust (10 m visibility). For reference, rainfall more intense than 25 mm/h is rare.

In another study [Yen et al., 2015], a Delphi 77 GHz radar, same as the one used later in this paper, is deployed in a collision warning system for a snow plow for operation during snow storms in the Sierra Nevada mountains. The Delphi radar is found to perform adequately for collision warning even in these extreme operating conditions, except when a thick layer of ice is accumulated on the radar fascia [Yen et al., 2015, Chap. 4], similar to the observation in [Hong et al., 2020]. Mining is another application that demands perception through cavities filled with dust and water vapor shortly after a blast. Reference [Brooker et al., 2007] conducts a theoretical analysis of the expected attenuation and backscatter for a 77 GHz radar in mining environments, and concludes that an attenuation of 10 dB/km may be expected in fog with only 4 m visibility, while no perceptible attenuation is expected from dust. The theoretical analysis suggests negligible backscatter from both dust and water droplets, and empirical results from successful deployment and long-term testing of millimeter-wave radar in the worst possible mining environments are shown to support the theoretical models. Similarly, the performance of a 77 GHz radar in blowing snow is presented in [Foessel et al., 1999], where empirical data collected in Antarctica show robust target measurements during snowfall. These empirical studies at 77 GHz reaffirm this paper’s claim of all-weather operation of automotive radars.

Millimeter-wave radar has also been deployed for tactical guidance and communications for several decades. Several measurements of the attenuation of electromagnetic radiation through adverse weather elements are aggregated in [Chen, 1975], ranging from the RF radiation to the visible spectrum. The study concludes that haze, fog, and clouds are transparent to the RF spectrum radiation of millimeter-wave radars. Heavy rain is concluded to be the dominant source of attenuation in the RF spectrum, but even in severe rain of 50 mm/h, a 77 GHz radar is predicted to attenuate by no more than 20 dB/km, or about 2 dB over 100 m, which is the typical maximum operating range necessary for mapping and localization. Similar results and conclusions have been made in other studies [Kobayashi, 1980, Wallace, 1988], with millimeter-wave attenuation no more than 15 dB/km in the most intense rain and snow conditions.

Overall, a thorough survey of the existing literature suggests that 77 GHz automotive radar experiences negligible degradation in adverse weather, except after build up of ice on the radar fascia itself. The major challenge with low-cost automotive-radar-based positioning is instead dealing with the poor angular resolution, as well as the sparse and cluttered point cloud data made available by these sensors in comparison to the high-resolution images and point clouds generated with cameras and lidars. This paper presents a sensor fusion engine that provides robust sub-lane-level accurate localization despite the poor automotive radar sensor characteristics.

3 Related Work

This section reviews a wide variety of literature on mapping and localization with radar and radar-inertial sensing. This includes prior work on point cloud alignment and image registration techniques, occupancy grid-based mapping and localization, random-finite-set-based mapping and localization, and inertial-aided mapping and localization.

Related work in point cloud alignment. A radar-based map can have many different representations. One obvious representation is to store all the radar measurements as a point cloud. Fig. 1b is an example of this representation. Localization within this map can be performed with point cloud registration techniques like the iterative closest point (ICP) algorithm. ICP is known to converge to local minima which may occur due to outlying points that do not have correspondences in the two point clouds being aligned. A number of variations and generalizations of ICP robust to such outliers have been proposed in the literature [Chetverikov et al., 2002, Ward and Folkesson, 2016, Holder et al., 2019, Tsin and Kanade, 2004, Jian and Vemuri, 2010, Myronenko and Song, 2010, Gao and Tedrake, 2019]. A few of these have been applied specifically to automotive radar data [Ward and Folkesson, 2016, Holder et al., 2019]. But the technique in [Ward and Folkesson, 2016] is only evaluated on a 5 min dataset, while [Holder et al., 2019] performs poorly on datasets larger than 1 km.

This paper steers away from ICP and its gradient-based variants because automotive radar data in urban areas exhibit another source of incorrect-but-plausible registration solutions which are not addressed in the above literature—repetitive structure, e.g., due to a series of parked cars, may result in multiple locally-optimal solutions within 2–3 m of the globally-optimal solution. Gradient-based techniques which iteratively estimate correspondences based on the distance between pairs of points are susceptible to converge to such locally-optimal solutions. Accordingly, the batch-based pose estimator proposed in this paper is designed to approximate the globally-optimal solution.

In contrast to ICP and its variants, globally-optimal point cloud registration can be achieved by performing global point correspondence based on distinctive feature descriptors [Cen and Newman, 2018, Cen and Newman, 2019, Barnes and Posner, 2020]. All of these works use a sophisticated mechanically-rotating radar unit that is not expected to be available on an AGV. Feature description and matching on the low-cost automotive radars used in this paper is likely to be fragile. Even when using the mechanically-rotating radar, [Barnes et al., 2019] shows that a correlation-based approach, such as the one developed in this paper, outperforms other feature-descriptor-based point cloud methods.

Related work in image registration and occupancy grid techniques. Occupancy grid mapping and localization techniques have been traditionally applied for lidar-based systems, and recent work in [Schuster et al., 2016, Schoen et al.,] has explored similar techniques with automotive radar data. In contrast to batch-based pose estimation de-

scribed in this paper, both [Schuster et al., 2016] and [Schoen et al.,] perform particle-filter based localization with individual scans, as is typical for lidar-based systems. These methods were only evaluated on small-scale datasets collected in a parking lot, and even so, the reported meter-level localization accuracy is not sufficient for lane-level positioning.

Occupancy grid maps are similar to camera-based top-down images, and thus may be aligned with image registration techniques, that may be visual-descriptor-based [Callmer et al., 2011, Hong et al., 2020] or correlation-based [Yoneda et al., 2018]. Reliable extraction and matching of visual features, e.g., SIFT or SURF, is significantly more challenging with automotive radar data. Correlation-based registration is more robust [Yoneda et al., 2018, Barnes et al., 2019], and forms the basis of one of the components in this paper. In contrast to prior work [Yoneda et al., 2018, Barnes et al., 2019], this paper provides a probabilistic interpretation for the correlation operation. The mechanically-rotating radar of [Barnes et al., 2019] allows correlation-based pose estimation based on a single scan of radar data. But for the low-cost automotive radars used in this paper, it becomes necessary to accumulate radar scans over time, which requires integration with other odometric sensors. This paper develops and demonstrates a complete sensor fusion pipeline around radar-based pose estimation and evaluates its performance on a large urban dataset.

Related work in random finite set mapping and localization. The occupancy grid representation commonly used in robotics is an approximation to the probability hypothesis density (PHD) function [Mahler, 2003, Erdinc et al., 2009]: a concept first introduced in the random finite set (RFS) based target tracking literature. Unsurprisingly, PHD- and RFS-based mapping and localization have been previously studied in [Mullane et al., 2011, Deusch et al., 2015, Stübler et al., 2017, Lundgren et al., 2014]. In contrast to occupancy grid-based methods, techniques in [Mullane et al., 2011, Deusch et al., 2015, Stübler et al., 2017, Lundgren et al., 2014] make the point target assumption where no target may generate more than one measurement in a single scan, and no target may occlude another target. However, in reality, planar and extended targets such as walls and building fronts are commonplace in the urban AGV environment. Mapping of ellipsoidal extended targets has recently been proposed in [Fatemi et al., 2017], but occlusions are not modeled and only simulation results are presented.

Related work in inertial-aided mapping and localization. This paper couples radar batch-based pose estimation with other all-weather automotive sensors such as IMU and GNSS. Inertial aiding has been widely applied in visual- and lidar-based mapping and localization [Qin et al., 2018, Mur-Artal and Tardós, 2017, Chiang et al., 2020, Forster et al., 2013, Steder et al., 2008, Ye et al., 2019, Li et al., 2014]. This paper extends inertial-aiding to sensors that can operate under harsh weather conditions. Recently, radar measurements have been applied to constrain IMU position drift in [Barra et al., 2019]. Radar-inertial odometry for indoor robots has been proposed in [Almalioglu et al., 2019, Kramer et al., 2020]. This paper is the first to integrate low-cost automotive radars with inertial sensing, GNSS, and

ground vehicle dynamics for lane-level accurate positioning in challenging urban environments.

4 Radar-Batch-Based Pose Estimation

This section describes the formulation of the radar-batch-based pose estimation method introduced in this paper. It first details the statistical motivation behind the method, and then develops an efficient approximation to the globally-optimal estimator. The output of this estimator acts as one of the measurements provided to the overall localization system presented later in Sec. 5.

4.1 Pose Estimation using Probability Hypothesis Density

For the purpose of radar-based pose estimation, an AGVs environment can be described as a collection of arbitrarily shaped radar reflectors in a specific spatial arrangement. Assuming sufficient temporal permanence of this environment, radar-equipped AGVs make sample measurements of the underlying structure over time.

4.1.1 The Probability Hypothesis Density Function

A probabilistic description of the radar environment is required to set up radar-based pose estimation as an optimization problem. This paper chooses the PHD function [Mahler, 2003] representation of the radar environment. The PHD at a given location gives the density of the expected number of radar reflectors at that location. For a static radar environment, the PHD $D(\mathbf{x})$ at a location $\mathbf{x} \in \mathcal{X}$ can be written as

$$D(\mathbf{x}) = I \cdot p(\mathbf{x})$$

where \mathcal{X} is the set of all locations in the environment, $p(\mathbf{x})$ is a probability density function such that $\int p(\mathbf{x})d\mathbf{x} = 1$, and I , the intensity, is the total number of radar reflectors in the environment. For a time-varying radar environment, both I and $p(\mathbf{x})$ are functions of time. For $\mathcal{A} \subset \mathcal{X}$, the expected number of radar reflectors in \mathcal{A} is given as

$$I_{\mathcal{A}} = \int_{\mathcal{A}} D(\mathbf{x})d\mathbf{x}$$

4.1.2 Estimating Vehicle State from PHDs

Let $D_m(\mathbf{x})$ denote the “map” PHD function representing the distribution of radar reflectors in an environment, estimated as a result of mapping with known vehicle poses. During localization, the vehicle makes a radar scan, or a series of consecutive radar scans. A natural solution to the pose estimation problem may be stated as the vehicle pose which maximizes the likelihood of the observed batch of scans, given that the scan was drawn from $D_m(\mathbf{x})$ [Myronenko and Song, 2010]. This maximum likelihood estimate (MLE) has many desirable properties such as asymptotic efficiency. However, the MLE solution is known to be sensitive to outliers that may occur if the batch of scans was sampled from a slightly different PHD, e.g., due to variations in the radar environment between mapping and localization [Jian and Vemuri, 2010].

A more robust solution to the PHD-based pose estimation problem may be stated as follows. Let Θ denote the vector of parameters of the rigid or non-rigid transformation \mathcal{T} between the vehicle’s prior belief of its pose, and its true pose. For example, in case of a two-dimensional rigid transformation, $\Theta = [\Delta x, \Delta y, \Delta \phi]^\top$, where Δx and Δy denote a two-dimensional position and $\Delta \phi$ denotes heading. Also, let $D_b(\mathbf{x}')$ denote a local “batch” PHD function estimated from a batch of scans during localization, defined over $\mathbf{x}' \in \mathcal{A} \subset \mathcal{X}$. This PHD is represented in the coordinate system consistent with vehicle’s prior belief, such that $\mathbf{x}' = \mathcal{T}_\Theta(\mathbf{x})$. Estimating the vehicle pose during localization is defined as estimating Θ such that some distance metric between the PHDs $D_m(\mathbf{x})$ and $D_b(\mathbf{x}')$ is minimized.

This paper chooses the L^2 distance between $D_m(\mathbf{x})$ and $D_b(\mathbf{x}')$ as the distance metric to be minimized. As compared to the MLE which minimizes Kullback-Leibler divergence, L^2 minimization trades off asymptotic efficiency for robustness to measurement model inaccuracy [Jian and Vemuri, 2010]. The L^2 distance $d_{L^2}(\Theta)$ to be minimized is given as

$$d_{L^2}(\Theta) = \int_{\mathcal{A}} (D_m(\mathbf{x}) - D_b(\mathcal{T}_\Theta(\mathbf{x})))^2 d\mathbf{x}$$

For rigid two-dimensional transformations, it can be shown as follows that minimizing the L^2 distance between the PHDs is equivalent to maximization of the cross-correlation between the PHDs.

$$\begin{aligned} \hat{\Theta} &= \underset{\Theta'}{\operatorname{argmin}} \int_{\mathcal{A}} (D_m(\mathbf{x}) - D_b(\mathcal{T}_{\Theta'}(\mathbf{x})))^2 d\mathbf{x} \\ &= \underset{\Theta'}{\operatorname{argmin}} \left[\int_{\mathcal{A}} D_m^2(\mathbf{x}) d\mathbf{x} + \int_{\mathcal{A}} D_b^2(\mathcal{T}_{\Theta'}(\mathbf{x})) d\mathbf{x} \right. \\ &\quad \left. - 2 \int_{\mathcal{A}} D_m(\mathbf{x}) D_b(\mathcal{T}_{\Theta'}(\mathbf{x})) d\mathbf{x} \right] \end{aligned}$$

Note that the first term above is fixed during optimization, while the second term is invariant under rigid transformation.

As a result, the above optimization is equivalent to maximizing the cross-correlation:

$$\hat{\Theta} = \operatorname{argmax}_{\Theta'} \int_{\mathcal{A}} D_m(\mathbf{x}) D_b(\mathcal{T}_{\Theta'}(\mathbf{x})) d\mathbf{x} \quad (1)$$

For differentiable D_m and D_b , the above optimization can be solved with gradient-based methods. However, the cross-correlation maximization problem in the urban AGV environment may have locally optimal solutions in the vicinity of the global minimum due to repetitive structure of radar reflectors. In applications with high integrity requirements, a search for the globally optimal solution is necessary. This paper notes that if the PHDs in (1) were to be discretized in \mathbf{x} , then the cross-correlation values can be evaluated exhaustively with computationally efficient techniques. Let \mathbf{x}_{pq} denote the location at the (p, q) translational offset in discretized \mathcal{A} . Then

$$\hat{\Theta} = \operatorname{argmax}_{\Theta'} \sum_{p=0}^{P-1} \sum_{q=0}^{Q-1} D_m(\mathbf{x}_{pq}) D_b(\lfloor \mathcal{T}_{\Theta'}(\mathbf{x}_{pq}) \rfloor) \quad (2)$$

where $\lfloor \cdot \rfloor$ denotes the nearest grid point in the discretized space.

The technique developed above relies on the PHDs D_m and D_b . The next subsections detail the recipe for estimating these PHDs from the radar observations.

4.2 Estimating the map PHD from measurements

This section addresses the procedure to estimate the map PHD $D_m(\mathbf{x})$ from radar measurements. This paper works with an occupancy grid map (OGM) approximation to the continuous PHD function. In [Erdinc et al., 2009], it has been shown that the PHD representation is a limiting case of the OGM as the grid cell size becomes vanishingly small. Intuitively, let c_{pq} denote the grid cell region with center \mathbf{x}_{pq} , and let δc_{pq} denote the area of this grid cell, which is small enough such that no more than one reflector may be found in any cell. Let $p_{pq}(O)$ denote the occupancy probability of c_{pq} , and let \mathcal{A} be defined as the region formed by the union of all c_{pq} whose centers \mathbf{x}_{pq} fall within \mathcal{A} . Then, the expected number of radar reflectors $\mathbb{E}[|\mathcal{A}|]$ in \mathcal{A} is given by

$$\begin{aligned} \mathbb{E}[|\mathcal{A}|] &= \sum_{c_{pq} \in \mathcal{A}} p_{pq}(O) = \sum_{c_{pq} \in \mathcal{A}} \frac{p_{pq}(O)}{\delta c_{pq}} \delta c_{pq} \\ &\triangleq \sum_{c_{pq} \in \mathcal{A}} \bar{D}(\mathbf{x}_{pq}) \delta c_{pq} \\ &= \int_{\mathcal{A}} \bar{D}(\mathbf{x}_{pq}) d\mathbf{x}, \quad \text{as } \lim_{\delta c_{pq} \rightarrow 0} \end{aligned}$$

where $\bar{D}(\mathbf{x}_{pq}) \equiv \frac{p_{pq}(O)}{\delta c_{pq}}$ can be considered to be an approximation of the PHD $D(\mathbf{x})$ for $\mathbf{x} \in c_{pq}$ since its integral over \mathcal{A} is equal to the expected number of reflectors in \mathcal{A} .

The advantage of working with an OGM approximation of the PHD is two-fold: first, since the OGM does not attempt to model individual objects, it is straightforward to represent arbitrarily-shaped objects, and second, in contrast to the “point target” measurement model assumption in standard PHD filtering, the OGM can straightforwardly model occlusions due to extended objects.

At this point, the task of estimating $D_m(\mathbf{x})$ has been reduced to estimating the occupancy probability of each grid cell in discretized \mathcal{A} . Each grid cell c_{pq} takes up one of two states: occupied (O) or free (F). Based on the radar measurement \mathbf{z}_k at each time k , the Bernoulli probability distribution of such binary state cells may be recursively updated with the binary Bayes filter. In particular, let $\mathbf{z}_{1:k}$ denote all radar measurements made up to time k , and let

$$l_{pq}^k(O) \equiv \log \frac{p_{pq}(O | \mathbf{z}_{1:k})}{1 - p_{pq}(O | \mathbf{z}_{1:k})} \quad (3)$$

denote the *log odds ratio* of c_{pq} being in state O . Also define $l_{pq}^0(O)$ as

$$l_{pq}^0(O) \equiv \log \frac{p_{pq}(O)}{1 - p_{pq}(O)}$$

with $p_{pq}(O)$ being the prior belief on the occupancy state of c_{pq} before any measurements are made. With these definitions, the binary Bayes filter update is given by [Thrun et al., 2005]

$$l_{pq}^k(O) = \log \frac{p_{pq}(O | \mathbf{z}_k)}{1 - p_{pq}(O | \mathbf{z}_k)} - l_{pq}^0(O) + l_{pq}^{k-1}(O) \quad (4)$$

where $p_{pq}(O | \mathbf{z}_k)$ is known as the *inverse* sensor model: it describes the probability of c_{pq} being in state O , given only the latest radar scan \mathbf{z}_k .

The required occupancy probability $p_{pq}(O | \mathbf{z}_{1:k})$ is easy to compute from the log odds ratio in (3). Observe that the inverse sensor model $p_{pq}(O | \mathbf{z}_k)$, in addition to the prior occupancy belief $p_{pq}(O)$, completely describes the procedure for estimating the OGM from radar measurements, and hence approximating the PHD. Adapting $p_{pq}(O | \mathbf{z}_k)$ to the characteristics of the automotive radar sensors, however, is not straightforward, and is discussed next.

4.3 Automotive Radar Inverse Sensor Model

This section addresses the challenge of adapting the inverse sensor model $p_{pq}(O \mid z_k)$ to the measurement characteristics of automotive radar sensors. Fig. 2 shows a simplified radar scan z_k of an underlying occupancy grid. For clarity of exposition, four distinct categories of grid cells in Fig. 2 are defined below:

- *Type A*: Grid cells in the vicinity of a radar range-azimuth return.
- *Type B*: Grid cells along the path between the radar sensor and *Type A* grid cells.
- *Type C*: Grid cells in the “viewshed” of the radar sensor, i.e., within the radar field-of-view and not shadowed by another object, but not of *Type A* or *Type B*.
- *Type D*: Grid cells outside the field-of-view of the radar (*Type D1*) or shadowed by other objects closer to the radar (*Type D2*).

The inverse sensor model must choose a $p_{pq}(O \mid z_k)$ value for each of these types of grid cells. In the following, the subscript pq is dropped for cleaner notation.




Figure 2: Schematic diagram showing four types of grid cells.

4.3.1 Conventional Choices for the Inverse Sensor Model

Since z_k provides no additional information on *Type D* grid cells, the occupancy in these cells is conditionally independent of z_k , that is

$$p^D(O | z_k) = p(O)$$

where $p(O)$ is the prior probability of occupancy defined earlier in Sec. 4.1.

Grid cells of *Type B* and *Type C* may be hypothesized to have low occupancy probability, since these grid cells were scanned by the sensor but no return was obtained. As a result, conventionally

$$p^B(O | z_k) \leq p(O)$$

and

$$p^C(O | z_k) \leq p(O)$$

Finally, grid cells of *Type A* may be hypothesized to have higher occupancy probability, since a return has been observed in the vicinity of these cells. Conventionally,

$$p^A(O | z_k) \geq p(O)$$

In the limit, if the OGM grid cell size is comparable to the sensor range and angle uncertainty, or if the number of scans is large enough such that the uncertainty is captured empirically, only the grid cells that contain the sensor measurement may be considered to be of *Type A*.

4.3.2 Automotive Radar Sensor Characteristics

Intense clutter properties and sparsity of the automotive radar data complicate the choice of the inverse sensor model.

Sparsity. First, sparsity of the radar scan implies that many occupied *Type A* grid cells in the radar environment might be incorrectly categorized as free *Type C* cells. This can be observed in Fig. 1. As evidenced by the batch of scans in Fig. 1d, the radar environment is “dense” in that many grid cells contain radar reflectors. However, any individual radar scan, such as the one shown in Fig. 1c, suggests a much more sparse radar environment. As a result, a grid cell which is occupied in truth will be incorrectly categorized as *Type C* in many scans, and correctly as *Type A* in a few scans.

The sparsity of radar returns also makes it challenging to distinguish *Type C* cells from cells of *Type D2*. Since many occluding obstacles are not detected in each scan, the occluded cells of *Type D2* are conflated with free cells of *Type C*.

In context of the inverse sensor model, as the radar scan becomes more sparse

$$p^C(O | z_k) \rightarrow p^D(O | z_k)^-$$

where the superscript $-$ denotes a limit approaching from below. Intuitively, approaching $p^D(O | z_k)$ implies that the measurement z_k is very sparse in comparison to the true occupancy, and thus does not provide much information on lack of occupancy.

Clutter. Second, there is the matter of clutter. The grid cells in the vicinity of a clutter measurement may be incorrectly categorized as *Type A*, and the grid cells along the path between the radar and clutter measurement may be incorrectly categorized as *Type B*.

In context of the inverse sensor model, as the radar scan becomes more cluttered

$$\begin{aligned} p^B(O | z_k) &\rightarrow p^D(O | z_k)^- \\ p^A(O | z_k) &\rightarrow p^D(O | z_k)^+ \end{aligned}$$

where the superscript $+$ denotes a limit approaching from above.

4.3.3 A Pessimistic Inverse Sensor Model

The results presented in Sec. 6 are based on a pessimistic sensor model, such that $p^B(O | z_k) = p^C(O | z_k) = p^D(O | z_k)$. This model assumes that the radar measurements provide no information about free space in the radar environment.

In particular, the inverse sensor model assumes

$$p^B(O | z_k) = p^C(O | z_k) = p^D(O | z_k) = p(O) = 0.1$$

and

$$p^A(O | z_k) = 0.2$$

4.4 Estimating the batch PHD from measurements

The procedure for generating an approximation to $D_b(\mathbf{x}')$ from a batch of radar measurements is identical to the procedure for generating $D_m(\mathbf{x})$ from mapping vehicle data, except that precise, absolute location and orientation data is not available during localization. Instead, pose estimates from the sensor fusion filter described in Sec. 5 are used to estimate the relative locations and orientations of each radar scan in the batch, and the scans are transformed into a common coordinate frame before updating the occupancy state of grid cells.

Once the map and batch PHDs have been approximated from radar measurements, the correlation-maximization technique developed in Sec. 4.1 can be applied to obtain the estimate $\hat{\Theta}$. This estimate is handed back to the sensor fusion filter as a pose offset measurement to constrain the odometric drift during absence of other sources of absolute localization, e.g., GNSS.

5 State Estimation with Sensor Fusion

Thus far, Sec. 4 has developed the theory and implementation of the radar batch correlation measurement, which provides an estimate $\hat{\Theta}$ of the 3-DoF (degrees-of-freedom) pose offset relative to the prior map. This section details a localization pipeline that incorporates the batch measurement update along with an array of other automotive all-weather sensing modalities to track the full 6-DoF vehicle pose trajectory. The high-rate pose estimates from this pipeline are also used to spatially organize individual scans to form the batch of radar scans used in the batch correlation update.

The choice of sensors available for all-weather localization is limited to radio-frequency sensors such as GNSS and automotive radars, and to proprioceptive sensors such as IMUs and wheel encoders. Any additional domain knowledge, such as properties of ground vehicle dynamics, may also be combined with these sensor measurements.

The localization pipeline in this paper is developed around a low-cost MEMS IMU. Fig. 3 shows a block diagram of the overall pipeline. The error-state multiplicative extended Kalman filter (EKF) makes use of cm-accurate CDGNSS position measurements whenever such measurements are available, e.g., in clear-sky GNSS environments. Radial velocity and bearing measurements from low-cost automotive radars are combined with nearly-zero sideslip and vertical speed constraints of a ground vehicle to continually track and limit the errors in inertial navigation. Smoothed batches of radar scans are correlated with a prior map to limit odometric position drift during CDGNSS outages. The following subsections outline the formulation of the estimator, the nonlinear state dynamics, the various measurement models, and the necessary calibration procedures.

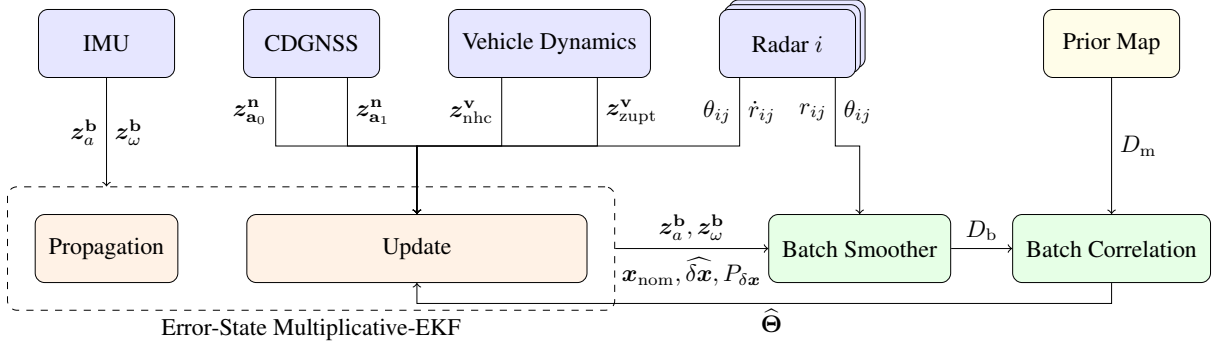


Figure 3: Block diagram of the localization pipeline. A low-cost MEMS IMU provides high-rate specific force and angular rate measurements. The error-state multiplicative extended Kalman filter (EKF) makes use of cm-accurate CDGNSS position measurements whenever such measurements are available, e.g., in clear-sky GNSS environments. Radial velocity and bearing measurements from low-cost automotive radars are combined with nearly-zero sideslip and vertical speed constraints of a ground vehicle to continually track and limit the errors in inertial navigation. Smoothed batches of radar scans are correlated with a prior map to limit odometric position drift during CDGNSS outages.

5.1 Sensor Platform & Coordinate Frames

To facilitate the discussion on measurement models and calibration, the sensor-instrumented vehicle and a few related coordinate frames are introduced here. An integrated perception platform called the *Sensorium*, shown schematically in Fig. 4, brings together the various low-cost automotive sensors considered in this paper. Many of these sensors provide measurements in their respective local frames, leading to a number of different coordinate frames that must be considered.

The IMU *body frame*, denoted \mathbf{b} , is the frame defined by the IMU’s accelerometer triad.

The *navigation frame*, denoted \mathbf{n} , is a local geographical reference frame, e.g., an ENU frame. The estimator wishes to track the pose trajectory of \mathbf{b} with respect to \mathbf{n} .

The *radar frames*, denoted \mathbf{r}_i for the i th radar, are local frames in which the radar sensors report range, range rate, and bearing to a number of targets.

The *vehicle frame*, denoted \mathbf{v} , is characterized by the direction in which the vehicle travels when the commanded steering angle is zero. This direction defines the y -axis of \mathbf{v} , as shown in Fig. 4. The origin of \mathbf{v} is located at the center of rotation of the vehicle.

The *Sensorium frame*, denoted \mathbf{s} , is defined by the physical structure of the Sensorium. It is essentially a convenience reference frame in which the nominal lever arm and orientation between different sensors are available per the mechanical specifications of the Sensorium. The origin of \mathbf{s} is arbitrarily chosen to be co-located with one of the GNSS




Figure 4: The University of Texas Sensorium is an integrated platform for automated and connected vehicle perception research. It includes three automotive radar units, one electronically-scanning radar (ESR) and two short-range radars (SRR2s); stereo visible light cameras; automotive- and industrial-grade inertial measurement units (IMUs); a dual-antenna, multi-frequency software-defined GNSS receiver; and an internal computer. An iXblue ATLANS-C CDGNSS-disciplined inertial navigation system (INS) (not shown) is mounted at the rear of the platform to provide the ground truth trajectory. The vehicle frame \mathbf{v} is located approximately at the center of the line connecting the rear axles.

antennas.

5.2 Error-State Filtering

The localization system of Fig. 3 estimates the following 16-element state vector:

$$\mathbf{x}_k = \left[\mathbf{p}_k^{\mathbf{n}\top}, \mathbf{v}_k^{\mathbf{n}\top}, \mathbf{q}_k^{\mathbf{nb}\top}, \mathbf{b}_{a,k}^{\mathbf{b}\top}, \mathbf{b}_{\omega,k}^{\mathbf{b}\top} \right]^\top$$

where $\mathbf{p}_k^{\mathbf{n}}$ is the 3×1 vector from \mathbf{n} to \mathbf{b} at time k expressed in \mathbf{n} , $\mathbf{v}_k^{\mathbf{n}}$ is the 3×1 velocity of \mathbf{b} relative to \mathbf{n} at time k expressed in the \mathbf{n} frame, $\mathbf{q}_k^{\mathbf{nb}}$ is the 4×1 unit quaternion that rotates a vector from \mathbf{b} to \mathbf{n} at time k , and $\mathbf{b}_{a,k}^{\mathbf{b}}$ and $\mathbf{b}_{\omega,k}^{\mathbf{b}}$ are the 3×1 accelerometer and gyroscope biases of the IMU at time k , expressed in \mathbf{b} .

Note that the vehicle orientation only has three effective degrees-of-freedom since $\mathbf{q}_k^{\mathbf{nb}}$ is constrained to be a unit quaternion. Enforcing such a constraint may result in a singular covariance matrix. This issue is typically dealt with an error-state filter [Sola, 2017] where the true state is split into a nominal-state vector

$$\mathbf{x}_{\text{nom},k} = \left[\hat{\mathbf{p}}_k^{\mathbf{n}\top}, \hat{\mathbf{v}}_k^{\mathbf{n}\top}, \hat{\mathbf{q}}_k^{\mathbf{nb}\top}, \hat{\mathbf{b}}_{a,k}^{\mathbf{b}\top}, \hat{\mathbf{b}}_{\omega,k}^{\mathbf{b}\top} \right]^\top$$

and an error-state vector $\delta \mathbf{x}_k$, related by the generalized addition operator \oplus as follows:

$$\mathbf{x}_k = \mathbf{x}_{\text{nom},k} \oplus \delta \mathbf{x}_k$$

where the error-state vector $\delta \mathbf{x}_k$ is the minimal 15-element state representation denoted component-wise as follows:

$$\delta \mathbf{x}_k = \left[\delta \mathbf{p}_k^{\mathbf{n}\top}, \delta \mathbf{v}_k^{\mathbf{n}\top}, \boldsymbol{\eta}_k^{\mathbf{n}\top}, \delta \mathbf{b}_{a,k}^{\mathbf{b}\top}, \delta \mathbf{b}_{\omega,k}^{\mathbf{b}\top} \right]^\top$$

The \oplus operator corresponds to usual vector addition for the position, velocity, and bias states. For the orientation state, \oplus is defined as

$$\begin{aligned} \mathbf{q}_k^{\mathbf{nb}} &= \tilde{\mathbf{q}}_k^{\mathbf{nb}} \oplus \boldsymbol{\eta}_k^{\mathbf{n}} \\ &= \exp_q \left(\frac{\boldsymbol{\eta}_k^{\mathbf{n}}}{2} \right) \odot \tilde{\mathbf{q}}_k^{\mathbf{nb}} \end{aligned}$$

where \exp_q denotes the exponential map from $\mathfrak{so}(3)$ to $SO(3)$ [Kok et al., 2017], represented as a quaternion, and \odot denotes quaternion multiplication. Note that $\boldsymbol{\eta}_k^{\mathbf{n}}$ is parametrized as an orientation deviation in \mathbf{n} . A similar formulation may be derived with the orientation deviation expressed in \mathbf{b} [Sola, 2017].

The nonlinear error-state is tracked with an error-state EKF. Owing to the multiplicative orientation dynamics and update, this filter is sometimes referred to as the multiplicative-EKF [Crassidis et al., 2007].

5.3 State Dynamics

Inertial measurements, collectively denoted \mathbf{u}_k , are interpreted as control inputs during the state propagation step. The true-state dynamics function $f_k(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k)$ is modeled as

$$\begin{aligned} \mathbf{p}_{k+1}^{\mathbf{n}} &= \mathbf{p}_k^{\mathbf{n}} + T \mathbf{v}_k^{\mathbf{n}} + \frac{T^2}{2} (R_k^{\mathbf{nb}} (\mathbf{z}_{a,k}^{\mathbf{b}} - \mathbf{b}_{a,k}^{\mathbf{b}} - \mathbf{w}_{a,k}^{\mathbf{b}}) + \mathbf{g}^{\mathbf{n}}) \\ \mathbf{v}_{k+1}^{\mathbf{n}} &= \mathbf{v}_k^{\mathbf{n}} + T (R_k^{\mathbf{nb}} (\mathbf{z}_{a,k}^{\mathbf{b}} - \mathbf{b}_{a,k}^{\mathbf{b}} - \mathbf{w}_{a,k}^{\mathbf{b}}) + \mathbf{g}^{\mathbf{n}}) \\ \mathbf{q}_{k+1}^{\mathbf{nb}} &= \mathbf{q}_k^{\mathbf{nb}} \odot \exp_q \left(\frac{T}{2} (\mathbf{z}_{\omega,k}^{\mathbf{b}} - \mathbf{b}_{\omega,k}^{\mathbf{b}} - R_k^{\mathbf{bn}} \boldsymbol{\omega}_e^{\mathbf{n}} - \mathbf{w}_{\omega,k}^{\mathbf{b}}) \right) \\ \mathbf{b}_{a,k+1}^{\mathbf{b}} &= \mathbf{b}_{a,k}^{\mathbf{b}} + \mathbf{w}_{b_a,k}^{\mathbf{b}} \\ \mathbf{b}_{\omega,k+1}^{\mathbf{b}} &= \mathbf{b}_{\omega,k}^{\mathbf{b}} + \mathbf{w}_{b_\omega,k}^{\mathbf{b}} \end{aligned}$$

where T is the propagation duration, R_k^{nb} is the rotation matrix representation of q_k^{nb} , $z_{a,k}^{\text{b}}$ and $z_{\omega,k}^{\text{b}}$ are the IMU specific force and angular rate measurements, respectively, $w_{a,k}$ and $w_{\omega,k}$ are the IMU specific force and angular rate white noise, respectively, $g^{\text{n}} \approx [0, 0, -9.8 \text{ m/s}^2]$ is the acceleration due to gravity after compensation for the centripetal force due to earth's rotation, and ω_e^{n} is the angular rate of the earth with respect to an inertial frame. The accelerometer and gyroscope biases are modeled as random walk processes driven by white noise $w_{b_a,k}^{\text{b}}$ and $w_{b_\omega,k}^{\text{b}}$, respectively, whose variances are derived from the IMU bias instability parameters [Woodman, 2007].

The nominal-state dynamics function $f_{\text{nom},k}(x_{\text{nom},k}, u_k, w_k)$ is similar to $f_k(x_k, u_k, w_k)$:

$$\begin{aligned}\tilde{p}_{k+1}^{\text{n}} &= \tilde{p}_k^{\text{n}} + T\tilde{v}_k^{\text{n}} + \frac{T^2}{2} \left(\tilde{R}_k^{\text{nb}} (z_{a,k}^{\text{b}} - \tilde{b}_{a,k}^{\text{b}}) + g^{\text{n}} \right) \\ \tilde{v}_{k+1}^{\text{n}} &= \tilde{v}_k^{\text{n}} + T \left(\tilde{R}_k^{\text{nb}} (z_{a,k}^{\text{b}} - \tilde{b}_{a,k}^{\text{b}}) + g^{\text{n}} \right) \\ \tilde{q}_{k+1}^{\text{nb}} &= \tilde{q}_k^{\text{nb}} \odot \exp_q \left(\frac{T}{2} (z_{\omega,k}^{\text{b}} - \tilde{b}_{\omega,k}^{\text{b}} - \tilde{R}_k^{\text{bn}} \omega_e^{\text{n}}) \right) \\ \tilde{b}_{a,k+1}^{\text{b}} &= \tilde{b}_{a,k}^{\text{b}} \\ \tilde{b}_{\omega,k+1}^{\text{b}} &= \tilde{b}_{\omega,k}^{\text{b}}\end{aligned}$$

The error-state dynamics function $f_{\text{err},k}(\delta x_k, u_k, w_k)$, is straightforwardly defined as

$$f_{\text{err},k} \triangleq f_k \ominus f_{\text{nom},k}$$

where \ominus denotes a generalized subtraction operator similar to \oplus defined earlier.

The linearized covariance propagation step of the EKF requires computation of the following Jacobians.

$$F_k = \left. \frac{\partial f_{\text{err},k}(\delta x_k, u_k, w_k)}{\partial \delta x_k} \right|_{\substack{\delta x_k=0 \\ w_k=0}} \quad (5)$$

$$G_k = \left. \frac{\partial f_{\text{err},k}(\delta x_k, u_k, w_k)}{\partial w_k} \right|_{\substack{\delta x_k=0 \\ w_k=0}} \quad (6)$$

This involves calculus of rotations. The interested reader is referred to [Sola, 2017, Kok et al., 2017] for further details.

The nontrivial sub-blocks of F_k and G_k are documented in Appendix A.

5.4 Measurement Models & Calibration

This section details the measurement models for the various measurements applied to the error-state EKF, along with the calibration procedures necessary for the application of these measurements.

5.4.1 Inertial Measurements

IMUs measure the specific force and angular rate experienced by \mathbf{b} relative to an inertial frame. If the centripetal force due to earth's rotation is absorbed in $\mathbf{g}^{\mathbf{n}}$, then the accelerometer and gyroscope measurements $z_{a,k}^{\mathbf{b}}$ and $z_{\omega,k}^{\mathbf{b}}$, respectively, are modeled as

$$\begin{aligned} z_{a,k}^{\mathbf{b}} &= R_k^{\mathbf{bn}}(\mathbf{a}_k^{\mathbf{n}} - \mathbf{g}^{\mathbf{n}}) + \mathbf{b}_{a,k}^{\mathbf{b}} + \mathbf{w}_{a,k}^{\mathbf{b}} \\ z_{\omega,k}^{\mathbf{b}} &= \boldsymbol{\omega}_k^{\mathbf{b}} + R_k^{\mathbf{bn}}\boldsymbol{\omega}_e^{\mathbf{n}} + \mathbf{b}_{\omega,k}^{\mathbf{b}} + \mathbf{w}_{\omega,k}^{\mathbf{b}} \end{aligned}$$

where $\mathbf{a}_k^{\mathbf{n}}$ is the true acceleration of the IMU in the \mathbf{n} frame, which double-integrates to position deviation, and $\boldsymbol{\omega}_k^{\mathbf{b}}$ is the true angular rate of the IMU in the \mathbf{n} frame, which integrates to orientation deviation. For low-quality IMUs, accelerometer and gyroscope scale factors may also need to be modeled. For the MEMS IMU used in this work, it was observed that modeling the scale factors did not yield any performance benefit.

The stochastic models for IMU white noise and random walk process are derived from the IMU specifications. In addition to such intrinsic calibration, extrinsic calibration of the IMU with respect to \mathbf{s} is necessary for the application of other measurements expressed in \mathbf{s} . The vector $\mathbf{p}_{\mathbf{s}\mathbf{b}}^{\mathbf{s}}$ from \mathbf{s} to \mathbf{b} is taken to be known from the mechanical specification since this is not strongly observable from the available measurements. It is, however, important to estimate any deviations from the mechanically specified orientation $\bar{\mathbf{q}}^{\mathbf{s}\mathbf{b}}$ between \mathbf{b} and \mathbf{s} , since even sub-degree errors in the IMU orientation relative to \mathbf{s} may lead to substantial errors when multiplied with the lever arm to another sensor.

The orientation deviation of $\bar{\mathbf{q}}^{\mathbf{s}\mathbf{b}}$ from truth, denoted $\boldsymbol{\eta}_{\mathbf{s}\mathbf{b}}^{\mathbf{s}}$, can be effectively estimated when CDGNSS measurements from multiple antennas are available to the EKF, as will be discussed in Sec. 5.4.2. Accordingly, the state vector $\delta\mathbf{x}_k$ is augmented with $\boldsymbol{\eta}_{\mathbf{s}\mathbf{b}}^{\mathbf{s}}$ during clear-sky periods. It must be noted, however, that since the IMU is mounted near the line connecting the Sensorium's two GNSS antennas, only two of the three elements in $\boldsymbol{\eta}_{\mathbf{s}\mathbf{b}}^{\mathbf{s}}$ are strongly observable. Any orientation deviation about the vector joining the two antennas is poorly unobservable, and must be constrained by construction. Also note that estimation of $\boldsymbol{\eta}_{\mathbf{s}\mathbf{b}}^{\mathbf{s}}$ only need be performed once as long as all sensors are rigidly mounted, and may not even be necessary if the mechanical tolerances are acceptably small.

5.4.2 CDGNSS Measurements

CDGNSS offers cm-accurate position measurements under all weather conditions, but typically offers reduced solution availability in deep urban environments. This paper takes the approach of incorporating CDGNSS measurements in the localization engine whenever they are available, while being capable of maintaining the required lane-level accuracy over long CDGNSS outages in deep urban canyons. In essence, the approach developed in this paper leverages CDGNSS for periodic or one-time intrinsic and extrinsic calibration of other on-board sensors, and relies on these sensors for accurate localization when CDGNSS is unavailable.

Signals captured from the two GNSS antennas on the Sensorium are processed together with those from a nearby reference station to provide nearly-independent three-dimensional position measurements of the antennas in the \mathbf{n} frame. The position measurement for antenna \mathbf{a}_i , $i \in \{0, 1\}$ is modeled as

$$\mathbf{z}_{\mathbf{a}_i, k}^{\mathbf{n}} = \mathbf{p}_k^{\mathbf{n}} + R_k^{\mathbf{nb}} R^{\mathbf{bs}} \mathbf{p}_{\mathbf{b}\mathbf{a}_i}^{\mathbf{s}} + \mathbf{e}_{\mathbf{a}_i, k} \quad (7)$$

where $\mathbf{e}_{\mathbf{a}_i, k}$ is the CDGNSS measurement noise. The vector $\mathbf{p}_{\mathbf{b}\mathbf{a}_i}^{\mathbf{s}}$ from \mathbf{b} to the antenna \mathbf{a}_i , expressed in \mathbf{s} , is available from the mechanical specification. As discussed above, $R^{\mathbf{bs}}$ may be taken to be the same as $\bar{R}^{\mathbf{bs}}$ from the mechanical specification, or may be further calibrated by augmenting the state with $\boldsymbol{\eta}_{\mathbf{sb}}^{\mathbf{s}}$.

Additionally, the error-state EKF requires the Jacobian of the measurement model with respect to the error state:

$$H_{\mathbf{a}_i, k} \triangleq \left. \frac{\partial \mathbf{z}_{\mathbf{a}_i, k}^{\mathbf{n}}}{\partial \delta \mathbf{x}_k} \right|_{\substack{\delta \mathbf{x}_k = 0 \\ \mathbf{e}_{\mathbf{a}_i, k} = 0}} = \left. \frac{\partial \mathbf{z}_{\mathbf{a}_i, k}^{\mathbf{n}}}{\partial \mathbf{x}_k} \right|_{\substack{\mathbf{x}_k = \mathbf{x}_{\text{nom}, k} \\ \mathbf{e}_{\mathbf{a}_i, k} = 0}} \cdot \left. \frac{\partial \mathbf{x}_k}{\partial \delta \mathbf{x}_k} \right|_{\substack{\delta \mathbf{x}_k = 0 \\ \mathbf{e}_{\mathbf{a}_i, k} = 0}}$$

The nontrivial sub-blocks of $H_{\mathbf{a}_i, k}$ are documented in Appendix A.

5.4.3 Radar Range Rate & Bearing Measurements

The range rate and bearing measurements from automotive radars provide a valuable velocity constraint for inertial navigation. Importantly, the frequency modulated continuous wave (FMCW) signal used in automotive radars provides instantaneous range rate measurements to the detected targets, i.e., target tracking and/or matching across cluttered radar scans is not necessary to obtain and apply this measurement.

The relative velocity of a stationary target with respect to \mathbf{r}_i is given by the negative of the velocity with respect to \mathbf{n} of the i th radar, expressed in \mathbf{r}_i , written $-\mathbf{v}_{\mathbf{r}_i, k}^{\mathbf{r}_i}$, as shown in Fig. 5. Assuming that the radar only detects targets in the




Figure 5: A visual description of the radar range rate measurement model. Quantities labeled in green are measured by the radar. The relative velocity of a stationary target with respect to \mathbf{r}_i is the negative of the velocity with respect to \mathbf{n} of the i th radar, expressed in \mathbf{r}_i , written $-\mathbf{v}_{\mathbf{r}_i}^{\mathbf{r}_i}$. The measured radial velocity \dot{r}_{ij} of the j th stationary target is the projection of $-\mathbf{v}_{\mathbf{r}_i}^{\mathbf{r}_i}$ onto the line-of-sight direction between the i th radar and the j th target.

two-dimensional plane of the linear phased array, the range rate measurement is modeled as

$$\dot{r}_{ij,k} = \begin{bmatrix} \sin \theta_{ij,k} \\ -\cos \theta_{ij,k} \\ 0 \end{bmatrix}^\top R^{\mathbf{r}_i \mathbf{s}} R^{\mathbf{s} \mathbf{b}} (R_k^{\mathbf{b} \mathbf{n}} \mathbf{v}_k^{\mathbf{n}} + (\boldsymbol{\omega}_k^{\mathbf{b}} \times R^{\mathbf{b} \mathbf{s}} \mathbf{p}_{\mathbf{br}_i}^{\mathbf{s}})) \quad (8)$$

where the vector $\mathbf{p}_{\mathbf{br}_i}^{\mathbf{s}}$ and the radar orientation $R^{\mathbf{r}_i \mathbf{s}}$ may be taken from the mechanical specifications. Note that unlike typical measurement models where the right-hand side is composed of quantities that are either known or are being estimated, 8 has *measured* quantities $\theta_{ij,k}$ on the right-hand side of the equation. This implies that any errors in the bearing measurements will not be accounted for if the range rate measurements are modeled in the EKF as shown.

The application of range rate constraints comes with two major challenges. First, individual radar scans contain a number of spurious targets as discussed in Sec. 1. Second, automotive phased-array radars exhibit poor bearing resolution and accuracy, and this is further exacerbated by the unusual range rate measurement model described above. Both of these challenges are addressed by pre-processing the range rate and bearing measurements with a RANSAC routine that estimates a best-fit two-dimensional radar velocity model to the radar measurements. In particular, with




Figure 6: Example results of the RANSAC operation on radar range rate and bearing measurements. The two yellow sinusoidal curves represent the RANSAC-predicted radial velocities for the port and starboard radars from Fig. 4 as a function of the bearing. With a threshold of 0.2 m/s, RANSAC considers violet dots as inliers and magenta dots as outliers. Note that the radial velocity magnitude is maximized at -30° and 30° for the port and starboard radars, respectively, in agreement with the mounting angles of these radars on the vehicle.

N detected targets, the RANSAC operation finds a robust solution to the following system of equations:

$$\begin{bmatrix} \dot{r}_{i0} \\ \vdots \\ \dot{r}_{iN} \end{bmatrix} = \begin{bmatrix} \sin \theta_{i0} & -\cos \theta_{i0} \\ \vdots & \vdots \\ \sin \theta_{iN} & -\cos \theta_{iN} \end{bmatrix} \begin{bmatrix} v_{\mathbf{r}_i, x}^{\mathbf{r}_i} \\ v_{\mathbf{r}_i, y}^{\mathbf{r}_i} \end{bmatrix} \quad (9)$$

while eliminating the $(\dot{r}_{ij}, \theta_{ij})$ pairs that may be outliers. Example results from the RANSAC procedure are shown in Fig. 6. Ultimately, the solution to 9 is applied as a measurement to the EKF with the following measurement model:

$$\begin{aligned} \mathbf{z}_{\mathbf{r}_i, k}^{\mathbf{r}_i} &\triangleq \begin{bmatrix} v_{\mathbf{r}_i, x}^{\mathbf{r}_i} \\ v_{\mathbf{r}_i, y}^{\mathbf{r}_i} \end{bmatrix}_k \\ &= [R^{\mathbf{r}_i \mathbf{s}} R^{\mathbf{s} \mathbf{b}} (R_k^{\mathbf{b} \mathbf{n}} \mathbf{v}_k^{\mathbf{n}} + (\boldsymbol{\omega}_k^{\mathbf{b}} \times R^{\mathbf{b} \mathbf{s}} \mathbf{p}_{\mathbf{b} \mathbf{r}_i}^{\mathbf{s}}))]_{[0,1]} + \mathbf{e}_{\mathbf{r}_i, k} \end{aligned}$$

where the subscript $[0, 1]$ denotes the first two elements of the three-element vector. Parts of the Jacobian of this measurement model with respect to the EKF error-state are documented in Appendix A.

5.4.4 Ground Vehicle Dynamics Constraints

Under nominal driving conditions, a ground vehicle respects dynamical constraints which can be leveraged as measurements to the EKF. This paper incorporates near-zero sideslip and vertical velocity constraints, commonly referred to as nonholonomic constraints (NHC), as well as zero-speed updates (ZUPT). The measurement models for these

constraints are described below.

Nonholonomic Constraints (NHC) Nonholonomic constraints have been previously studied for limiting IMU drift during GNSS outages, e.g., in [Dissanayake et al., 2001]. Typically, however, the near-zero sideslip and vertical velocity constraints are applied to the IMU body \mathbf{b} frame. This works well when the vehicle drives on straight-aways. During turns, however, the lateral velocity of the IMU cannot be assumed to be zero unless the IMU is serendipitously located at the vehicle center of rotation. This paper introduces a novel technique for online estimation of the vehicle center of rotation by taking advantage of periods of precise GNSS availability, as described below. Subsequently, the NHC are applied to this estimated center of rotation rather than at the IMU position, making the constraints valid during straight-aways as well as turns.

The application of NHC is based on the following assumptions:

1. There exists a fixed center of rotation, taken to be the origin of \mathbf{v} , about which the vehicle rotates when a steering control input is applied.
2. When a zero steering input is applied, the vehicle only moves in the \mathbf{v}_y direction. This holds by definition of \mathbf{v} .
3. The vehicle does not slip sideways or leave the surface of the road.

When the above assumptions hold, it follows that the velocity of the vehicle, when expressed in \mathbf{v} , is zero in the \mathbf{v}_x and \mathbf{v}_z directions at all times. In practice, however, these assumptions only hold approximately. Accordingly, the zero sideslip and vertical velocity constraints are applied as *soft* constraints in the form of measurements with an associated measurement error covariance. The NHC is modeled as

$$\mathbf{0}_{2 \times 1} \triangleq \mathbf{z}_{\text{nhc},k}^{\mathbf{v}} \quad (10)$$

$$\begin{aligned} &= [\mathbf{v}_k^{\mathbf{v}}]_{[0,2]} + \mathbf{e}_{\text{nhc},k} \\ &= [R^{\mathbf{vs}} R^{\mathbf{sb}} (R_k^{\mathbf{bn}} \mathbf{v}_k^{\mathbf{n}} + (\boldsymbol{\omega}_k^{\mathbf{b}} \times R^{\mathbf{bs}} \mathbf{p}_{\mathbf{bv}}^{\mathbf{s}}))]_{[0,2]} + \mathbf{e}_{\text{nhc},k} \end{aligned} \quad (11)$$

where $\mathbf{p}_{\mathbf{bv}}^{\mathbf{s}} = \mathbf{p}_{\mathbf{bs}}^{\mathbf{s}} + \mathbf{p}_{\mathbf{sv}}^{\mathbf{s}}$ and $R^{\mathbf{vs}}$ are parts of the extrinsic calibration between \mathbf{v} and \mathbf{s} . Precise manual measurement of $\mathbf{p}_{\mathbf{sv}}^{\mathbf{s}}$ and $R^{\mathbf{vs}}$ is challenging. First, it is not obvious where the origin of \mathbf{v} lies, though the center of line connecting the two rear axles might be a reasonable guess. Second, it would be challenging to measure, for example, the pitch of the Sensorium relative to the plane of the vehicle chassis. Accordingly, this paper takes a data-driven approach to extrinsic calibration of \mathbf{v} .

Once again, the extrinsic calibration technique relies on clear-sky periods with good CDGNSS availability, such that the nominal state estimates of \mathbf{v}_k^n , \mathbf{q}_k^{nb} , and $\mathbf{b}_{\omega,k}^{\text{b}}$ are close to their true values. Furthermore, calibration begins with coarse initial guesses of R^{vs} and $\mathbf{p}_{\text{sv}}^{\text{s}}$, denoted \bar{R}^{vs} and $\bar{\mathbf{p}}_{\text{sv}}^{\text{s}}$, respectively, and attempts to estimate the orientation deviation $\boldsymbol{\eta}_{\text{vs}}^{\text{s}}$ and lever arm deviation $\delta \mathbf{p}_{\text{sv}}^{\text{s}}$ with respect to these. With other quantities assumed known, 11 may be rewritten as

$$\begin{aligned} \mathbf{e}_{\text{nhc},k} &= \left[(\bar{R}^{\text{vs}} \oplus \boldsymbol{\eta}_{\text{vs}}^{\text{s}}) (\mathbf{v}_k^{\text{s}} + (\boldsymbol{\omega}_k^{\text{s}} \times (\bar{\mathbf{p}}_{\text{bv}}^{\text{s}} + \delta \mathbf{p}_{\text{bv}}^{\text{s}}))) \right]_{[0,2]} \\ &\triangleq \mathbf{h}_{\text{nhc},k}(\boldsymbol{\eta}_{\text{vs}}^{\text{s}}, \delta \mathbf{p}_{\text{bv}}^{\text{s}}) \end{aligned}$$

This model is nonlinear in $\boldsymbol{\eta}_{\text{vs}}^{\text{s}}$, and may be solved as a nonlinear least squares problem, e.g., with the Gauss-Newton method. The Jacobian of $\mathbf{h}_{\text{nhc},k}$ evaluated at $\boldsymbol{\eta}_{\text{vs}}^{\text{s}} = \mathbf{0}$ and $\delta \mathbf{p}_{\text{bv}}^{\text{s}} = \mathbf{0}$ is composed of

$$\begin{aligned} \frac{\partial \mathbf{h}_{\text{nhc},k}}{\partial \boldsymbol{\eta}_{\text{vs}}^{\text{s}}} &= \left[(\mathbf{v}_k^{\text{s}} + \boldsymbol{\omega}_k^{\text{s}} \times \bar{\mathbf{p}}_{\text{bv}}^{\text{s}})^{\top} \otimes [\bar{R}^{\text{vs}}]_{[(0,2),(\cdot)]} \right] \begin{bmatrix} [-\hat{\mathbf{i}}]_{\times} \\ [-\hat{\mathbf{j}}]_{\times} \\ [-\hat{\mathbf{k}}]_{\times} \end{bmatrix} \\ \frac{\partial \mathbf{h}_{\text{nhc},k}}{\partial \delta \mathbf{p}_{\text{bv}}^{\text{s}}} &= [\bar{R}^{\text{vs}}]_{[(0,2),(\cdot)]} [\boldsymbol{\omega}_k^{\text{s}}]_{\times} \end{aligned}$$

where \otimes denotes the Kronecker product, subscript $[(0,2),(\cdot)]$ denotes selection of the first and third rows of a matrix, $[\cdot]_{\times}$ denotes the skew-symmetric cross-product matrix corresponding to the 3-element argument, and $\hat{\mathbf{i}}$, $\hat{\mathbf{j}}$, and $\hat{\mathbf{k}}$ denote the cardinal unit vectors. To make the system observable, measurements from multiple epochs must be stacked and solved as a batch. Additionally, the nonlinear problem must be iteratively linearized and solved until convergence.

Zero-Speed Update (ZUPT) The ZUPT constraint is another valuable measurement that limits odometric drift, especially in situations where the platform makes frequent stops. The measurement model for ZUPT is trivially written as

$$\begin{aligned} \mathbf{0}_{3 \times 1} &\triangleq \mathbf{z}_{\text{zupt},k}^{\text{v}} \\ &= R^{\text{vs}} R^{\text{sb}} R_k^{\text{bn}} \mathbf{v}_k^{\text{n}} + \mathbf{e}_{\text{zupt},k} \end{aligned} \tag{12}$$

The primary challenge of applying ZUPT is detection of epochs where this constraint is valid. Importantly, this condition must be detected independently from the EKF state estimate, e.g., by inspection of the raw IMU measurements. In theory, it is not possible to make any claims about zero speed based on acceleration and/or angular rate data, since

IMU measurements of a vehicle moving with a constant velocity and orientation must be indistinguishable from those of a stationary vehicle. In practice, however, the IMU measurements exhibit a distinct behavior when the vehicle is in motion, e.g., due to road roughness and vehicle vibrations. Prior work has made use of these *artifacts* to detect stationary periods. This paper follows the angular rate energy method from [Skog et al., 2010] for ZUPT detection. In practice, if wheel odometry data are available from the vehicle CAN bus, as is common in most modern vehicles, then ZUPT detection can be performed trivially and with high reliability.

An observant reader might wonder why ZUPT is not applied directly to v_k^n in 12. The advantage of applying ZUPT in v is that a tighter zero-speed constraint can be reliably applied in the lateral and vertical directions.

5.5 Batch Smoothing & Update

Real-time estimates of the vehicle pose trajectory obtained from the EKF may be used to string together individual scans and perform a radar batch measurement update. However, since these data are processed batches, it is desirable to perform backward smoothing over the short duration of the batch. Backward smoothing enforces the dynamics function backwards in time, ironing out any large jumps that may have occurred in the EKF forward pass.

Accordingly, the batch smoother component in Fig. 3 stacks all inertial measurements and snapshots of the estimator state over the duration of the batch. When the batch is ready to be processed for correlation, backward smoothing is enforced with the inertial measurements as control inputs. The smoothing formulation in this case is somewhat more complicated than usual [Särkkä, 2013] due to nonlinear backward dynamics and the error-state formulation. Details on nonlinear error-state Rauch-Tung-Striebel smoothing are provided in Appendix B.

The correlation peak search region is taken to be ± 5 m and $\pm 3^\circ$. The 3-DoF pose offset $\hat{\Theta}$ from radar batch correlation is applied as horizontal position and heading measurements to the EKF. Outliers from batch correlation are excluded in the EKF based on a χ^2 -test on the normalized innovation squared (NIS) [Bar-Shalom et al., 2001].

6 Experimental Results

The radar-inertial positioning system of Fig. 3 was evaluated experimentally using the dataset described in [Narula et al., 2020b], collected during approximately 1.5 h of driving on two separate days in and around the urban center of Austin, TX. This section presents the evaluation results.

6.1 Dataset

Fig. 7 shows the route followed by the sensor-instrumented vehicle on Thursday, May 9, 2019 (in blue) and Sunday, May 12, 2019 (in red). The test route combs through every street in the Austin, TX downtown area, since such environments are the most challenging for CDGNSS-based positioning [Humphreys et al., 2020] and would benefit the most from multi-sensor all-weather positioning. The route was driven once on a weekday and again on the weekend to evaluate robustness of the proposed map-based approach to changes in the traffic and parking patterns. The reader may refer to [Narula et al., 2020b, Fig. 5] for a visual description of the radar and GNSS environment navigated in this test route.

It must be noted that these data were collected in clear weather conditions. The results presented in this section are nevertheless expected to extend to adverse weather conditions, as detailed in Sec. 2. Importantly, this paper does not claim that camera- or lidar-based localization would fail in the conditions in which these data were collected. Instead, the focus of this work is on the development of a system that is robust to the sparse and cluttered automotive radar measurements, under the assumption that the radar measurements do not significantly degrade in adverse weather.




Figure 7: Test route through The University of Texas west campus and Austin downtown. These areas are the most challenging for precise GNSS-based positioning and thus would benefit the most from radar-based positioning. The route was driven once on a weekday and again on the weekend to evaluate robustness of the radar map to changes in traffic and parking patterns. Red is the mapping run (May 12), blue is the localization run (May 9). Note that a short part of the route in the north-west section, as well as the final part of the route in the north-east section, were different during the mapping and localization runs (i.e., the red and blue trajectories do not overlap) due to street closures, preventing the use of a map-based positioning approach. These sections of the test route have been omitted from the evaluation results.

6.1.1 Sensors

The Sensorium, shown in Fig. 4, features two types of automotive radars: one Delphi electronically-scanning radar (ESR) in the middle and two Delphi short-range radars (SRR2s) on the two sides. Both the ESR and the SRR2 are commercially available; similar radars are available on economy-class consumer vehicles. The ESR provides simultaneous sensing in a narrow ($\pm 10^\circ$) long-range (175 m) coverage area and a wider ($\pm 45^\circ$) medium-range (60 m) area. The SRR2 units each have a coverage area of $\pm 75^\circ$ and 80 m (see [Narula et al., 2020a, Fig. 6]). Each SRR2 is installed facing outward from the center-line at an angle of 30° . The Sensorium’s onboard computer timestamps and logs the radar returns from the three radar units.

The LORD MicroStrain 3DM-GX5-25 MEMS IMU is an industrial-grade inertial sensor that acts as the core sensor of the localization pipeline. The IMU provides temperature-compensated accelerometer and gyroscope readings at 100 Hz. Two Antcom G8Ant-3A4TNB1 high performance GNSS patch antennas pull in signals from all three GNSS frequency bands and include a 40 dB active low-noise amplifier.

6.1.2 Ground-Truth Trajectory

The ground-truth position and orientation trajectory for the data are generated with the iXblue ATLANS-C, a high-performance CDGNSS coupled fiber-optic gyroscope INS. The post-processed position solution obtained from the ATLANS-C is decimeter-accurate throughout the dataset.

6.1.3 Dataset Splits

With a limited amount of field data available for development and evaluation, it is critical to ensure that the proposed positioning technique does not overfit this particular dataset. Accordingly, the data used in the development of the algorithms were restricted to a fixed 30 min segment, where the prior radar map was constructed with radar measurements from May 9 and localization was performed with radar, inertial, and CDGNSS measurements from May 12. In contrast, during evaluation the full 62 min of data were used, and the mapping and localization datasets were inverted, i.e., the prior map was constructed with radar measurements from May 12, and localization was performed with all sensor data from May 9. The algorithms have not been modified to maximize the performance over the evaluation dataset.

6.2 Prior Radar Mapping

The first step to radar-map-based localization is the generation of a radar map point cloud. Radar scans collected from the May 12, 2019 drive were aggregated to create a map with the benefit of the ATLANS-C ground-truth trajectory. In a practical system, the radar map may be generated during favorable conditions for optical sensors such as cameras and lidar, such that the mapping vehicle can accurately track its pose. Additionally, the mapping process may be crowd-sourced from consumer vehicles [Narula et al., 2018a, Narula et al., 2018b]. The map point cloud is stored in a k-d tree for efficient querying during localization.

Two implementation notes are in order here. First, automotive radar clutter is especially intense when the vehicle is stationary. Accordingly, radar range measurements obtained when the vehicle was moving slower than 1 m/s were discarded for both mapping and localization. This implies that radar correlation measurements were only available during periods when the vehicle was moving faster than 1 m/s. Second, it was observed that radar returns far from the vehicle are mostly clutter and have negligible resemblance to the surrounding structure. Radar returns with range larger than 50 m were discarded for both the map and batch PHDs. It is noted that these two parameters have not been optimized to produce the smallest estimation errors; instead they have been fixed based on visual inspection.

6.3 Offline Calibration

Extrinsic calibration among the IMU frame \mathbf{b} , the Sensorium frame \mathbf{s} , and the vehicle frame \mathbf{v} was performed offline with 125 s of sensor data with CDGNSS availability. While it is possible to estimate the calibration parameters online, it may not be desirable to do so if these parameters are not expected to change over time.

The orientation deviation $\eta_{\mathbf{s}\mathbf{b}}^{\mathbf{s}}$ between the IMU body frame and the Sensorium frame was calibrated for the localization dataset, as described in Sec. 5.4.1. With two GNSS antennas, only two out of the three DoFs in $\eta_{\mathbf{s}\mathbf{b}}^{\mathbf{s}}$ are observable. Accordingly, the orientation deviation around \mathbf{b}_x , which is mostly unobservable, was tightly constrained to the initial guess of zero. The deviations around \mathbf{b}_y and \mathbf{b}_z rapidly converged to sub-degree offsets from the mechanical specification.

Extrinsic calibration between \mathbf{v} and \mathbf{s} was similarly estimated over the 125 s period as detailed in Sec. 5.4.4.

The commercial automotive radars on the Sensorium do not offer any mechanism to synchronize their scans with an external reference clock. Analysis of the radar range rate residuals in the EKF showed clear evidence of latency between the data logging timestamp and the true scan times. Accordingly, radar latency calibration was performed offline with a best fit approach.

6.4 Implementation Notes

A few implementation- and dataset-specific notes relating to the localization pipeline are documented below.

CDGNSS Measurements & Outages The CDGNSS position measurements used in this evaluation are in fact the output of the post-processed ground-truth system, i.e., these measurements have not been obtained from an unaided CDGNSS receiver. While this is not ideal for realistic evaluation, the evaluation results presented herein do not mislead because, first, CDGNSS measurements are only applied for a 125 s period for initial calibration, and second, any commercial CDGNSS receiver would be able to generate similar cm-accurate position solutions in the clear-sky region where the CDGNSS measurements were applied.

Measurement Noise Correlation Observations from the field data revealed that the measurement noise in the radar range rate measurements is not independent between consecutive radar scans. This is problematic since the EKF applied assumes each measurement to have errors that are uncorrelated in time. Accordingly, the radar range rate measurements were decimated to 1 Hz such that the measurements were spaced out by roughly the decorrelation time of the measurement noise. A more principled approach to this problem is to augment the state vector with states to pre-whiten the measurements. But this approach was empirically observed to not outperform the straightforward measurement decimation, while introducing additional complexity and tuning parameters.

Similarly, the NHC and ZUPT measurements can in theory be applied at every applicable IMU epoch. But to prevent correlated errors in these constraints (e.g., due to sideslip experienced while cornering) from making the EKF inconsistent, they are only applied at 1 Hz.

Filter Tuning Parameters The process noise covariance used in the EKF is derived from the IMU datasheet parameters [Woodman, 2007, [LORD Sensing MicroStrain](#),]. The measurement noise covariance associated with CDGNSS measurements is available directly from the ATLANS-C receiver. A few other measurement noise standard deviations and tuning parameters are documented in Table 1.

6.5 Localization Results

This section presents empirical error statistics obtained from field evaluation of the proposed method. The test scenario evaluated in this section is an extreme one: the vehicle starts off in a clear-sky environment with 125 s of CDGNSS availability, and subsequently all CDGNSS measurements are cut off for the next 3600 s of driving, during which the system must rely on radar and inertial sensing along with vehicle dynamical constraints to maintain an accurate

Table 1: A List of Parameters Involved in the Localization Pipeline

Minimum speed for valid radar range	1 m/s
Maximum valid radar range	50 m
Minimum RANSAC inliers	10
Minimum fraction of RANSAC inliers	0.65
$v_{\mathbf{r}_i, x}^{\mathbf{r}_i}$ (broadside) standard deviation	0.2 m/s
$v_{\mathbf{r}_i, y}^{\mathbf{r}_i}$ (boresight) standard deviation	0.1 m/s
$v_{\text{nhc}, x}^{\mathbf{v}}$ (lateral) standard deviation	0.1 m/s
$v_{\text{nhc}, z}^{\mathbf{v}}$ (vertical) standard deviation	0.2 m/s

estimate of its pose.




Figure 8: This figure shows an interesting example of radar-based urban positioning with the proposed method. Panel (a) shows the occupancy grid estimated from the prior map point cloud. Panel (b) shows the same for a 5 s batch of scans collected in the same region. For ease of visualization, the batch occupancy grid has already been aligned with the map occupancy grid. Panel (c) shows the cross-correlation between the batch and map occupancy grids at $\Delta\phi = 0^\circ$. Given that no rotational or translational offset error has been applied to the batch, the correlation peak should appear at $(0, 0)$. The offset of the peak in panel (c) from $(0, 0)$ is the translational estimate error of the proposed method. Also note the increased positioning uncertainty in the along-track direction, and the two local correlation peaks (marked with red squares in panel (c)) due to the repeating periodic pattern of radar reflectors in the map and the batch (marked with red rectangles in panels (a) and (b)).

Before diving into the quantitative analysis, it is interesting to inspect the example of a radar batch update shown in Fig. 8. For ease of visualization, the batch point cloud to be localized has already been adjusted for any translational or rotational offset from the ground truth. The occupancy grid estimated from the 5 s batch of scans is shown in Fig. 8b. Similarly, Fig. 8a shows the occupancy grid estimated from the map point cloud retrieved from the map database. Fig. 8c shows the cross-correlation between the batch and map occupancy grids. Given that the batch is already aligned with ground truth, one should expect the correlation peak to appear at $(0, 0)$ in Fig. 8c. The offset of the peak from $(0, 0)$ in this case would be the translational estimate error.

Two interesting features of the cross-correlation in Fig. 8c are worth noting. First, the correlation peak decays slower in the along-track direction—in this case approximately aligned with the south-southwest direction. This is a general feature observed throughout the dataset, since most of the radar reflectors are aligned along the sides of the streets. Second, there emerge two local correlation peaks offset by ≈ 4 m along the direction of travel. These local peaks are due to the repeating periodic structure of radar reflectors in both the map and the batch occupancy grids. In other words, shifting the batch occupancy grid forward or backward along the vehicle trajectory by ≈ 4 m aligns the periodically-repeating reflectors in an off-by-one manner, leading to another plausible solution. Importantly, the uncertainty envelope of the initial position estimate can span several meters, encompassing both the global optimum and one or more local optima. This explains why gradient-based methods, which seek the nearest optimum, are poorly suited for use in the urban automotive radar environment.

6.5.1 Performance with 4 s Radar Batches

Fig. 9 shows the east and north position error time histories from the test scenario described above. For the results presented in Fig. 9 and 10, a 4 s radar batch duration is chosen. In the first 125 s of clear-sky conditions with CDGNSS availability, the east and north position errors with respect to the ground truth are sub-decimeter, as expected. Over the subsequent 60 min of driving in and around the urban center of the city, the proposed method maintains sub-35-cm horizontal position errors (95%). The horizontal position estimation errors are consistent with the predicted standard deviation from the EKF. This is a remarkable result which shows that, given a prior radar map, lane-level-accurate horizontal positioning is achievable under GNSS-denied conditions with the types of all-weather sensors that are already available on commercial vehicles. Vertical position errors are not shown in Fig. 9 since these are not constrained by the two-dimensional radar batch correlation update. For ground vehicle applications, a digital elevation map can effectively constrain errors in altitude, if necessary.

Vehicle orientation estimation errors for the same scenario are shown in Fig. 10. Heading estimation error, shown in the bottom panel, is most important for ground vehicle applications. The proposed technique maintains vehicle heading estimates to within 0.5° of the ground truth throughout most of the dataset, and the errors are consistent with the predicted uncertainty. Roll and pitch estimation errors are smaller and stay within 0.2° of the ground truth. Better estimation of roll and pitch is expected since these are directly observable with the accelerometer measurements. The same phenomenon explains the substantially shorter decorrelation times for roll and pitch errors as compared to the heading error. Finally, it is noted that the EKF is mildly inconsistent in regards to roll and pitch estimation errors. This suggests that the accelerometer white noise and bias stability characteristics claimed in the IMU datasheet [LORD Sensing MicroStrain,] may be optimistic in field application.




Figure 9: East and north position error time histories from field evaluation. In the first 125 s of clear-sky conditions with CDGNSS availability, the east and north position errors with respect to the ground truth are sub-decimeter, as expected. Over the subsequent 60 min of driving in and around the urban center of the city, the proposed method maintains sub-35-cm (95%) horizontal position errors. The horizontal position estimation errors are consistent with the predicted standard deviation from the EKF.

Of the related work reported in Sec. 3, results in [Yoneda et al., 2018, Ward and Folkesson, 2016, Lundgren et al., 2014] were reported on a similar scenario to this paper, where a prior radar map is generated based on a ground-truth pose estimation system, followed by localization within the prior radar map. While none of the above methods have an open-source implementation for direct comparison on this paper’s dataset, it is possible to compare the positioning accuracy results reported in each work. Even though the method in [Yoneda et al., 2018] assumes perfect, hypothetical local odometry as well as heading to create the radar batches, an RMS horizontal positioning error of 25 cm is reported in the most favorable conditions, which is worse than the 35 cm 95-percentile accuracy of the end-to-end sensor fusion pipeline of this paper. Without camera-based lane marking measurements, [Lundgren et al., 2014, Table II] reports 10 % longitudinal positioning errors larger than 1 m, and 77 % lateral positioning errors larger than 20 cm, even though [Lundgren et al., 2014] makes use of GPS, wheel speed sensors, and a gyroscope. The current paper significantly outperforms [Lundgren et al., 2014] with no GNSS or wheel speed measurements. Similarly, [Ward and Folkesson, 2016] uses speed and yaw-rate from a ground-truth reference system along with a prior radar map to achieve RMS longitudinal positioning error of 37.7 cm, which is worse than the 95-percentile horizontal positioning error reported in this paper without any assistance from the ground-truth reference system.

6.5.2 Choosing a Radar Batch Length

The problem of choosing the duration of a radar batch during localization presents an interesting trade-off. On the one hand, longer batch durations are preferable because, intuitively, cross-correlation using a larger *patch* of the radar environment is more likely to produce a strong and unambiguous correlation peak. Fig. 11 shows results from an




Figure 10: Vehicle orientation estimation errors from field evaluation. The proposed technique maintains vehicle heading estimates to within 0.5° of the ground truth throughout most of the dataset, and the errors are consistent with the predicted uncertainty. Roll and pitch estimation errors are smaller and stay within 0.2° of the ground truth.




Figure 11: CCDFs for different batch lengths between 1 s and 8 s. The 50-percentile errors are similar for shorter and longer batch lengths, but the difference becomes more noticeable at higher percentiles.

empirical test of this intuition. In this test, radar batches of different durations between 1 s and 8 s were generated with ground-truth odometry and correlated against a prior map to obtain the estimated offset from the ground-truth pose. The complementary cumulative distribution function (CCDF) of the horizontal position estimation errors is shown in 11. It is interesting to note that up to the 70th percentile, errors are similar for different batch lengths. The difference between the CCDFs becomes more pronounced at higher percentiles, implying that errors for shorter batch lengths have heavy tails. Recall that in the overall localization pipeline of Fig. 3, these errors will act as measurement errors in $\hat{\Theta}$. An EKF models measurement errors to be Gaussian, which is not a good model for heavy-tailed distributions. Accordingly, longer batch durations would appear preferable.

On the other hand, longer batches have several disadvantages. First, longer durations between batch measurement updates leads to larger odometric drift *between* updates, as well as poorer reconstruction of the radar batch itself.

Second, some of the worst outliers due to shorter batch lengths may be rejected in the EKF based on the χ^2 NIS test, thus blunting the relative advantage of longer batches. Shorter batch lengths allow for a larger number of measurement updates to be performed per unit time, even if a few of those measurements may have to be rejected as outliers.

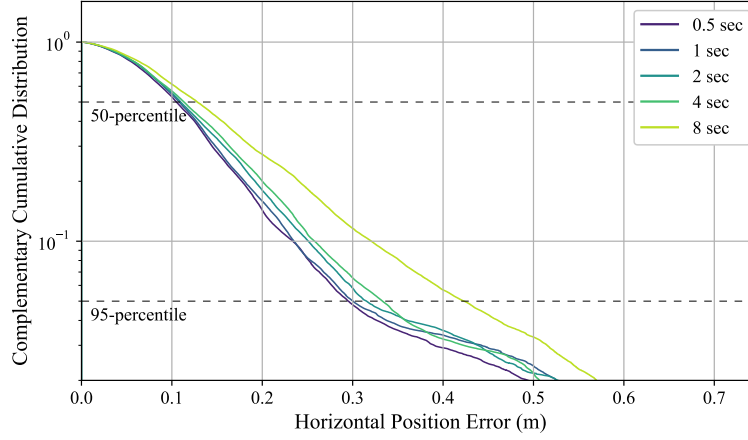


Figure 12: End-to-end effect of different batch lengths on horizontal positioning performance. Other than the longest batch length of 8 s, most batch lengths appear to perform similarly well, with 95th-percentile horizontal position errors near 30 cm.

Fig. 12 reveals the end-to-end effect of different batch lengths. For a given batch length, its measurement error standard deviation was obtained from the corresponding CCDF in Fig. 11, i.e., the $\hat{\Theta}$ measurement standard deviation is smaller for longer batches. Interestingly, other than the longest batch length of 8 s, most batch lengths appear to perform similarly well, with 95-percentile horizontal position errors near 30 cm. Given the heavy-tailed nature of measurement noise distributions when working with very short batches (from Fig. 11), batch lengths from 2 to 4 s may be taken to be a good compromise.

6.5.3 Sensitivity to Map Discrepancies

As it includes mapping and localization data collected on different days—one on a weekday and the other on a weekend—the dataset processed in this work allows evaluation of the robustness of the system to discrepancies between the prior map and the radar batch. Note that these data were collected prior to the COVID-19 pandemic and exhibit significant variation in traffic and parking patterns between the weekday and weekend collection, especially in the university and downtown area. Street parking, in particular, is drastically different between the prior map and the localization data. Even so, as described above, the proposed system exhibits remarkable accuracy and robustness. It must be noted, however, that the correlation-based approach of this paper assumes sufficient consistency between the prior map and the current batch such that the cross-correlation is maximized for the true translation and heading offset within the peak search window. The system would fail in the extreme case that a preponderance of features in

the surrounding environment changed between mapping and localization.

7 Conclusion

A robust pipeline for all-weather sub-50-cm urban ground vehicle positioning has been proposed. The positioning engine is based on commercially-available low-cost automotive radars, MEMS IMU, ground vehicle dynamics constraints, and, when available, precise GNSS measurements. Remarkably, it has been shown that given a prior radar map, lane-level-accurate horizontal positioning is achievable with the types of all-weather sensors that are already available on commercial vehicles. In comparison with a post-processed ground truth trajectory, it was shown that during 60 min of GNSS-denied driving in the urban center of Austin, TX during clear weather, the proposed pipeline has 95th-percentile errors of 35 cm in horizontal position and 0.5° in heading. Based on several decades of empirical studies of millimeter-wave radar in rain, snow, fog, and dust, these results are expected to extend to adverse weather conditions without significant degradation. This is a significant development in the field of AGV localization, which has traditionally been based on sensors such as lidar and cameras that perform poorly in bad weather conditions.

A Partial Derivatives

A.1 Linearized Forward Dynamics

A few block components of F_k and G_k from (5) and (6) are listed below.

$$\begin{aligned}\frac{\partial \delta \mathbf{p}_{k+1}^{\mathbf{n}}}{\partial \boldsymbol{\eta}_k^{\mathbf{n}}} \Big|_{\substack{\delta \mathbf{x}_k=0 \\ \mathbf{w}_k=0}} &= -\frac{T^2}{2} \left(\tilde{R}_k^{\mathbf{nb}} \left[\mathbf{z}_{a,k}^{\mathbf{b}} - \tilde{\mathbf{b}}_{a,k}^{\mathbf{b}} \right]_{\times} \right) \\ \frac{\partial \delta \mathbf{p}_{k+1}^{\mathbf{n}}}{\partial \delta \mathbf{b}_{a,k}^{\mathbf{b}}} \Big|_{\substack{\delta \mathbf{x}_k=0 \\ \mathbf{w}_k=0}} &= -\frac{T^2}{2} \tilde{R}_k^{\mathbf{nb}}\end{aligned}$$

The partial derivatives of $\delta \mathbf{v}_{k+1}^{\mathbf{n}}$ with respect to $\delta \mathbf{x}_k$ follow similarly.

$$\begin{aligned}\frac{\partial \boldsymbol{\eta}_{k+1}^{\mathbf{n}}}{\partial \boldsymbol{\eta}_k^{\mathbf{n}}} \Big|_{\substack{\delta \mathbf{x}_k=0 \\ \mathbf{w}_k=0}} &\approx I_{3 \times 3} \\ \frac{\partial \boldsymbol{\eta}_{k+1}^{\mathbf{n}}}{\partial \delta \mathbf{b}_{\omega,k}^{\mathbf{b}}} \Big|_{\substack{\delta \mathbf{x}_k=0 \\ \mathbf{w}_k=0}} &\approx -T \tilde{R}_{k+1}^{\mathbf{nb}} \mathbf{J}_r \left(\frac{T}{2} \left(\mathbf{z}_{\omega,k}^{\mathbf{b}} - \tilde{\mathbf{b}}_{\omega,k}^{\mathbf{b}} - \tilde{R}_k^{\mathbf{bn}} \boldsymbol{\omega}_e^{\mathbf{n}} \right) \right) \\ &\approx -T \tilde{R}_{k+1}^{\mathbf{nb}}\end{aligned}$$

where

$$\mathbf{J}_r(\boldsymbol{\theta}) = \mathbf{I}_{3 \times 3} - \frac{1 - \cos \|\boldsymbol{\theta}\|}{\|\boldsymbol{\theta}\|^2} [\boldsymbol{\theta}]_{\times} + \frac{\|\boldsymbol{\theta}\| - \sin \|\boldsymbol{\theta}\|}{\|\boldsymbol{\theta}\|^3} [\boldsymbol{\theta}]_{\times}^2$$

is the right Jacobian of $SO(3)$ [Sola, 2017].

$$\begin{aligned} \left. \frac{\partial \boldsymbol{\eta}_{k+1}^{\mathbf{n}}}{\partial \mathbf{w}_{\omega,k}} \right|_{\substack{\delta \mathbf{x}_k=0 \\ \mathbf{w}_k=0}} &\approx T \tilde{R}_{k+1}^{\mathbf{nb}} \mathbf{J}_r \left(\frac{T}{2} \left(\mathbf{z}_{\omega,k}^{\mathbf{b}} - \tilde{\mathbf{b}}_{\omega,k}^{\mathbf{b}} - \tilde{R}_k^{\mathbf{bn}} \boldsymbol{\omega}_{\mathbf{e}}^{\mathbf{n}} \right) \right) \\ &\approx T \tilde{R}_{k+1}^{\mathbf{nb}} \end{aligned}$$

A.2 Linearized Measurement Models

The partial derivative of the measurement $\mathbf{z}_{\mathbf{a}_i,k}^{\mathbf{n}}$ from (7) can be expressed as

$$\left. \frac{\partial \mathbf{z}_{\mathbf{a}_i,k}^{\mathbf{n}}}{\partial \delta \mathbf{x}_k} \right|_{\substack{\delta \mathbf{x}_k=0 \\ \mathbf{e}_{\mathbf{a}_i,k}=0}} = \left. \frac{\partial \mathbf{z}_{\mathbf{a}_i,k}^{\mathbf{n}}}{\partial \mathbf{x}_k} \right|_{\substack{\delta \mathbf{x}_k=0 \\ \mathbf{e}_{\mathbf{a}_i,k}=0}} \cdot \left. \frac{\partial \mathbf{x}_k}{\partial \delta \mathbf{x}_k} \right|_{\substack{\delta \mathbf{x}_k=0 \\ \mathbf{e}_{\mathbf{a}_i,k}=0}}$$

where the non-trivial block matrices are as follows:

$$\begin{aligned} \left. \frac{\partial \mathbf{z}_{\mathbf{a}_i,k}^{\mathbf{n}}}{\partial \mathbf{q}_k^{\mathbf{nb}}} \right|_{\substack{\delta \mathbf{x}_k=0 \\ \mathbf{e}_{\mathbf{a}_i,k}=0}} &= \frac{\partial (\mathbf{q}_k^{\mathbf{nb}} \odot R^{\mathbf{bs}} \mathbf{p}_{\mathbf{ba}_i}^{\mathbf{s}} \odot \mathbf{q}_k^{\mathbf{bn}})}{\partial \mathbf{q}_k^{\mathbf{nb}}} \\ \left. \frac{\partial \mathbf{q}_k^{\mathbf{nb}}}{\partial \boldsymbol{\eta}_k^{\mathbf{n}}} \right|_{\substack{\delta \mathbf{x}_k=0 \\ \mathbf{e}_{\mathbf{a}_i,k}=0}} &= \frac{1}{2} \begin{bmatrix} -q_x & -q_y & -q_z \\ q_w & q_z & -q_y \\ -q_z & q_w & q_x \\ q_y & -q_x & q_w \end{bmatrix} \end{aligned}$$

with $\tilde{\mathbf{q}}_k^{\mathbf{nb}} = [q_w, q_x, q_y, q_z]$. The expression for derivative of the rotation with respect to the quaternion can be found in [Sola, 2017, Sec. 4.3.2].

For the radar range rate measurement $\mathbf{z}_{\mathbf{r}_i,k}^{\mathbf{r}_i}$

$$\begin{aligned} \frac{\partial \mathbf{z}_{\mathbf{r}_i,k}^{\mathbf{r}_i}}{\partial \mathbf{v}_k^{\mathbf{n}}} &= \left[R^{\mathbf{r}_i \mathbf{s}} R^{\mathbf{sb}} \tilde{R}_k^{\mathbf{bn}} \right]_{[(0,1),(\cdot)]} \\ \frac{\partial \mathbf{z}_{\mathbf{r}_i,k}^{\mathbf{r}_i}}{\partial \mathbf{q}_k^{\mathbf{nb}}} &= \left[R^{\mathbf{r}_i \mathbf{s}} R^{\mathbf{sb}} \frac{\partial (\mathbf{q}_k^{\mathbf{bn}} \odot \tilde{\mathbf{v}}_k^{\mathbf{n}} \odot \mathbf{q}_k^{\mathbf{nb}})}{\partial \mathbf{q}_k^{\mathbf{nb}}} \right]_{[(0,1),(\cdot)]} \\ \frac{\partial \mathbf{z}_{\mathbf{r}_i,k}^{\mathbf{r}_i}}{\partial \mathbf{b}_{\omega,k}^{\mathbf{b}}} &= \left[-R^{\mathbf{r}_i \mathbf{s}} R^{\mathbf{sb}} [R^{\mathbf{bs}} \mathbf{p}_{\mathbf{br}_i}^{\mathbf{s}}]_{\times} \right]_{[(0,1),(\cdot)]} \end{aligned}$$

where $[0, 1][\cdot]$ denotes the first two rows of the matrix. The partial derivatives of $\mathbf{z}_{\text{nhc},k}^{\mathbf{v}}$ and $\mathbf{z}_{\text{zupt},k}^{\mathbf{v}}$ follow similarly.

B Nonlinear Error-State Rauch-Tung-Striebel Smoother

The conventional expression for the extended Rauch-Tung-Striebel (RTS) smoother is given as [Särkkä, 2013, Chap. 9]

$$\begin{aligned}\mathbf{x}_k^* &= \hat{\mathbf{x}}_k + C_k (\mathbf{x}_{k+1}^* - f_k(\hat{\mathbf{x}}_k)) \\ P_k^* &= P_k + C_k (P_{k+1}^* - F_k P_k F_k^\top - G_k Q_k G_k^\top) C_k^\top\end{aligned}$$

with

$$C_k = P_k F_k^\top (F_k P_k F_k^\top + G_k Q_k G_k^\top)^{-1}$$

where $*$ indicates the smoothed estimate and $\hat{\cdot}$ indicates the filtered estimate. This expression is derived by linearizing the dynamics at the *filtered* state estimate during the *backward smoothing* pass.

In contrast, this paper prefers to linearize the dynamics at the predicted smoothed estimate $\bar{\mathbf{x}}_k^*$ instead

$$\bar{\mathbf{x}}_k^* \triangleq f_k^{-1}(\mathbf{x}_{k+1}^*, \mathbf{u}_k, \mathbf{0})$$

This formulation results in a similar but slightly modified expression for the extended RTS smoother

$$\begin{aligned}\mathbf{x}_k^* &= \hat{\mathbf{x}}_k + C_k^* F_k^* (\bar{\mathbf{x}}_k^* - \hat{\mathbf{x}}_k) \\ P_k^* &= P_k + C_k^* (P_{k+1}^* - F_k^* P_k F_k^{*\top} - G_k^* Q_k G_k^{*\top}) C_k^{*\top}\end{aligned}$$

with

$$C_k^* = P_k F_k^{*\top} (F_k^* P_k F_k^{*\top} + G_k^* Q_k G_k^{*\top})^{-1}$$

where F_k^* and G_k^* denote linearized forward dynamics around $\bar{\mathbf{x}}_k^*$.

Acknowledgments

This work has been supported by Honda R&D Americas through The University of Texas Situation-Aware Vehicular Engineering Systems (SAVES) Center (<http://utsaves.org/>), an initiative of the UT Wireless Networking and Communications Group; by the National Science Foundation under Grant No. 1454474 (CAREER); by the Data-supported Transportation Operations and Planning Center (DSTOP), a Tier 1 USDOT University Transportation Center; and by Sandia National Laboratories.

References

- Almalioglu, Y., Turan, M., Lu, C. X., Trigoni, N., and Markham, A. (2019). Milli-RIO: Ego-motion estimation with low-cost millimetre-wave radar. *arXiv preprint arXiv:1909.05774*.
- Bar-Shalom, Y., Li, X. R., and Kirubarajan, T. (2001). *Estimation with Applications to Tracking and Navigation*. John Wiley and Sons, New York.
- Barnes, D. and Posner, I. (2020). Under the radar: Learning to predict robust keypoints for odometry estimation and metric localisation in radar. *arXiv preprint arXiv:2001.10789*.
- Barnes, D., Weston, R., and Posner, I. (2019). Masking by moving: Learning distraction-free radar odometry from pose information. *arXiv preprint arXiv:1909.03752*.
- Barra, J., Lesecq, S., Zarudniev, M., Debicki, O., Mareau, N., and Ouvry, L. (2019). Localization system in GPS-denied environments using radar and IMU measurements: Application to a smart white cane. In *2019 18th European Control Conference (ECC)*, pages 1201–1206. IEEE.
- Bijelic, M., Gruber, T., and Ritter, W. (2018). A benchmark for lidar sensors in fog: Is detection breaking down? In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 760–767. IEEE.
- Brooker, G., Hennessey, R., Lobsey, C., Bishop, M., and Widzyk-Capehart, E. (2007). Seeing through dust and water vapor: Millimeter wave radar sensors for mining applications. *Journal of Field Robotics*, 24(7):527–557.
- Callmer, J., Törnqvist, D., Gustafsson, F., Svensson, H., and Carlbom, P. (2011). Radar SLAM using visual features. *EURASIP Journal on Advances in Signal Processing*, 2011(1):71.
- Cen, S. H. and Newman, P. (2018). Precise ego-motion estimation with millimeter-wave radar under diverse and challenging conditions. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE.
- Cen, S. H. and Newman, P. (2019). Radar-only ego-motion estimation in difficult settings via graph matching. *arXiv preprint arXiv:1904.11476*.
- Chen, C. C. (1975). Attenuation of electromagnetic radiation by haze, fog, clouds, and rain. Technical report, RAND Corp, Santa Monica, CA.
- Chetverikov, D., Svirko, D., Stepanov, D., and Krsek, P. (2002). The trimmed iterative closest point algorithm. In *Object recognition supported by user interaction for service robots*, volume 3, pages 545–548. IEEE.

- Chiang, K.-W., Tsai, G.-J., Li, Y.-H., Li, Y., and El-Sheimy, N. (2020). Navigation engine design for automated driving using INS/GNSS/3D LiDAR-SLAM and integrity assessment. *Remote Sensing*, 12(10):1564.
- Choi, J., Va, V., Gonzalez-Prelcic, N., Daniels, R., Bhat, C. R., and Heath, R. W. (2016). Millimeter-wave vehicular communication to support massive automotive sensing. *IEEE Communications Magazine*, 54(12):160–167.
- Crassidis, J. L., Markley, F. L., and Cheng, Y. (2007). Survey of nonlinear attitude estimation methods. *Journal of guidance control and dynamics*, 30(1):12.
- Deutsch, H., Reuter, S., and Dietmayer, K. (2015). The labeled multi-Bernoulli SLAM filter. *IEEE Signal Processing Letters*, 22(10):1561–1565.
- Dissanayake, G., Sukkarieh, S., Nebot, E., and Durrant-Whyte, H. (2001). The aiding of a low-cost strapdown inertial measurement unit using vehicle model constraints for land vehicle applications. *IEEE transactions on robotics and automation*, 17(5):731–747.
- Erdinc, O., Willett, P., and Bar-Shalom, Y. (2009). The bin-occupancy filter and its connection to the PHD filters. *IEEE Transactions on Signal Processing*, 57(11):4232–4246.
- Fajardo, D., Au, T.-C., Waller, S., Stone, P., and Yang, D. (2011). Automated intersection control: Performance of future innovation versus current traffic signal control. *Transportation Research Record: Journal of the Transportation Research Board*, (2259):223–232.
- Fatemi, M., Granström, K., Svensson, L., Ruiz, F. J., and Hammarstrand, L. (2017). Poisson multi-bernoulli mapping using Gibbs sampling. *IEEE Transactions on Signal Processing*, 65(11):2814–2827.
- Foessel, A., Chheda, S., and Apostolopoulos, D. (1999). Short-range millimeter-wave radar perception in a polar environment.
- Forster, C., Lynen, S., Kneip, L., and Scaramuzza, D. (2013). Collaborative monocular slam with multiple micro aerial vehicles. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3962–3970.
- Gao, W. and Tedrake, R. (2019). FilterReg: Robust and efficient probabilistic point-set registration using gaussian filter and twist parameterization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11095–11104.
- Holder, M., Hellwig, S., and Winner, H. (2019). Real-time pose graph SLAM based on radar. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 1145–1151. IEEE.
- Hong, Z., Petillot, Y., and Wang, S. (2020). RadarSLAM: Radar based large-scale SLAM in all weathers. *arXiv preprint arXiv:2005.02198*.

- Humphreys, T. E. (2017). Interference. In *Springer Handbook of Global Navigation Satellite Systems*, pages 469–503. Springer International Publishing.
- Humphreys, T. E., Murrian, M. J., and Narula, L. (2020). Deep-urban unaided precise global navigation satellite system vehicle positioning. *IEEE Intelligent Transportation Systems Magazine*, 12(3):109–122.
- Jian, B. and Vemuri, B. C. (2010). Robust point set registration using Gaussian mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1633–1645.
- Jokela, M., Kutila, M., and Pyykönen, P. (2019). Testing and validation of automotive point-cloud sensors in adverse weather conditions. *Applied Sciences*, 9(11):2341.
- Kennedy, S., Hamilton, J., and Martell, H. (2006). Architecture and system performance of SPAN—NovAtel’s GPS/INS solution. In *Position, Location, And Navigation Symposium, 2006 IEEE/ION*, page 266. IEEE.
- Kobayashi, H. K. (1980). Atmospheric effects on millimeter radio waves. Technical report, Army Electronics Research and Development Command, White Sands Missile Range, NM.
- Kok, M., Hol, J. D., and Schön, T. B. (2017). Using inertial sensors for position and orientation estimation. *arXiv preprint arXiv:1704.06053*.
- Kramer, A., Stahoviak, C., Santamaria-Navarro, A., Agha-Mohammadi, A.-A., and Heckman, C. (2020). Radar-inertial ego-velocity estimation for visually degraded environments. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.
- Kutila, M., Pyykönen, P., Holzhüter, H., Colomb, M., and Duthon, P. (2018). Automotive LiDAR performance verification in fog and rain. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 1695–1701. IEEE.
- Kyle Stock (2018). Self-driving cars can handle neither rain nor sleet nor snow. <https://www.bloomberg.com/news/articles/2018-09-17/self-driving-cars-still-can-t-handle-bad-weather>.
- LaChapelle, D., Humphreys, T. E., Narula, L., Iannucci, P. A., and Moradi-Pari, E. (2020). Automotive collision risk estimation under cooperative sensing. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Barcelona, Spain.
- Li, R., Liu, J., Zhang, L., and Hang, Y. (2014). Lidar/mems imu integrated navigation (slam) method for a small uav in indoor environments. In *2014 DGON Inertial Sensors and Systems (ISS)*, pages 1–15. IEEE.
- Li, Y., Duthon, P., Colomb, M., and Ibanez-Guzman, J. (2020). What happens for a ToF LiDAR in fog? *IEEE Transactions on Intelligent Transportation Systems*.

- LORD Sensing MicroStrain. 3DM-GX5-25 Attitude and Heading Reference System. <https://bit.ly/32CKIaO>. Accessed 2020-08-31.
- Lundgren, M., Stenborg, E., Svensson, L., and Hammarstrand, L. (2014). Vehicle self-localization using off-the-shelf sensors and a detailed map. In *2014 IEEE Intelligent Vehicles Symposium Proceedings*, pages 522–528. IEEE.
- Mahler, R. P. (2003). Multitarget Bayes filtering via first-order multitarget moments. *IEEE Transactions on Aerospace and Electronic systems*, 39(4):1152–1178.
- Mohammed, A. S., Amamou, A., Ayevide, F. K., Kelouwani, S., Agbossou, K., and Zioui, N. (2020). The perception system of intelligent ground vehicles in all weather conditions: a systematic literature review. *Sensors*, 20(22):6532.
- Mullane, J., Vo, B.-N., Adams, M. D., and Vo, B.-T. (2011). A random-finite-set approach to Bayesian SLAM. *IEEE Transactions on Robotics*, 27(2):268–282.
- Mur-Artal, R. and Tardós, J. D. (2017). Visual-inertial monocular SLAM with map reuse. *IEEE Robotics and Automation Letters*, 2(2):796–803.
- Myronenko, A. and Song, X. (2010). Point set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2262–2275.
- Narula, L., Iannucci, P. A., and Humphreys, T. E. (2020a). Automotive-radar-based 50-cm urban positioning. In *Proceedings of the IEEE/ION PLANSx Meeting*.
- Narula, L., LaChapelle, D. M., Murrian, M. J., Wooten, J. M., Humphreys, T. E., Lacambre, J.-B., de Toldi, E., and Morvant, G. (2020b). TEX-CUP: The University of Texas Challenge for Urban Positioning. In *Proceedings of the IEEE/ION PLANSx Meeting*.
- Narula, L., Murrian, M. J., and Humphreys, T. E. (2018a). Accuracy limits for globally-referenced digital mapping using standard GNSS. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3075–3082. IEEE.
- Narula, L., Wooten, J. M., Murrian, M. J., LaChapelle, D. M., and Humphreys, T. E. (2018b). Accurate collaborative globally-referenced digital mapping with standard GNSS. *Sensors*, 18(8).
- Petovello, M., Cannon, M., and Lachapelle, G. (2004). Benefits of using a tactical-grade IMU for high-accuracy positioning. *Navigation, Journal of the Institute of Navigation*, 51(1):1–12.
- Qin, T., Li, P., and Shen, S. (2018). VINS-mono: A robust and versatile monocular visual-inertial state estimator. *IEEE Transactions on Robotics*, 34(4):1004–1020.

- Rasshofer, R. H. and Gresser, K. (2005). Automotive radar and lidar systems for next generation driver assistance functions. *Advances in Radio Science*, 3(B. 4):205–209.
- Reina, G., Underwood, J., Brooker, G., and Durrant-Whyte, H. (2011). Radar-based perception for autonomous outdoor vehicles. *Journal of Field Robotics*, 28(6):894–913.
- Ryde, J. and Hillier, N. (2009). Performance of laser and radar ranging devices in adverse environmental conditions. *Journal of Field Robotics*, 26(9):712–727.
- Särkkä, S. (2013). *Bayesian filtering and smoothing*, volume 3. Cambridge University Press.
- Scherzinger, B. M. (2006). Precise robust positioning with inertially aided RTK. *Navigation*, 53(2):73–83.
- Schoen, M., Horn, M., Hahn, M., and Dickmann, J. Real-time radar SLAM.
- Schuster, F., Keller, C. G., Rapp, M., Haueis, M., and Curio, C. (2016). Landmark based radar SLAM using graph optimization. In *Intelligent Transportation Systems (ITSC), 2016 IEEE 19th International Conference on*, pages 2559–2564. IEEE.
- Skog, I., Handel, P., Nilsson, J.-O., and Rantakokko, J. (2010). Zero-velocity detection—An algorithm evaluation. *IEEE transactions on biomedical engineering*, 57(11):2657–2666.
- Sola, J. (2017). Quaternion kinematics for the error-state Kalman filter. *arXiv preprint arXiv:1711.02508*.
- Steder, B., Grisetti, G., Stachniss, C., and Burgard, W. (2008). Visual SLAM for flying vehicles. *IEEE Transactions on Robotics*, 24(5):1088–1093.
- Stübler, M., Reuter, S., and Dietmayer, K. (2017). A continuously learning feature-based map using a Bernoulli filtering approach. In *2017 Sensor Data Fusion: Trends, Solutions, Applications (SDF)*, pages 1–6. IEEE.
- Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic robotics*. MIT press.
- Tsin, Y. and Kanade, T. (2004). A correlation-based approach to robust point set registration. In *European conference on computer vision*, pages 558–569. Springer.
- Wallace, H. B. (1988). Millimeter-wave propagation measurements at the Ballistic Research Laboratory. *IEEE transactions on geoscience and remote sensing*, 26:253–258.
- Ward, E. and Folkesson, J. (2016). Vehicle localization with low cost radar sensors. In *2016 IEEE Intelligent Vehicles Symposium (IV)*, pages 864–870. IEEE.

- Woodman, O. (2007). An introduction to inertial navigation. *University of Cambridge, Computer Laboratory, Tech. Rep. UCAMCL-TR-696*.
- Ye, H., Chen, Y., and Liu, M. (2019). Tightly coupled 3d lidar inertial odometry and mapping. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3144–3150. IEEE.
- Yen, K. S., Shankwitz, C., Newstrom, B., Lasky, T. A., and Ravani, B. (2015). Evaluation of the University of Minnesota GPS Snowplow Driver Assistance Program. Technical report, California Department of Transportation.
- Yoneda, K., Hashimoto, N., Yanase, R., Aldibaja, M., and Sukanuma, N. (2018). Vehicle localization using 76GHz omnidirectional millimeter-wave radar for winter automated driving. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 971–977. IEEE.
- Zang, S., Ding, M., Smith, D., Tyler, P., Rakotoarivelo, T., and Kaafar, M. A. (2019). The impact of adverse weather conditions on autonomous vehicles: how rain, snow, fog, and hail affect the performance of a self-driving car. *IEEE vehicular technology magazine*, 14(2):103–111.
- Zhang, H. T. (2006). Performance comparison on kinematic GPS integrated with different tactical-grade IMUs. Master’s thesis, The University of Calgary.