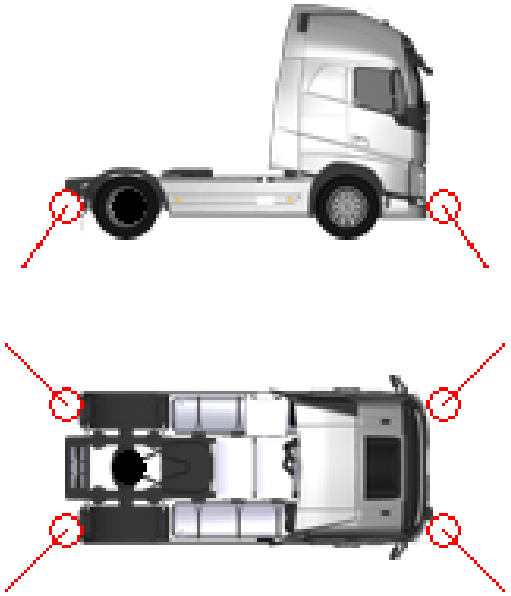

Ego-Motion and Road Attitude Estimation Using Ground Speed Radars and an IMU

Master's thesis in Systems, Control, and Mechatronics

SIMON BROMERT & ISAK ASFEDAI LARSSON

DEPARTMENT OF ELECTRICAL ENGINEERING

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2025
www.chalmers.se

Master's thesis 2025

Ego-Motion and Road Attitude Estimation Using Ground Speed Radars and an IMU

SIMON BROMERT
ISAK ASFEDAI LARSSON

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
Chalmers University of Technology
Gothenburg, Sweden 2025

Ego-Motion and Road Attitude Estimation Using Ground Speed Radars and an IMU
SIMON BROMERT
ISAK ASFEDAI LARSSON

Supervisor: Murat Kumru, Volvo Trucks
Examiner: Lars Hammarstrand, Department of Electrical Engineering

Master's Thesis 2025
Department of Electrical Engineering
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Screenshot from IPG TruckMaker. Used with permission from IPG Automotive.

Typeset in LaTeX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2025

Ego-Motion and Road Attitude Estimation Using Ground Speed Radars and an IMU

SIMON BROMERT, ISAK ASFEDAI LARSSON
Department of Electrical Engineering
Chalmers University of Technology

Abstract

Accurate ego-motion estimation is essential for autonomous and assisted driving in heavy-duty vehicles. This task becomes particularly challenging in scenarios where traditional sensors, such as wheel speed encoders, are unreliable. Additionally, a good understanding of the road characteristics, such as slope, is critical to ensure safe and efficient vehicle operation. This thesis addresses the mentioned challenges through a method that fuses measurements from ground-speed radars and an Inertial Measurement Unit (IMU). The proposed method jointly estimates vehicle motion and road attitude, while accounting for issues such as IMU bias and multirate sensor data.

The main contributions of this thesis are the estimation of road slope angle and a method for estimating IMU accelerometer biases. The road slope is estimated through pseudo-measurements derived from the vehicle's translational velocity. Meanwhile, the IMU accelerometer biases are estimated by employing a tilted IMU configuration. This setup enhances the observability of the biases, leading to improved filter convergence and estimation accuracy during typical driving maneuvers.

The proposed method is evaluated using simulated data generated with IPG TruckMaker (TM), a high-fidelity simulation environment that enables controlled, repeatable experiments across diverse driving scenarios. A Monte Carlo (MC) analysis is performed using multiple sensor configurations, allowing statistical assessment of system consistency and robustness. The results demonstrate that the filter reliably estimates vehicle velocity, attitude, and the road slope angle despite the presence of sensor noise and bias. These findings support the effectiveness of using a minimal radar setup in combination with a tilted IMU to enable robust and accurate ego-motion estimation in heavy-duty vehicle applications.

Keywords: Ego-Motion Estimation, Ego-Pose Estimation, Sensor Fusion, Ground-Speed Radar, Inertial Navigation, Extended Kalman Filter, Inertial Measurement Unit, Bias Estimation, Multirate Updates, Road Attitude Estimation.

Acknowledgements

We would like to begin by expressing our gratitude to Volvo GTT for providing us the opportunity to conduct our master's thesis within their organization. This experience has been invaluable and has contributed significantly to both our academic and personal development.

We are especially thankful to our industrial supervisor, Murat Kumru, for his ongoing support, guidance, and engagement throughout the project. We would also like to thank Axel Ceder for taking the time to discuss the limitations and potential improvements related to IMU usage. Furthermore, we are grateful to Aron Enliden and Anton Carlsson for their constructive feedback and peer reviews, which helped strengthen the quality of our work.

We would also like to thank the team at IPG Automotive, in particular Alexander Hägglund and Shreyas Kogalur, for their introduction to the simulation tools and for their prompt, engaged support whenever questions or issues with TruckMaker arose.

Finally, we extend our sincere thanks to our academic examiner at Chalmers University of Technology, Lars Hammarstrand, for his critical insights and thoughtful feedback on our design decisions. His input challenged us to improve our methodology and contributed to a more rigorous and well-rounded thesis.

Simon Bromert, Isak Asfedai Larsson, Gothenburg, June 2025

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

DH	Denavit–Hartenberg
DoF	Degrees of Freedom
ECEF	Earth-Centered, Earth-Fixed
ECI	Earth-Centered Inertial
EKF	Extended Kalman Filter
ENU	East–North–Up
FFT	Fast Fourier Transform
FMCW	Frequency-Modulated Continuous Wave
FoM	Figures of Merit
IMU	Inertial Measurement Unit
INS	Inertial Navigation System
KF	Kalman Filter
MC	Monte Carlo
NED	North-East-Down
RMSE	Root Mean Square Error
RSI	Raw Signal Interfaces
SoG	Speed Over Ground
TM	TruckMaker

Nomenclature

Below is the nomenclature of indices, sets, parameters, and variables that have been used throughout this thesis.

Notable Operations

$\mathbf{A}$	Vector notation (bold indicates a vector)
$\hat{\mathbf{A}}$	Unit vector in the direction of $\mathbf{A}$
$Pr_{\hat{\mathbf{A}}}(\mathbf{B})$	Projection of vector $\mathbf{B}$ onto unit vector $\hat{\mathbf{A}}$
$[\cdot]_B^A$	Expression of a quantity in frame A, with regards to frame B
$[\cdot]^{(A)}$	Expression of a quantity in frame A
$[\cdot]_{k-1\|k-1}$	Expression of a quantity at time $k-1$ given information for time $k-1$.
$[\cdot]_{k\|k-1}$	Expression of a quantity at time k given information for time $k-1$.
$[\cdot]_{k\|k}$	Expression of a quantity at time k given information for time k.
$[\cdot]_k$	Expression of a quantity at time k.

Constants

n_{radar}	Number of radar sensors
g	Gravitational acceleration constant ($9.81\,[\mathrm{m/s^2}]$)
Δt_{IMU}	IMU update period [s]
Δt_{Radar}	Radar update period [s]

Coordinate Frame Indices

i	Inertial Frame
n	Navigation Frame

r	Road Frame
b	Body Frame

Variables

t	Time
ϕ	Roll angle
θ	Pitch angle
ψ	Yaw angle
$\mathbf{R}$	Rotation matrix
$\mathbf{R}_{axis,angle}$	Rotation matrix about a particular axis
$\mathbf{\Phi}$	Euler angle vector (ϕ, θ, ψ)
$\boldsymbol{\omega}$	Angular velocity vector
$\mathbf{v}$	Translational velocity vector
$b_{sensor_{axis}}$	Bias of a given sensor (e.g., b_{a_x} for accelerometer x-axis)
η	Noise parameter

Kalman Filter Parameters

$\mathbf{x}$	State vector
$\mathbf{u}$	Control input vector
$\mathbf{P}$	State covariance matrix
$\mathbf{\Sigma}_x$	Process noise covariance matrix
$\mathbf{\Sigma}_y$	Measurement noise covariance matrix
$f(\mathbf{x}, \mathbf{u})$	Nonlinear process model
$F(\mathbf{x}, \mathbf{u})$	Jacobian of the process model
$h(\mathbf{x})$	Measurement model
$H(\mathbf{x})$	Jacobian of the measurement model
$\mathbf{y}_a$	Measured acceleration vector from IMU (accelerometer)
$\mathbf{y}_\omega$	Measured angular velocity vector from IMU (gyroscope)
$\mathbf{y}_r$	Measured position or velocity from Radar

Contents

List of Figures

List of Tables

1
Introduction

Reliable estimation of a vehicle's ego-motion and ego-pose is essential for safety-critical systems in modern heavy-duty vehicles. Traditional approaches, relying on wheel encoders, are unreliable in scenarios involving wheel slip or surface irregularities. To overcome these limitations, this thesis proposes a sensor fusion framework that combines measurements from ground-speed radars and an Inertial Measurement Unit (IMU).

To address these challenges, this work investigates a multi-sensor fusion approach grounded in well-established estimation theory. By leveraging the complementary characteristics of radars and an IMU sensor, the proposed method aims to deliver a robust and accurate estimation of vehicle motion and attitude. The system is evaluated with a range of dynamic driving scenarios, including slope transitions, turning maneuvers, and acceleration events.

This chapter outlines the motivation and background of the thesis, states the research objectives and questions, and defines the scope and limitations of the study.

1.1 Background

Modern vehicle control systems depend heavily on accurate real-time knowledge of ego-motion relative to the environment. This information is crucial for enabling safety systems such as anti-lock braking systems, adaptive cruise control, and advanced driver-assistance systems.

Traditional ego-motion estimation techniques use wheel encoders, which estimate ego-motion based on mechanical rotation. While effective in controlled conditions, these systems degrade under dynamic scenarios where issues such as wheel slip or lateral displacement are present. Because these methods measure relative wheel rotation rather than actual ground-referenced ego-motion, their reliability diminishes in low-friction or irregular situations [1].

Ground-speed radars address these limitations by directly measuring the vehicle's motion relative to the ground using the Doppler shift phenomena. These sensors operate independently of wheel dynamics, offering improved robustness under extreme maneuvers and challenging surface conditions.

A sensor that can provide complementary measurements is the IMU which provides high-frequency measurements of the translational acceleration and angular velocity. Although IMUs offer essential dynamic information, over time they suffer from drift and cumulative error in the absence of external references.

Together, radar and IMU sensors offer complementary strengths as the radar provides drift-free velocity data, while the IMU enables fast, orientation-aware measurements. When fused, these data sources support accurate ego-motion and ego-attitude estimation.

This thesis develops a sensor fusion framework that combines ground-speed radar and IMU measurements to estimate vehicle motion and orientation in dynamic driving scenarios, including turning, acceleration, braking, and slope transitions.

1.2 Objective and Purpose

The main objective of this thesis is to design and evaluate a sensor fusion framework for estimating the ego-motion and ego-attitude of a heavy-duty vehicle using multiple ground-speed radars and an IMU with six Degrees of Freedom (DoF). In addition to vehicle kinematics, the framework is designed to estimate road-related features, such as the road slope angle.

The proposed solution must be resilient to sensor imperfections, including measurement noise, static bias, and fixed internal misalignment. Moreover, it aims to provide probabilistic estimations that include confidence intervals, to allow for performance evaluation.

1.3 Research Questions

This thesis addresses the following key research questions:

1. What are the strengths and weaknesses of combining radar Doppler velocity and IMU measurements for ego-motion and ego-attitude estimation?

2. How can fused estimates of ego-motion and ego-attitude be used to derive road geometry characteristics such as slope and bank angles?

3. In what way do sensor imperfections, such as bias and noise affect the performance, accuracy, and robustness of the proposed estimation framework?

1.4 Scope and Limitations

This study is focused on the development and simulation-based evaluation of a sensor fusion framework that combines radar and IMU measurements for vehicle state

estimation. The system is intended for use in heavy-duty vehicles.

- **Data Generation:** All evaluations are based on synthetic data generated using a high-fidelity simulation environment. Doppler velocity inputs are derived from known ground-truth translational and rotational motion.

- **Sensor Configuration:** Sensor placement and orientation are assumed to be fixed and known. No physical redesign or optimization of hardware is performed.

- **Environmental Conditions:** The scope includes common road geometries such as slopes and banked curves. External environmental factors like weather or road surface variation are not considered.

- **Sensor Models:** IMU and radar measurements include additive white noise and static biases. Calibration errors, misalignment, and scale factor deviations are excluded. The radar measurements include Gaussian noise but not systematic effects such as Doppler bias or beam misalignment.

- **Radar Measurement Scope:** The Doppler velocity is the only information used from the radar sensors. Thus are Doppler range measurements disregarded.

1.5 Related Work

A wide range of methods have been proposed for ego-motion estimation using radar sensors. In [2], the authors extend the single-radar approach of [3] by employing multiple object-based radars to extract Doppler velocities from various spatial points on the vehicle. Stationary targets are filtered using RANSAC, and motion is estimated using Levenberg–Marquardt optimization. While effective, these methods are susceptible to velocity-induced degradation and environmental dependencies. Additionally, they rely on object detection, which differs from the ground-speed radar approach employed in this thesis.

Radar-IMU fusion has been explored in small robotic systems using discrete-time filtering methods [4]. These approaches typically focus on three DoF velocity estimation in visually degraded environments and are not designed for the full pose estimation or complex dynamics of heavy-duty vehicles.

A continuous-time radar-IMU fusion approach using B-spline interpolation is proposed in [5], which relies on radar detections of static landmarks for state updates. While effective in improving temporal resolution, it introduces sensitivity to azimuth errors and requires complex modeling of external reference features. These challenges are avoided in this thesis by using ground-speed radars that only detect the Doppler velocity of the road relative to the vehicle and disregards external land-

marks.

Recent work [6] presents radar-only ego-motion estimation using Doppler and range data. Although effective in certain structured environments, the method lacks IMU integration, making it more vulnerable to angular inaccuracies and environmental variability. Furthermore, the lack of an IMU makes attitude estimation with regard to the inertial frame impossible. This thesis extends such work by introducing a radar-IMU fusion method capable of estimating vehicle motion, orientation, and road slope with improved robustness and accuracy.

1.6 Thesis Contributions

This thesis advances radar–IMU-based ego-motion estimation by introducing a fusion framework tailored to the dynamics and operational requirements of heavy-duty vehicles. The proposed estimation architecture is modular and scalable.

It combines multiple ground-speed radars and an IMU to estimate vehicle motion and attitude. Unlike object-based or simplified velocity approaches, this technique relies on direct Doppler velocity and inertial measurements to produce drift-resilient state estimates.

The method also estimates road geometry features such as road slope, extending its functionality beyond motion estimation.

2
Theory

This chapter presents the theoretical foundations used in the thesis, including the modeling, sensor data processing, and estimation techniques used throughout the thesis. It also introduces mathematical tools for sensor fusion, reference frame transformations, inertial navigation dynamics, and radar measurement modeling.

The main concepts used throughout the thesis are Bayesian inference for state estimation, the handling of reference frames and kinematics for transforming sensor data, and the treatment of sensor errors such as bias and noise. The chapter also addresses system observability analysis and simplifications made to make the problem feasible within the scope of this thesis.

2.1 Foundations of Sensor Fusion

Sensor fusion refers to the process of combining measurements from multiple sensors to estimate a system's state. In this thesis, it is central to combining data from ground-speed radars and an IMU to estimate the ego-motion and ego-attitude of heavy-duty vehicles.

This section introduces the theoretical basis for sensor fusion, beginning with Bayesian inference and proceeding to the Kalman Filter (KF) and its nonlinear extension, the Extended Kalman Filter (EKF).

2.1.1 Bayesian Statistics in a Sensor Fusion Context

Bayesian statistics takes, previously known information into account when producing the posterior probability of an outcome.

Bayesian inference updates a prior belief about the system state based on new measurements using Bayes' theorem:

$$p(\mathbf{x}|\mathbf{z}) = \frac{p(\mathbf{z}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{z})}, \tag{2.1}$$

where $p(\mathbf{x}|\mathbf{z})$ is the posterior probability of the state $\mathbf{x}$ given the measurement $\mathbf{z}$, $p(\mathbf{z}|\mathbf{x})$ is the likelihood, $p(\mathbf{x})$ is the prior, and $p(\mathbf{z})$ is the marginal likelihood [7].

2.1.2 Kalman Filter

The KF is a recursive state estimation algorithm for linear, discrete-time systems subject to Gaussian noise. It provides an estimate of the system states by sequentially predicting and correcting the mean and covariance of the posterior distribution. These two steps are taken using mathematical model and sensor measurements [7].

A linear state-space model is expressed as

$$\mathbf{x}_k = \mathbf{F}\mathbf{x}_{k-1} + \mathbf{B}\mathbf{u}_{k-1} + \mathbf{w}_{k-1}, \tag{2.2}$$
$$\mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{v}_k, \tag{2.3}$$

where:
- $\mathbf{x}$: state vector
- $\mathbf{F}$: state transition matrix
- $\mathbf{B}$: control input matrix
- $\mathbf{u}$: control input
- $\mathbf{w} \sim \mathcal{N}(0, \mathbf{\Sigma}_x)$: process noise
- $\mathbf{z}$: measurement vector
- $\mathbf{H}$: observation matrix
- $\mathbf{v} \sim \mathcal{N}(0, \mathbf{\Sigma}_y)$: measurement noise

Based on these models, the KF operates in two main phases, prediction and update:

Prediction step

$$\hat{\mathbf{x}}_{k|k-1} = \mathbf{F}\hat{\mathbf{x}}_{k-1|k-1} + \mathbf{B}\mathbf{u}_{k-1}, \tag{2.4}$$
$$\mathbf{P}_{k|k-1} = \mathbf{F}\mathbf{P}_{k-1|k-1}\mathbf{F}^\top + \mathbf{\Sigma}_x, \tag{2.5}$$

Update step

$$\mathbf{y}_k = \mathbf{z}_k - \mathbf{H}\hat{\mathbf{x}}_{k|k-1}, \tag{2.6}$$
$$\mathbf{S}_k = \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^\top + \mathbf{\Sigma}_y, \tag{2.7}$$
$$\mathbf{K}_k = \mathbf{P}_{k|k-1}\mathbf{H}^\top\mathbf{S}_k^{-1}, \tag{2.8}$$
$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k\mathbf{y}_k, \tag{2.9}$$
$$\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_k\mathbf{H})\mathbf{P}_{k|k-1}, \tag{2.10}$$

where $\mathbf{K}$ is the Kalman gain and $\mathbf{y}$ is the innovation, $\mathbf{S}$ is the innovation covariance, $\hat{\mathbf{x}}$ is the state estimate, and $\mathbf{P}$ is the error covariance matrix.

2.1.3 Extended Kalman Filter

The EKF generalizes the KF to non-linear systems by linearizing the process and measurement models using first-order Taylor expansion. This allows the KF framework to be applied in settings where linear models are insufficient. The EKF performance depends, however, on the quality of the linearization and can degrade if the system exhibits strong non-linearities [7].

The system is modeled as

$$\mathbf{x}_k = f(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}) + \mathbf{w}_{k-1}, \tag{2.11}$$
$$\mathbf{z}_k = h(\mathbf{x}_k) + \mathbf{v}_k, \tag{2.12}$$

where $f(\cdot)$ and $h(\cdot)$ are nonlinear functions for the system process and measurement models, respectively. As in the linear case, the noise is Gaussian:

$$\mathbf{w}_{k-1} \sim \mathcal{N}(0, \mathbf{\Sigma}_x), \quad \mathbf{v}_k \sim \mathcal{N}(0, \mathbf{\Sigma}_y).$$

The EKF requires linearization using Jacobians:

- $\mathbf{F}_k = \left.\frac{\partial f}{\partial \mathbf{x}}\right|_{\hat{\mathbf{x}}_{k-1|k-1}}$: Jacobian of the process model
- $\mathbf{H}_k = \left.\frac{\partial h}{\partial \mathbf{x}}\right|_{\hat{\mathbf{x}}_{k|k-1}}$: Jacobian of the measurement model

The two main phases the EKF operates in are the same as for the KF, prediction and update:

Prediction step

$$\hat{\mathbf{x}}_{k|k-1} = f(\hat{\mathbf{x}}_{k-1|k-1}, \mathbf{u}_{k-1}), \tag{2.13}$$
$$\mathbf{P}_{k|k-1} = \mathbf{F}_k \mathbf{P}_{k-1|k-1} \mathbf{F}_k^\top + \mathbf{\Sigma}_x. \tag{2.14}$$

Update step

$$\mathbf{y}_k = \mathbf{z}_k - h(\hat{\mathbf{x}}_{k|k-1}), \tag{2.15}$$
$$\mathbf{S}_k = \mathbf{H}_k \mathbf{P}_{k|k-1} \mathbf{H}_k^\top + \mathbf{\Sigma}_y, \tag{2.16}$$
$$\mathbf{K}_k = \mathbf{P}_{k|k-1} \mathbf{H}_k^\top \mathbf{S}_k^{-1}, \tag{2.17}$$
$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k \mathbf{y}_k, \tag{2.18}$$
$$\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_{k|k-1}. \tag{2.19}$$

2.2 Coordinate Systems and Kinematics

Precise definitions and consistent handling of coordinate systems are fundamental to interpreting sensor measurements and modeling vehicle motion. Since different sensors report quantities relative to distinct reference frames, this section introduces the coordinate conventions used throughout the thesis. It describes the orientations and relationships of these frames and explains how to transform quantities between them.

2.2.1 Coordinate Frame Definitions

A reference frame defines the spatial context of a certain component of the system, in which quantities such as velocity and orientation are expressed. In three

dimensions, this is established by three mutually orthogonal unit vectors, typically denoted as the x, y, and z axes. These vectors are used to construct a rotation matrix that defines the frame's orientation with respect to an inertial or local frame [8].

Common reference frames in dynamic systems include inertial, body-fixed, and instrument frames. These may be static or dynamic, depending on whether their position and orientation change over time. The following provides a detailed description of each frame type.

Earth-based inertial frames include the Earth-Centered Inertial (ECI) frame, fixed at Earth's center with the z-axis aligned to the North Pole, and the Earth-Centered, Earth-Fixed (ECEF) frame, which shares the same origin but rotates with the Earth. For local navigation, a vehicle-fixed inertial frame can be defined in ECEF coordinates using either North-East-Down (NED) or East–North–Up (ENU) conventions. This frame's orientation remains fixed while its position follows the vehicle's movement [9]. These frames are illustrated in Figure 2.1.

The body-fixed frame is rigidly attached to the vehicle and defined relative to a point on the chassis. It moves with the vehicle, meaning its orientation and position vary over time according to vehicle motion [8].

The instrument frame represents the sensor's position and orientation relative to the body frame. Its alignment may be static or dynamic depending on the mounting configuration. Because it is physically mounted on the vehicle, its motion is governed by both its local mounting and the motion of the vehicle as a whole [8].

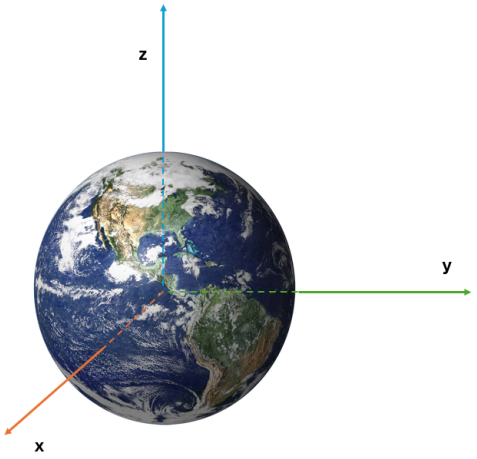

(a) ECI Frame

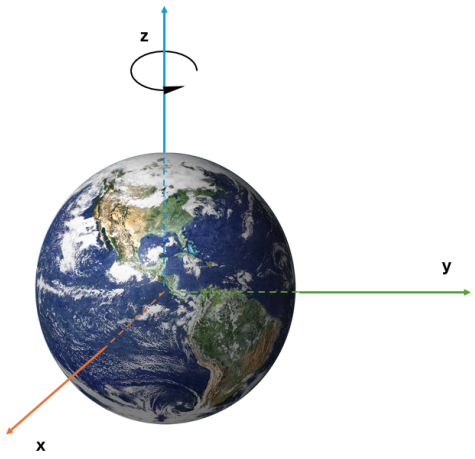

(b) ECEF Frame

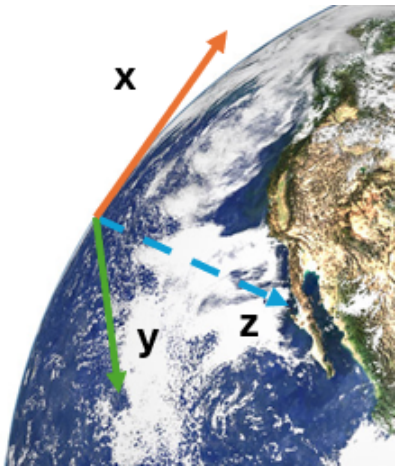

(c) NED Frame

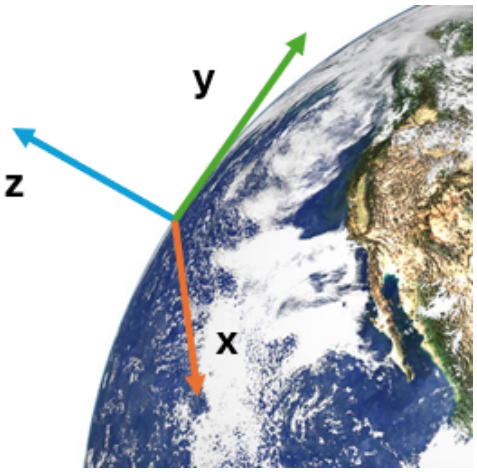

(d) ENU Frame

Figure 2.1: Visualization of the Earth-based reference frames. Top: global inertial and Earth-fixed frames (ECI and ECEF). Bottom: local-level navigation frames (NED and ENU). The image of Earth by DonkeyHotey is licensed under CC BY 3.0.

The relationships between these reference frames are described using rotation matrices and translation vectors, which may be constant or time-varying depending on the nature of the frames involved. For static frames, such as fixed sensor mounts, these transformations are defined by fixed kinematic parameters like mounting position and orientation. In contrast, for dynamic frames the transformations evolve over time based on the motion of the vehicle [8].

2.2.2 Vehicular Frame Standards

In vehicular navigation, several coordinate frame conventions are used to represent a vehicle and its motion relative to the surrounding environment. This thesis follows a widely adopted standard, ISO8855:2011, with an additional extension introduced to support road estimation.

Standard ISO 8855:2011 Convention

The ISO8855:2011 standard, introduced in [10], defines three hierarchical coordinate frames:

- **Inertial**: A fixed global reference frame aligned with the ENU convention. It serves as the inertial frame for absolute positioning and orientation.
- **Intermediate**: A local frame aligned with the road heading but excluding road slope and bank. Its yaw may differ from the Earth frame depending on the road heading.
- **Vehicle**: A body-fixed frame attached to the middle of the vehicle's rear axle. It shares the yaw of the intermediate frame but may differ in pitch and roll.

Extended Frame Structure

To support estimation of road characteristics, an extended frame structure was used. This results in four reference frames:

- **Inertial (i)**: A fixed global reference frame aligned with the ENU convention. It is equivalent to the Inertial frame in ISO8855:2011.
- **Navigation (n)**: A local frame aligned with the road heading but excluding road slope and bank. It corresponds to the ISO8855:2011 intermediate frame.
- **Road (r)**: A body-fixed frame aligned with the pitch and roll of the road surface. It has the same yaw as the navigation frame.
- **Body (b)**: A body-fixed frame aligned with the pitch and roll of the vehicle. It has the same yaw as the navigation frame.

2.2.3 Rotation Matrices

Rotation matrices provide a framework for representing both rotational and translational relationships between frames. These transformations are with regard to the frame from which they originate. When using Euler angles, the rotation matrix is typically constructed from a sequence of basic rotations about each axis. These are defined as:

$$\mathbf{R}_{x,\phi} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & c_\phi & -s_\phi \\ 0 & s_\phi & c_\phi \end{bmatrix}, \quad \mathbf{R}_{y,\theta} = \begin{bmatrix} c_\theta & 0 & s_\theta \\ 0 & 1 & 0 \\ -s_\theta & 0 & c_\theta \end{bmatrix}, \quad \mathbf{R}_{z,\psi} = \begin{bmatrix} c_\psi & -s_\psi & 0 \\ s_\psi & c_\psi & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{2.20}$$

where ϕ, θ, ψ are the roll, pitch, and yaw angles, respectively, and $s_x = \sin(x)$ and $c_x = \cos(x)$ are shorthand notations for sine, cosine, and tangent [11].

Denavit–Hartenberg Convention

The Denavit–Hartenberg (DH) convention provides a way to describe the relationship between adjacent coordinate frames, primarily used in robotics. It also applies well to vehicle systems, particularly when modeling sensors rigidly mounted to a body. The DH method uses four parameters to describe each transformation:

Table 2.1: DH parameters.

Parameter	Description
θ_i	Joint angle (rotation around z_{i-1})
d_i	Link offset (translation along z_{i-1})
a_i	Link length (translation along x_i)
α_i	Link twist (rotation around x_i)

When the relationship between two joints are dynamic, θ_i and d_i vary. θ_i varies for rotational joints and d_i varies for translational joints [8]. A contextually relevant visualization of how the DH parameters is shown in Figure 2.2.

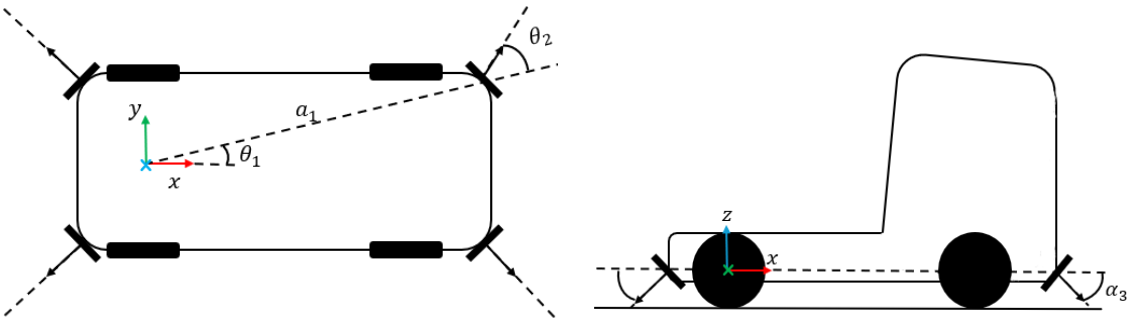

(a) Top view: DH parameters **(b)** Side view: DH parameters

Figure 2.2: Visualization of the DH parameters from two perspectives.

2.2.4 Angular Rate Transformations

The gyroscope in the IMU measures angular velocity in the body frame, denoted as $\boldsymbol{\omega}^{(b)}$. However, these body-frame angular rates cannot be directly interpreted as changes in Euler angles since the angular rates do not correspond to the same axes as the Euler rates. Accurate orientation estimation therefore requires converting

body-frame angular rates into Euler rates using a transformation that depends on the current orientation.

To perform this conversion, the relationship between body angular velocity and the Euler rates must be established. Assuming a yaw–pitch–roll rotation order, this transformation involves sequentially rotating the angular velocity vectors into the appropriate frames. The roll rate $\dot{\phi}$ is defined directly in the body frame. The pitch rate $\dot{\theta}$ is determined by rotating the corresponding angular velocity component about the body x-axis. Finally, the yaw rate $\dot{\psi}$ is obtained after rotating the angular velocity about both the x and y axes.

Following the formulation in [11], the body-frame angular velocity is expressed as:

$$\boldsymbol{\omega}^{(b)} = \begin{bmatrix} \dot{\phi} \\ 0 \\ 0 \end{bmatrix} + \mathbf{R}_{x,\phi}^{\top} \begin{bmatrix} 0 \\ \dot{\theta} \\ 0 \end{bmatrix} + \mathbf{R}_{x,\phi}^{\top} \mathbf{R}_{y,\theta}^{\top} \begin{bmatrix} 0 \\ 0 \\ \dot{\psi} \end{bmatrix}. \tag{2.21}$$

This relation can be compactly expressed using a inverse transformation matrix $\mathbf{T}_{\Phi}^{-1}$:

$$\boldsymbol{\omega}^{(b)} = \mathbf{T}_{\Phi}^{-1}(\boldsymbol{\Phi}) \begin{bmatrix} \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \end{bmatrix}, \tag{2.22}$$

where $\boldsymbol{\Phi} = \begin{bmatrix} \phi & \theta & \psi \end{bmatrix}^{\top}$ and the matrix $\mathbf{T}_{\Phi}^{-1}$ is defined as:

$$\mathbf{T}_{\Phi}^{-1}(\boldsymbol{\Phi}) = \begin{bmatrix} 1 & 0 & -s_{\theta} \\ 0 & c_{\phi} & c_{\theta} s_{\phi} \\ 0 & -s_{\phi} & c_{\theta} c_{\phi} \end{bmatrix}. \tag{2.23}$$

To extract Euler rates from body-frame angular velocity, the regular transformation is used:

$$\begin{bmatrix} \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \end{bmatrix} = \mathbf{T}_{\Phi}(\boldsymbol{\Phi}) \boldsymbol{\omega}^{(b)}, \tag{2.24}$$

where

$$\mathbf{T}_{\Phi}(\boldsymbol{\Phi}) = \begin{bmatrix} 1 & s_{\phi} t_{\theta} & c_{\phi} t_{\theta} \\ 0 & c_{\phi} & -s_{\phi} \\ 0 & \frac{s_{\phi}}{c_{\theta}} & \frac{c_{\phi}}{c_{\theta}} \end{bmatrix}. \tag{2.25}$$

2.3 Sensor Errors and Limitations

Sensor imperfections are a significant source of uncertainty in measurement systems, often limiting the overall accuracy and robustness of estimation and control. These errors originate from both physical and electronic sources, and their impact varies depending on the sensor type and operating conditions.

Table 2.2 shows three common imperfections found in inertial sensors and Frequency-Modulated Continuous Wave (FMCW). These imperfections will be addressed in this thesis, and therefore the simulation method for each is also displayed.

Table 2.2: Common error sources in inertial sensors and FMCW radar systems.

Error Type	Radar (FMCW)	Accelerometer	Gyroscope	Simulation Method
Noise	Thermal noise, phase noise	White noise	White noise	Gaussian Noise
Jitter	Sampling timing variation	Minor timing uncertainty	Minor timing uncertainty	Gaussian Noise
Bias	DC offset, calibration drift	Static thermal bias	Long-term drift	Added DC-offset

2.3.1 Noise

Noise may have many sources in electrical systems. These range from quantitation noise in signal converters, which occurs when sampling data from sensors, to white noise stemming from the cables, sensor components imperfections, or electromagnetic interference [12].

2.3.2 Jitter

Jitter, or clock jitter, is a type of error that occurs due to irregular sampling. This irregular sampling is usually the result of an internal system clock that causes sampling times to be non-uniform. This has a large impact on systems that have signals with high amplitudes and frequencies since clock jitter may cause the system to detect a very different value than the ground truth. However, over the Nyquist region of the signal's spectrum, the jitter is regarded to be white. It is unavoidable within the Nyquist region [12].

2.3.3 Bias

Sensor bias is an offset of the output of a sensor. This affects all sensors and can be detrimental to the performance of systems reliant on this bias. It is therefore important to accurately estimate the bias to be able to compensate for it. There are several methods of estimating bias, amongst them the least square method and the Kalman Filter [12].

2.4 Sensor Physics and Models

Accurate modeling of sensor behavior requires an understanding in their function. Inertial sensors, such as accelerometers and gyroscopes, rely on mechanical dynamics to detect motion. Radar systems use electromagnetic wave propagation and reflection, in combination with advanced signal processing techniques, to estimate distance and velocity.

This section outlines the function of the sensors as well as the mathematical models

used for each sensor type in this thesis. In the accelerometer and gyroscope subsections, focus is placed on the distinction between the true and inertial parts of the measurements. The radar subsection presents the function of FMCW radars, along with the role of signal processing in extracting target range and velocity.

2.4.1 Accelerometer

An accelerometer measures the specific force acting on a body. This includes both the true inertial acceleration and the apparent force resulting from gravity, as experienced in the sensor frame. However, the accelerometer cannot distinguish between these contributions, causing the measurement to be a combination of these and various sensor disturbances.

In accordance with Newton's second law of motion, the specific force $\mathbf{f}^{(b)}$ in the body frame is given by

$$\mathbf{f}^{(b)} = \frac{1}{m}\sum \mathbf{F}_{\text{specific}} = \mathbf{a}^{(b)} + \mathbf{g}^{(b)}, \tag{2.26}$$

where $\mathbf{a}^{(b)}$ is the true inertial acceleration and $\mathbf{g}^{(b)}$ is the gravitational acceleration, both expressed in the body frame. However, since the measurement is not perfect, a realistic model of the measured acceleration $\mathbf{y}_a$ includes not only the specific force expressed in (2.26), but also sensor errors

$$\mathbf{y}_a = \mathbf{a}^{(b)} + \mathbf{g}^{(b)} + \mathbf{b}_a + \mathbf{n}_a, \tag{2.27}$$

where $\mathbf{b}_a$ represents sensor bias and $\mathbf{n}_a$ denotes additive noise, commonly modeled as zero-mean Gaussian white noise [13].

The gravitational acceleration $\mathbf{g}^{(b)}$ is typically transformed from the navigation frame using a rotation matrix

$$\mathbf{g}^{(b)} = \mathbf{R}_b^n \mathbf{g}^{(n)}, \tag{2.28}$$

where $\mathbf{g}^{(n)} = \begin{bmatrix} 0 & 0 & g \end{bmatrix}^\top$ with $g \approx 9.81\,\mathrm{m/s^2}$, and $\mathbf{R}_b^n$ is the rotation matrix from the navigation frame to the body frame.

2.4.2 Gyroscope

A gyroscope measures the angular velocity of a body, representing its rate of rotation about one or more axes in the sensor's frame. Gyroscopes operate by sensing the Coriolis force induced on a vibrating proof mass. When the sensor undergoes rotation, the Coriolis force causes a displacement of the mass that is orthogonal to both the vibration direction and the axis of rotation. This displacement is proportional to the angular velocity and can be detected electrically.

In practical applications, the gyroscopic measurement is affected by sensor errors. This is modeled by

$$\mathbf{y}_\omega = \boldsymbol{\omega}^{(b)} + \mathbf{b}_\omega + \mathbf{n}_\omega, \tag{2.29}$$

where $\mathbf{y}_\omega$ is the gyroscopic measurement, $\boldsymbol{\omega}^{(b)}$ is the angular velocity, $\mathbf{b}_\omega$ is the gyroscope bias and $\mathbf{n}_\omega$ is additive noise, generally assumed to be zero-mean Gaussian white noise [13].

Since angular velocity is integrated over time to determine orientation, even small errors in gyroscope measurements can lead to significant long-term drift. This makes the gyroscopic sensor errors detrimental to the performance of most systems.

2.4.3 Radar

Ground-speed radars are downward-facing sensors that estimate a platform's velocity relative to the underlying terrain. Unlike front-facing automotive radars, that detects objects in front of the vehicle, ground-speed radars rely on the sensors relative velocity to the ground.

Both radar configurations use FMCW radar technology, that transmits a frequency varied signal. The reflected signal is compared with the transmitted waveform to produce a beat frequency f_b, which encodes the range through the signal delay $\tau = \frac{2d}{c}$, where d is the distance and c is the propagation speed of the signal [14].

For downward-facing radars with fixed mounting angles, the measured beat frequency can be directly mapped to the platform's relative ground speed. However, performance may degrade in the presence of low-reflectivity terrain, steep tilt angles, or adverse environmental conditions.

Although FMCW radar signal processing typically involves operations such as 2D Fast Fourier Transforms (FFT), this thesis omits these details as it does not delve into processing the raw radar signals. The radar is modeled as a velocity sensor. An illustrative example of FMCW signal behavior is shown in Figure 2.3 where f_b is the beat frequency, f_d is the Doppler frequency, and T_m is the modulation period [15].

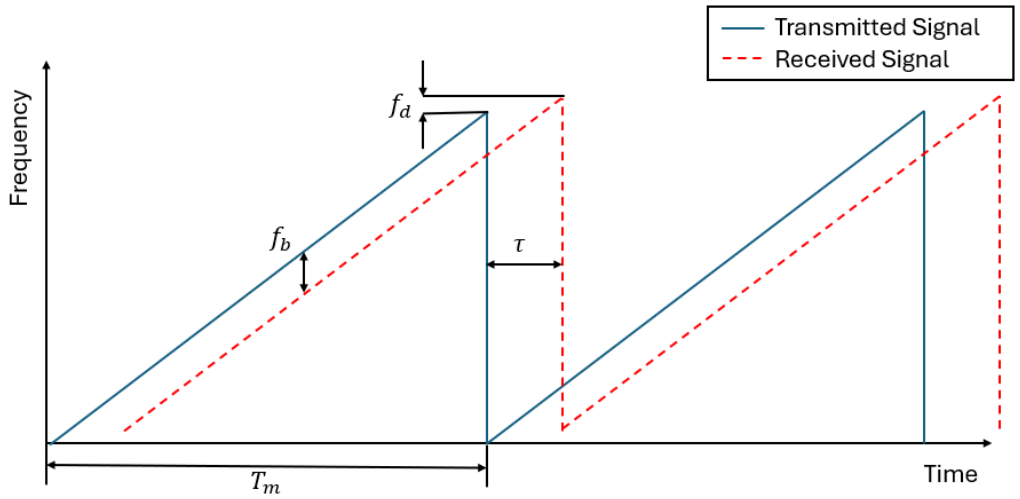

Figure 2.3: Example of the transmitted and received FMCW signals.

2.4.4 Doppler Effect and Doppler Shift

The Doppler effect describes the change in frequency of a wave, that has been transmitted and then received, due to relative motion between the transmitter and the reflective surface. In radar systems, this results in a frequency shift called the Doppler shift. This shift is directly related to the reflective surface's relative radial velocity with respect to the radar sensor.

The Doppler shift is, for the simple case where a single frequency is present, defined as the difference between the received and transmitted frequency

$$f_d = f_{\text{Rx}} - f_{\text{Tx}}, \tag{2.30}$$

where f_d is the Doppler shift, f_{Rx} is the frequency of the received signal, and f_{Tx} is the frequency of the transmitted signal.

In FMCW radar systems, the transmitted signal is a sequence of chirps, signals whose frequency increases linearly over time, like the example signal in Figure 2.3. The echoes of these chirps are sampled to form a two-dimensional time-domain matrix with two axes, the fast and short time axes. The fast time axis represents samples within a single chirp, while the slow time axis represents the sequence of chirps over time [15].

To retrieve both range and relative velocity information, a two-dimensional FFT is applied to the fast-slow time matrix. The FFT along the slow time axis extracts Doppler velocity components, and the FFT along the fast time axis resolves range. This results in a Range–Doppler map, which reveals both the distance and relative speed of the reflective surface.

2.5 Concepts for Dynamic Modeling

This section examines the principles behind the modeling of system's dynamics. It first examines fictitious forces and accelerations that arise in non-inertial frames, such as those affecting vehicle mounted sensors, with a focus on centripetal and centrifugal accelerations. It also explores the Moore-Penrose inverse which is used throughout the thesis.

2.5.1 Fictitious and Inertial Forces

Fictitious accelerations arise when motion is described relative to a non-inertial reference frame which is, one that is accelerating or rotating with respect to an inertial frame. These accelerations do not result from physical forces but are introduced to preserve Newton's second law when applied in the non-inertial frame.

The most commonly encountered fictitious accelerations and their associated forces include:

- **Centrifugal Acceleration**: An apparent outward acceleration experienced in a rotating frame with constant angular velocity. It balances the real centripetal acceleration observed in the inertial frame.
- **Euler Acceleration**: The apparent acceleration experienced in a rotating frame when the angular velocity is changing.
- **Rectilinear Force**: The apparent backward force felt in a linearly accelerating frame, such as when a vehicle suddenly speeds up.
- **Coriolis Force**: A fictitious force observed within a rotating frame that acts on moving objects, causing their paths to curve relative to the surface of the rotating system.

In this thesis, only the centrifugal and centripetal accelerations are considered, as they are the most relevant in the context of rigidly mounted vehicle sensors.

Centripetal acceleration is a real force observed in the inertial frame whenever an object follows a curved path. It acts perpendicular to the object's velocity vector and always points toward the center of the turn. If the object moves in a straight line, there is no change in direction, and the centripetal acceleration is zero [16]. From the perspective of a rotating frame, the effect of centripetal acceleration is counteracted by an outward fictitious acceleration known as centrifugal acceleration [17].

The centrifugal acceleration acting on a point at position $\mathbf{r}$ relative to the axis of rotation in a rotating frame with angular velocities $\boldsymbol{\omega}^{(b)}$ is given by:

$$\mathbf{a}^{(b)}_{\text{centrifugal}} = -\boldsymbol{\omega}^{(b)} \times (\boldsymbol{\omega}^{(b)} \times \mathbf{r}). \tag{2.31}$$

In this thesis, sensor positions are assumed fixed relative to the rotating vehicle frame, meaning $\mathbf{r}$ is constant. The corresponding velocity of the point in the inertial frame is:

$$\mathbf{v}^{(b)} = \boldsymbol{\omega}^{(b)} \times \mathbf{r}. \tag{2.32}$$

Substituting this into the centrifugal acceleration expression, the real centripetal acceleration experienced in the inertial frame can be expressed as:

$$\mathbf{a}^{(b)}_{\text{centripetal}} = \boldsymbol{\omega}^{(b)} \times \mathbf{v}^{(b)}, \tag{2.33}$$

since

$$\mathbf{a}^{(b)}_{\text{centripetal}} = -\mathbf{a}^{(b)}_{\text{centrifugal}}. \tag{2.34}$$

2.5.2 Moore–Penrose Inverse

The Moore–Penrose inverse, denoted $\mathbf{A}^\dagger$ and commonly referred to as the pseudo-inverse, extends the concept of matrix inversion to rectangular matrices.

Given a matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$, the pseudo-inverse $\mathbf{A}^\dagger \in \mathbb{R}^{n\times m}$ is the unique matrix that satisfies the four Penrose conditions [18]:

$$\mathbf{A}\mathbf{A}^\dagger\mathbf{A} = \mathbf{A}, \tag{2.35}$$
$$\mathbf{A}^\dagger\mathbf{A}\mathbf{A}^\dagger = \mathbf{A}^\dagger, \tag{2.36}$$
$$(\mathbf{A}\mathbf{A}^\dagger)^\top = \mathbf{A}\mathbf{A}^\dagger, \tag{2.37}$$
$$(\mathbf{A}^\dagger\mathbf{A})^\top = \mathbf{A}^\dagger\mathbf{A}. \tag{2.38}$$

These conditions are necessary for $\mathbf{A}^\dagger$ to behave like an ordinary inverse when $\mathbf{A}$ is square and full-rank, while remaining well-defined for all real or complex matrices.

The pseudo-inverse provides the least-squares solution to the system $\mathbf{A}\mathbf{x} = \mathbf{b}$, given by:

$$\mathbf{x} = \mathbf{A}^\dagger\mathbf{b}. \tag{2.39}$$

2.6 System Evaluation Theory

This section illustrates the theory behind the methods used to evaluate system performance. The discussion begins with the principles underlying observability analysis, which determines whether a system's states can be inferred from the frameworks state space model. Specifically, the tools for analyzing observability, including null-space analysis and the observability Gramian, are presented to characterize estimation limitations. Thereafter, the section summarizes the simulation tools used for model validation, that is the TruckMaker (TM) environment. It then concludes by outlining the function of Monte Carlo (MC) simulations.

2.6.1 Observability in State Estimation

Observability refers to the ability to uniquely determine the state of a system from its outputs over time.

For linear time-invariant systems in discrete time, consider the standard state-space model:

$$\mathbf{x}_k = \mathbf{F}\mathbf{x}_{k-1} + \mathbf{w}_{k-1}, \tag{2.40}$$

$$\mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{v}_k, \tag{2.41}$$

where $\mathbf{x}_k$ is the system state, $\mathbf{z}_k$ the observation, and $\mathbf{F}, \mathbf{H}$ the system and observation matrices. The system is observable if the observability matrix

$$\mathcal{O} = \begin{bmatrix} \mathbf{H} \\ \mathbf{HF} \\ \vdots \\ \mathbf{HF}^{n-1} \end{bmatrix} \tag{2.42}$$

has full rank, equal to the number of state variables n.

For nonlinear systems, observability is often evaluated by linearizing the system around a operating point. This simplifies the analysis and makes it possible to evaluate whether changes in the states lead to noticeable effects in the measurements. In linear systems, observability is assessed by verifying that the observability matrix has full rank. In both linear and linearized cases, observability is a property of the system model even though the result for a linearized system depends on the operating point [19].

2.6.2 Null-space Analysis of Observability

While classical observability analysis determines whether a system is observable as a whole, it does not indicate which specific directions in the state space are unobservable. Null-space analysis addresses this by identifying the subspace of states that are indistinguishable from sensor measurements.

In the context of non-linear systems estimated with an EKF, local observability is typically analyzed using a linearized observability matrix $\mathcal{O}$. If there exists a non-zero vector $\mathbf{x}_{null}$ such that

$$\mathcal{O}\,\mathbf{x}_{null} = 0, \tag{2.43}$$

then $\mathbf{x}_{null}$ lies in the null-space of $\mathcal{O}$ and corresponds to an unobservable direction. Changes along this direction cannot be distinguished from the measurements.

For discrete-time, time-varying nonlinear systems, an alternative approach involves the local observability Gramian:

$$\mathbf{W}_o = \mathcal{O}_k^\top \mathcal{O}_k, \tag{2.44}$$

$$\mathbf{W}_o\,\mathbf{x}_{null} = 0, \tag{2.45}$$

where $\mathbf{W}_o$ is the observability Gramian and k is the current time step.

If $\mathbf{W}_o$ is not full rank, there exists a null-space for $\mathbf{W}_o$ [20]. The rank deficiency of the observability Gramian reveals the number of unobservable directions. However, the significance of these directions depend on how the system evolves. If the state remains orthogonal to the null-space directions, these may have minimal practical effect. In contrast, if the system evolves along an unobservable direction, the state components most aligned with the direction, cannot be reliably estimated.

This type of analysis supports robust filter design by revealing structural limitations of a sensor configuration. It highlights conditions under which states may become unobservable due to limited excitation or insufficient sensor diversity [21].

2.6.3 TruckMaker

TM, developed by IPG Automotive, is a simulation environment used for modeling and testing the behavior of commercial vehicles. It provides a high-fidelity virtual platform in which dynamic driving scenarios can be simulated with accurate representations of physical vehicle behavior.

The environment supports both Software-in-the-Loop and Hardware-in-the-Loop configurations, making it suitable for the development and verification of control algorithms and sensor fusion systems.

The most feature used is the realistic simulation of any given situation with a realistic representation of a heavy-duty vehicle. This provides true values for the road characteristics, velocities, and attitude of the vehicle. It does this with realistic sensor models and complex vehicle-environment interaction [22].

2.6.4 Monte Carlo Simulations

MC simulations are a class of numerical methods that use repeated random sampling to analyze the behavior of systems with uncertainty elements. In the context of sensor fusion and state estimation, the method provides a way to evaluate the performance of algorithms. This is particularly important in this thesis since the sensors experience measurement noise and unknown sensor biases [7].

In theoretical terms, MC simulations allow an empirical assessment of a framework given randomly chosen biases and noise variance. The approach involves generating multiple independent realizations of the same simulation, each influenced by randomized noise and biases. The method is reliable in cases where the random elements are the only varying elements, while the deterministic variables are known. By analyzing the resulting state estimates, it can be observed how closely the predicted results align with the true values over several realizations of the simulation [23].

MC methods are especially useful when analytical tools such as observability analysis are not practical to use due to system non-linearity or complexity.

3

Method

The methodology that was used in this thesis was structured into a set of development stages, covering both the construction of the sensor fusion framework and the preparation of simulation data for evaluation. The filter design process was divided into three main phases, while the simulation setup and the data collection were addressed in two parallelly developed phases. An overview of the development workflow is presented in Figure 3.1.

To construct the complete fusion filter, the sensors were first modeled to link the measurements to real phenomena. The radar sensors were modeled using a kinematic approach, producing both measurement equations and a pseudo measurement. This modeling is described in Section 3.1.

Simultaneously, an EKF was implemented for processing data from the IMU. This intermediate filter functioned as a standalone Inertial Navigation System (INS) filter, used to estimate motion states from IMU measurements alone. It served as the base for the full fusion filter, and is presented in Section 3.2. After the completion of the sensor modeling, the complete fusion filter was constructed by combining the radar and IMU processing pipelines. The full filter structure and formulation is described in Section 3.3.

Simulation data was used throughout the development process to test and validate the various system components. This data was generated using the TM simulation environment, which required custom design of test tracks and subsequent export of simulation outputs. The procedure for designing these test tracks is presented in Section 3.4.4, and the corresponding data collection and evaluation strategy is detailed in Section 3.4.5.

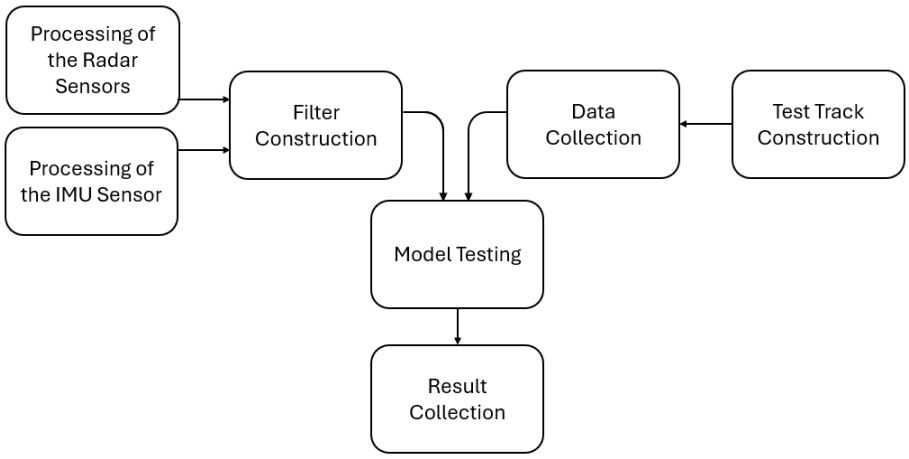

Figure 3.1: System development workflow and structure.

3.1 Radar Sensor Processing

Accurate interpretation of radar measurements depend on a thorough understanding of how the sensor was configured and how it interacted with the vehicle's motion. The theoretical background supporting this modeling is presented in Subsection 2.4.3.

The radar system employed in this thesis was an array of FMCW radars, where the individual radars were mounted on the vehicle and oriented toward the ground. This setup allowed the radar to measure both the range and the Doppler velocity of a point on the ground intersected by the radar's bore axis, that is the axis that extruded straight out of the sensor. Under the assumption of proper signal processing, only the radar's mounting position and orientation relative to the vehicle body frame affect the radar measurements.

3.1.1 Radar Kinematics and Placement

To transform radar measurements into the body frame, the relative positioning of the radar sensors had to be defined. This was achieved by using the DH convention. Furthermore, the use of the DH convention allowed for future extensions where the cab may be modeled as a dynamic system.

Each radar sensor was described as a three-joint system. One from the body frame to the sensor mount, and another from the sensor mount to the sensor head. The corresponding DH parameters are listed in Table 3.1.

Table 3.1: DH parameters for each radar sensor, modeled as sequential link transformations from the vehicle body frame to the sensor head.

Sensor	**Link**	$\boldsymbol{a}$[m]	$\boldsymbol{\alpha}$[°]	$\boldsymbol{d}$[m]	$\boldsymbol{\theta}$[°]	**Description**
Front Left	1	5.1	0.0	0.0	11.3	Body → Sensor Mount
	2	0.0	0.0	0.0	33.7	Sensor Mount → Sensor Head Yaw
	3	0.0	45.0	0.0	0.0	Sensor Head Yaw → Sensor Head
Front Right	1	5.1	0.0	0.0	−11.3	Body → Sensor Mount
	2	0.0	0.0	0.0	−33.7	Sensor Mount → Sensor Head Yaw
	3	0.0	45.0	0.0	0.0	Sensor Head Yaw → Sensor Head
Rear Right	1	1.4	0.0	0.0	−135.0	Body → Sensor Mount
	2	0.0	0.0	0.0	0.0	Sensor Mount → Sensor Head Yaw
	3	0.0	45.0	0.0	0.0	Sensor Head Yaw → Sensor Head
Rear Left	1	1.4	0.0	0.0	135.0	Body → Sensor Mount
	2	0.0	0.0	0.0	0.0	Sensor Mount → Sensor Head Yaw
	3	0.0	45.0	0.0	0.0	Sensor Head Yaw → Sensor Head

These parameters were used to compute the position of the radar sensor head and the ground point it observes relative to the body frame. These distances were computed, within the body frame coordinate system, as follows:

- Position of the sensor mount relative to the body frame:

$$\mathbf{r}_s = \begin{bmatrix} a_1 \sin(\frac{\pi}{2} - \alpha_1)\cos(\theta_1) \\ a_1 \sin(\frac{\pi}{2} - \alpha_1)\sin(\theta_1) \\ a_1 \cos(\frac{\pi}{2} - \alpha_1) \end{bmatrix}. \tag{3.1}$$

- Position of the road intersection point relative to the sensor mount:

$$\mathbf{r}_g = \begin{bmatrix} a_g \sin(\frac{\pi}{2} - \alpha_3)\cos(\theta_2) \\ a_g \sin(\frac{\pi}{2} - \alpha_3)\sin(\theta_2) \\ a_g \cos(\frac{\pi}{2} - \alpha_3) \end{bmatrix}. \tag{3.2}$$

Note that a_g denotes the projection distance along the radar's bore axis and was not a DH parameter, as it canceled out when calculating the magnitude of the projected Doppler velocity. The resulting sensor positions based upon these parameters are visualized in Figure 3.2.

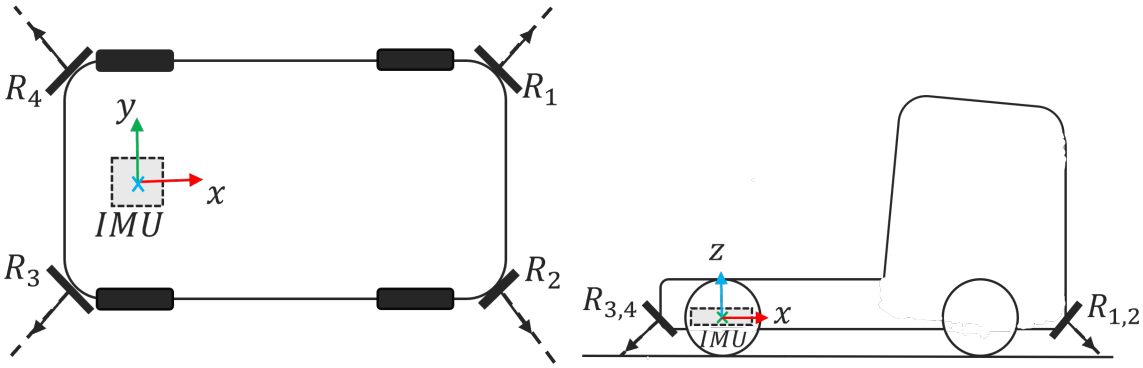

(a) Top view: Sensor Positioning **(b)** Side view: Sensor Positioning

Figure 3.2: Visualization of the sensor positions from two perspectives.

3.1.2 Doppler Velocity Equation

The relationship between the measured Doppler velocity and the vehicle's translational and rotational velocities was given by

$$\mathbf{v}_d^{(b)} = -\mathbf{v}^{(b)} - \boldsymbol{\omega}^{(b)} \times (\mathbf{r}_s + \mathbf{r}_g), \tag{3.3}$$

where $\mathbf{v}_d^{(b)}$ is the true relative velocity vector, $\mathbf{v}^{(b)}$ is the body frame translational velocity, and $\boldsymbol{\omega}^{(b)}$ is the body frame angular velocity vector.

To acquire the radar's actual measurement, this vector had to be projected along the radar's bore axis

$$Pr_{\hat{\mathbf{r}}_g}(\mathbf{v}_d) = \frac{\mathbf{v}_d^{(b)} \cdot \hat{\mathbf{r}}_g}{\|\hat{\mathbf{r}}_g\|^2}. \tag{3.4}$$

This resulted in a model that directly linked body frame velocities to the radar measurements.

For all n_{radar} sensors, the relationship was expressed in matrix form as

$$Pr_{\hat{\mathbf{r}}_g}(\mathbf{v}_d) = \mathbf{C}_{n_{radar} \times 6} \begin{bmatrix} \mathbf{v}^{(b)} \\ \boldsymbol{\omega}^{(b)} \end{bmatrix} + \boldsymbol{\eta}_r, \tag{3.5}$$

where $\mathbf{C}$ is an $n_{radar} \times 6$ matrix of projection coefficients. To recover the state from radar measurements, the Moore-Penrose inverse of $\mathbf{C}$ was used,

$$\begin{bmatrix} \mathbf{v}^{(b)} \\ \boldsymbol{\omega}^{(b)} \end{bmatrix} = \mathbf{C}^{\dagger}_{n_{radar} \times 6} Pr_{\hat{\mathbf{r}}_g}(\mathbf{v}_d) + \boldsymbol{\eta}_r. \tag{3.6}$$

(3.5) were employed in both the radar measurement generation and the final filter implementation, and (3.6) were used for testing and validation.

3.1.3 Pseudo Measurement of Vertical Velocity

To estimate the road bank and slope with regard to the navigation frame (ϕ_n^r, θ_n^r), a pseudo measurement was introduced based on the assumption of negligible vertical motion in the body frame when its z-axis was parallel with the z-axis of the road frame. This pseudo measurement was used to estimate the body attitude with regard to the road frame (ϕ_r^b, θ_r^b), which was used together with the body attitude with regard to the navigation frame to calculate the road bank and slope with regard to the navigation frame.

The velocity in the body frame was rotated to force its z-axis parallel with the road frame using

$$\mathbf{v}^{(r)} = \mathbf{R}_{y,\theta_r^b}\, \mathbf{R}_{x,\phi_r^b}\, \mathbf{v}^{(b)}, \tag{3.7}$$

where $\mathbf{R}_{x,\phi_r^b}$ and $\mathbf{R}_{y,\theta_r^b}$ are standard rotation matrices, as described in Section 2.2.3, for roll and pitch, respectively.

When assuming zero vertical velocity in the road frame, the pseudo measurement was expressed by

$$0 = v_z^{(r)} + \eta_{pseudo}, \tag{3.8}$$

where $v_z^{(r)}$ is the z component of the velocity vector in the road frame. This equation acted as a nonlinear model used to estimate the vehicle attitude with respect to road in the EKF. Note there is no noise term in this quantity since it is included in the estimated variables.

3.2 IMU Processing and the INS Filter

The IMU used in this thesis was a six DoF sensor with three orthogonal accelerometers and three orthogonal gyroscopes. Several approaches exist for fusing measurements from IMUs, including complementary filters [24] and various KF implementations [25, 26, 27]. In this thesis, an EKF was used, since it can handle systems with moderate non-linearities.

The INS filter served as the base component of the sensor fusion framework. As such, detailed tuning and noise modeling are deferred to the fusion filter. Therefore, this section omits explicit covariance matrices, as the INS filter was intended primarily for validation and development purposes.

3.2.1 INS State Definitions

The choice of state variables in the INS filter was determined by the quantities to be estimated, which in this case, were the vehicle's velocity, acceleration, attitude, angular rates, and IMU biases. These quantities are visualized in Figure 3.3, however, the figure includes all quantities that the final filter ought to measure and therefore has more variables percent than in the INS state vector realization.

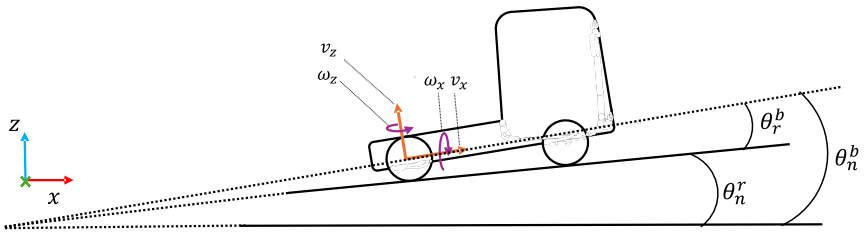

(a) Side view: State variables

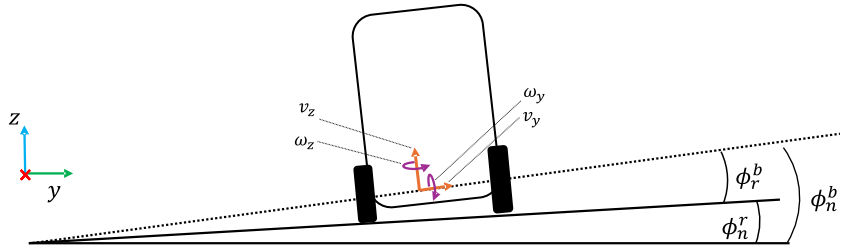

(b) Front view: State variables

Figure 3.3: Visualization of the state variables from two perspectives.

Since the system operated without direct control inputs, the state vector was defined as:

$$\mathbf{x}_{ins} = \begin{bmatrix} v_x^{(b)} \\ v_y^{(b)} \\ v_z^{(b)} \\ \dot{v}_x^{(b)} \\ \dot{v}_y^{(b)} \\ \dot{v}_z^{(b)} \\ \phi_n^b \\ \theta_n^b \\ \omega_x^{(b)} \\ \omega_y^{(b)} \\ \omega_z^{(b)} \\ b_{a_x} \\ b_{a_y} \\ b_{a_z} \\ b_{g_x} \\ b_{g_y} \\ b_{g_z} \end{bmatrix} \tag{3.9}$$

where $v_{x,y,z}^{(b)}$ are the body frame translational velocities, $\dot{v}_{x,y,z}^{(b)}$ the translational accelerations, ϕ_n^b and θ_n^b the roll and pitch angles with respect to the navigation frame, $\omega_{x,y,z}^{(b)}$ the rotational velocity components, and $b_{a_{x,y,z}}$ and $b_{g_{x,y,z}}$ the accelerometer and gyroscope biases, respectively.

3.2.2 INS Process Model

The process model was structured to independently capture the dynamics of translational and rotational motion. The translational states were modeled with the assumption of constant acceleration, while the attitude dynamics were modeled using non-linear relationships between angular rates and Euler angles. This Euler angle propagation model was based on the transformation between angular and Euler rates detailed in Section 2.2.4. The discrete-time prediction model was expressed by

$$\mathbf{v}_k^{(b)} = \mathbf{v}_{k-1}^{(b)} + \dot{\mathbf{v}}_{k-1}^{(b)}\,\Delta t + \boldsymbol{\eta}_v, \tag{3.10}$$

$$\dot{\mathbf{v}}_k^{(b)} = \dot{\mathbf{v}}_{k-1}^{(b)} + \boldsymbol{\eta}_{\dot{v}}, \tag{3.11}$$

$$\begin{aligned}\phi_{n_k}^b &= \phi_{n_{k-1}}^b + \omega_{x_{k-1}}^{(b)}\Delta t + \omega_{y_{k-1}}^{(b)}\sin\left(\phi_{n_{k-1}}^b\right)\tan\left(\theta_{n_{k-1}}^b\right)\Delta t \\ &\quad + \omega_{z_{k-1}}^{(b)}\cos\left(\phi_{n_{k-1}}^b\right)\tan\left(\theta_{n_{k-1}}^b\right)\Delta t + \eta_{\phi_n^b},\end{aligned} \tag{3.12}$$

$$\theta_{n_k}^b = \theta_{n_{k-1}}^b + \omega_{y_{k-1}}^{(b)}\cos\left(\phi_{n_{k-1}}^b\right)\Delta t - \omega_{z_{k-1}}^{(b)}\sin\left(\phi_{n_{k-1}}^b\right)\Delta t + \eta_{\theta_n^b}, \tag{3.13}$$

$$\boldsymbol{\omega}_k^{(b)} = \boldsymbol{\omega}_{k-1}^{(b)} + \boldsymbol{\eta}_\omega, \tag{3.14}$$

$$\mathbf{b}_{a_k} = \mathbf{b}_{a_{k-1}} + \boldsymbol{\eta}_{b_a}, \tag{3.15}$$

$$\mathbf{b}_{g_k} = \mathbf{b}_{g_{k-1}} + \boldsymbol{\eta}_{b_g}. \tag{3.16}$$

3.2.3 INS Measurement Model

The measurement model was utilized in the update step which was used to update state estimates based on the available IMU measurements, specifically the measured angular velocities from the gyroscopes and translational accelerations from the accelerometers. Due to sensor imperfections, these measurements deviated from the true values.

The primary sources of deviation from the true values were the gravitational acceleration, the centripetal acceleration due to rotation, and the sensor biases. All three affected the accelerometers whilst only the sensor biases affected the gyroscopes. Effects such as Euler and rectilinear accelerations were considered negligible in this context and are therefore omitted. The Coriolis force, which is only apparent in non-inertial reference frames when observed from an inertial frame, was not measured by the sensor and was thus also excluded. The accelerometer and gyroscope models were therefore expressed by

$$\mathbf{y}_a = \dot{\mathbf{v}}^{(b)} + \begin{bmatrix} -\sin(\theta_n^b) \\ \sin(\phi_n^b)\cos(\theta_n^b) \\ \cos(\phi_n^b)\cos(\theta_n^b) \end{bmatrix} 9.81 - \begin{bmatrix} v_y^{(b)}\omega_z^{(b)} - v_z^{(b)}\omega_y^{(b)} \\ v_z^{(b)}\omega_x^{(b)} - v_x^{(b)}\omega_z^{(b)} \\ v_x^{(b)}\omega_y^{(b)} - v_y^{(b)}\omega_x^{(b)} \end{bmatrix} + \begin{bmatrix} b_{a_x} \\ b_{a_y} \\ b_{a_z} \end{bmatrix} + \boldsymbol{\eta}_a, \tag{3.17}$$

$$\mathbf{y}_\omega = \begin{bmatrix} \omega_x^{(b)} \\ \omega_y^{(b)} \\ \omega_z^{(b)} \end{bmatrix} + \begin{bmatrix} b_{g_x} \\ b_{g_y} \\ b_{g_z} \end{bmatrix} + \boldsymbol{\eta}_\omega. \tag{3.18}$$

Accordingly, the complete INS measurement model used during the EKF update step was

$$\mathbf{y}_a = \dot{\mathbf{v}}^{(b)} + \mathbf{g}^{(b)}(\phi_n^b, \theta_n^b) - \mathbf{v}^{(b)} \times \boldsymbol{\omega}^{(b)} + \mathbf{b}_a + \boldsymbol{\eta}_a, \tag{3.19}$$

$$\mathbf{y}_\omega = \boldsymbol{\omega}^{(b)} + \mathbf{b}_g + \boldsymbol{\eta}_\omega. \tag{3.20}$$

3.3 Fusion Filter Design

The fusion filter was based on the previously introduced INS filter. In contrast to the INS filter setup, the fusion filter incorporated both radar and IMU measurements. However, since radar measurements were received at a lower rate than the IMU measurements, the filter had to accommodate multirate data availability and maintain state estimates at a frequency equivalent to that of the IMU. This requirement resulted in the filter structure illustrated in Figure 3.4.

The switch in Figure 3.4 represents whether radar measurements were available at a given time step. When available, radar measurements were used alongside IMU measurements to update the state estimates. Otherwise, the INS update was performed.

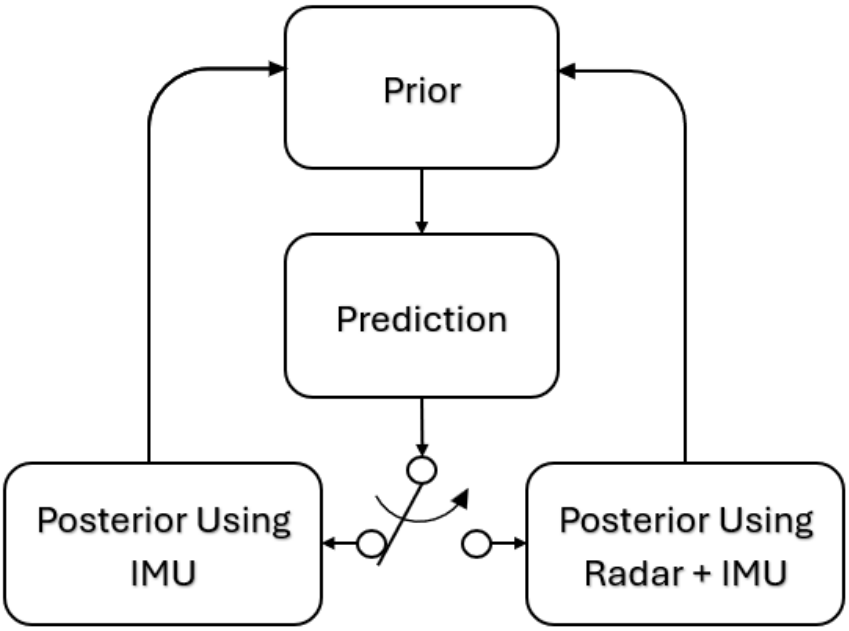

Figure 3.4: Flowchart for the filter structure.

3.3.1 State Vector and Process Model

The fusion filter's state vector expanded upon the INS filter by introducing two additional variables: the body roll and pitch angles with respect to the road frame, denoted ϕ_r^b and θ_r^b. Due to observability issues described in the results, ϕ_r^b was removed and modeled as a Gaussian random variable, $\phi_r^b \sim \mathcal{N}(0, 0.03)$. The pitch angle was modeled as a random walk variable and included in the state vector. The state vector thus became

$$\mathbf{x}_{complete} = \begin{bmatrix} v_x^{(b)} \\ v_y^{(b)} \\ v_z^{(b)} \\ \dot{v}_x^{(b)} \\ \dot{v}_y^{(b)} \\ \dot{v}_z^{(b)} \\ \phi_n^b \\ \theta_n^b \\ \theta_r^b \\ \omega_x^{(b)} \\ \omega_y^{(b)} \\ \omega_z^{(b)} \\ b_{a_x} \\ b_{a_y} \\ b_{a_z} \\ b_{g_x} \\ b_{g_y} \\ b_{g_z} \end{bmatrix}. \tag{3.21}$$

These variables are visualized in Figure 3.3.

The prediction model closely mirrored that of the INS filter. The two radar orientation angles were treated as constant during prediction, consistent with the random walk assumption. The discrete-time process model thus became

$$\mathbf{v}_k^{(b)} = \mathbf{v}_{k-1}^{(b)} + \dot{\mathbf{v}}_{k-1}^{(b)} \Delta t + \boldsymbol{\eta}_v, \tag{3.22}$$

$$\dot{\mathbf{v}}_k^{(b)} = \dot{\mathbf{v}}_{k-1}^{(b)} + \boldsymbol{\eta}_{\dot{v}}, \tag{3.23}$$

$$\begin{aligned} \phi_{n_k}^b = {} & \phi_{n_{k-1}}^b + \omega_{x_{k-1}}^{(b)} \Delta t + \omega_{y_{k-1}}^{(b)} \sin\left(\phi_{n_{k-1}}^b\right) \tan\left(\theta_{n_{k-1}}^b\right) \Delta t \\ & + \omega_{z_{k-1}}^{(b)} \cos\left(\phi_{n_{k-1}}^b\right) \tan\left(\theta_{n_{k-1}}^b\right) \Delta t + \eta_{\phi_n^b}, \end{aligned} \tag{3.24}$$

$$\theta_{n_k}^b = \theta_{n_{k-1}}^b + \omega_{y_{k-1}}^{(b)} \cos\left(\phi_{n_{k-1}}^b\right) \Delta t - \omega_{z_{k-1}}^{(b)} \sin\left(\phi_{n_{k-1}}^b\right) \Delta t + \eta_{\theta_n^b}, \tag{3.25}$$

$$\theta_{r_k}^b = \theta_{r_{k-1}}^b + \eta_{\theta_r^b}, \tag{3.26}$$

$$\boldsymbol{\omega}_k^{(b)} = \boldsymbol{\omega}_{k-1}^{(b)} + \boldsymbol{\eta}_\omega, \tag{3.27}$$

$$\mathbf{b}_{a_k} = \mathbf{b}_{a_{k-1}} + \boldsymbol{\eta}_{b_a}, \tag{3.28}$$

$$\mathbf{b}_{g_k} = \mathbf{b}_{g_{k-1}} + \boldsymbol{\eta}_{b_g}. \tag{3.29}$$

3.3.2 Multirate Estimation Strategy

To account for the multirate arrival of sensor measurements, the fusion filter employed two forked update steps:

- **INS update (high rate):** Applied when only IMU data was available. The prediction step was followed by a state update using the same model as in the INS filter.

- **Joint Radar–IMU update (low rate):** Triggered when radar data was available. Both radar and IMU measurements were used for the update, correcting drift and improving velocity and attitude estimates.

This forked update structure ensured that the filter output estimates at the rate of the IMU. It did this without significantly compromising performance due to the drift-free nature of radar measurements [7, 28].

3.3.3 Measurement Model

Two measurement models were implemented:

- **IMU model:** Used when radar data was not available. The IMU was assumed to be mounted at a α° pitch downward and β° roll to the right. α and β were set depending on how the IMU was to be configured. In this thesis, two configurations were used, one where $\alpha = \beta = 0^\circ$ and one where $\alpha = \beta = -45^\circ$. The rotation matrix was thus

$$\mathbf{R}_{\beta\alpha} = \mathbf{R}_{y,\beta}\mathbf{R}_{x,\alpha}. \tag{3.30}$$

  The accelerometer and gyroscope measurements were modeled as

$$\mathbf{y}_a = \mathbf{R}_{\beta\alpha}\left(\dot{\mathbf{v}}^{(b)} + \mathbf{g}^{(b)}(\phi_n^b, \theta_n^b) - \mathbf{v}^{(b)} \times \boldsymbol{\omega}^{(b)}\right) + \mathbf{b}_a + \boldsymbol{\eta}_a, \tag{3.31}$$

$$\mathbf{y}_\omega = \mathbf{R}_{\beta\alpha}\boldsymbol{\omega}^{(b)} + \mathbf{b}_g + \boldsymbol{\eta}_\omega. \tag{3.32}$$

- **Radar–IMU model:** Used when radar data was available. In addition to (3.31) and (3.32), radar Doppler measurements and a pseudo vertical velocity constraint were included resulting the extension of two additional equations

$$\mathbf{y}_r = \mathbf{C}\begin{bmatrix}\mathbf{v}^{(b)} \\ \boldsymbol{\omega}^{(b)}\end{bmatrix} + \boldsymbol{\eta}_r, \tag{3.33}$$

$$0 = v_z^{(r)} + \eta_{pseudo}. \tag{3.34}$$

3.3.4 Road Orientation Estimation and Covariance Propagation

The relative orientation between the road surface and the inertial frame was captured by the angles ϕ_n^r and θ_n^r, representing the road roll and pitch with regard to the navigation frame. The nominal values of these were the results from rotations around the x and y axes, respectively, and given by:

$$\mathbf{R}_{\phi_n^r} = \mathbf{R}_{\phi_r^b}^{-1}\mathbf{R}_{\phi_n^b}, \tag{3.35}$$

$$\mathbf{R}_{\theta_n^r} = \mathbf{R}_{\theta_r^b}^{-1}\mathbf{R}_{\theta_n^b}, \tag{3.36}$$

whereafter $\mathbf{R}_{\phi_n^r}$ and $\mathbf{R}_{\theta_n^r}$ were converted into Euler angles.

Since the body attitude with regard to the navigation frame (ϕ_n^b, θ_n^b) and the body attitude with regard to the road frame (ϕ_r^b, θ_r^b) were estimated quantities, it was necessary to take their respective uncertainties into account when computing the road attitude with regard to the navigation frame. To find the uncertainty of the derived road attitude angles, the covariances of the estimated attitudes were combined. The formula used considered both the variance of the individual angles and the cross-covariances between them.

To acquire the formula for (3.35) and (3.36), the respective Jacobians $\mathbf{H}_\phi$ and $\mathbf{H}_\theta$ are needed as well as the respective covariance matrices $\mathbf{P}_\phi$ and $\mathbf{P}_\theta$.

If the elements of the rotation matricies are defined as

$$\mathbf{R}_{\phi_n^r} = \begin{bmatrix} \mathbf{R}_{\phi_n^r}(1,1) & \mathbf{R}_{\phi_n^r}(1,2) & \mathbf{R}_{\phi_n^r}(1,3) \\ \mathbf{R}_{\phi_n^r}(2,1) & \mathbf{R}_{\phi_n^r}(2,2) & \mathbf{R}_{\phi_n^r}(2,3) \\ \mathbf{R}_{\phi_n^r}(3,1) & \mathbf{R}_{\phi_n^r}(3,2) & \mathbf{R}_{\phi_n^r}(3,3) \end{bmatrix}, \tag{3.37}$$

$$\mathbf{R}_{\theta_n^r} = \begin{bmatrix} \mathbf{R}_{\theta_n^r}(1,1) & \mathbf{R}_{\theta_n^r}(1,2) & \mathbf{R}_{\theta_n^r}(1,3) \\ \mathbf{R}_{\theta_n^r}(2,1) & \mathbf{R}_{\theta_n^r}(2,2) & \mathbf{R}_{\theta_n^r}(2,3) \\ \mathbf{R}_{\theta_n^r}(3,1) & \mathbf{R}_{\theta_n^r}(3,2) & \mathbf{R}_{\theta_n^r}(3,3) \end{bmatrix}. \tag{3.38}$$

Then the equations used to find the Jacobians were

$$\mathbf{H}_\phi = \begin{bmatrix} \frac{\partial arctan\left(\frac{\mathbf{R}_{\phi_n^r}(3,2)}{\mathbf{R}_{\phi_n^r}(2,2)}\right)}{\partial \phi_n^b} & \frac{\partial arctan\left(\frac{\mathbf{R}_{\phi_n^r}(3,2)}{\mathbf{R}_{\phi_n^r}(2,2)}\right)}{\partial \phi_r^b} \end{bmatrix}, \tag{3.39}$$

$$\mathbf{H}_\theta = \begin{bmatrix} \frac{\partial arctan\left(\frac{\mathbf{R}_{\theta_n^r}(1,3)}{\mathbf{R}_{\theta_n^r}(1,1)}\right)}{\partial \theta_n^b} & \frac{\partial arctan\left(\frac{\mathbf{R}_{\theta_n^r}(1,3)}{\mathbf{R}_{\theta_n^r}(1,1)}\right)}{\partial \theta_r^b} \end{bmatrix}, \tag{3.40}$$

The equations used to find the covariances were

$$\mathbf{P}_\phi = \begin{bmatrix} \sigma_{\phi_n^b} & \sigma_{\phi_n^b,\phi_r^b} \\ \sigma_{\phi_r^b,\phi_n^b} & \sigma_{\phi_r^b} \end{bmatrix} \tag{3.41}$$

$$\mathbf{P}_\theta = \begin{bmatrix} \sigma_{\theta_n^b} & \sigma_{\theta_n^b,\theta_r^b} \\ \sigma_{\theta_r^b,\theta_n^b} & \sigma_{\theta_r^b} \end{bmatrix} \tag{3.42}$$

Using these variables, the covariance matrices for the rotated variables were

$$\mathbf{P}_{\phi_n^r} = \mathbf{H}_\phi \mathbf{P}_\phi \mathbf{H}_\phi^\top, \tag{3.43}$$

$$\mathbf{P}_{\theta_n^r} = \mathbf{H}_\theta \mathbf{P}_\theta \mathbf{H}_\theta^\top. \tag{3.44}$$

The resulting matrices were subsequently used to evaluate the uncertainty of the road attitude with regard to the navigation frame. It important to note that since, in this thesis' filter realization, ϕ_r^b was unobservable, thus only (3.44) was used.

3.3.5 Noise Modeling

Three noise covariance matrices were defined and manually tuned:

$$\Sigma_{\mathbf{x}} = \text{Process noise covariance}$$
$$\Sigma_{\mathbf{y_1}} = \text{Measurement noise covariance with IMU}$$
$$\Sigma_{\mathbf{y_2}} = \text{Measurement noise covariance with radar and IMU}$$

These were empirically tuned, based on the expected measurement noise levels and presented in Appendix B.

3.4 Simulation Framework and Evaluation Setup

The simulation framework developed in this thesis served as the base for evaluating the performance and robustness of the proposed state estimation filter. Since real-world sensor experiments were not within the scope of this study, a simulation-based approach was adopted using TM to emulate vehicle dynamics and sensor behavior under controlled and repeatable conditions.

This section outlines the methodology used to generate noisy radar and IMU measurements, evaluate different sensor configurations, and assess the filter's performance using a MC-based analysis. It also details the test scenarios used to simulate different driving conditions and describes the metrics applied to quantify estimation performance.

3.4.1 Sensor Modeling and Simulation Setup

To generate realistic sensor measurements for evaluation, radar and IMU measurements were found using the ground-truth motion states provided by TM. Although TM supports virtual sensor emulation through RSIs, prior investigations [6] have demonstrated that RSI-based radar modeling is inaccurate when radars are pitched toward the ground to measure surface-relative motion. Therefore, a hybrid simulation strategy was adopted to address this limitation.

At each simulation time-step, the true vehicle motion was projected onto the radar's bore axis direction using (3.5), as explained in Section 2.4.4. This ensured that the projected Doppler velocity outputs were consistent with both translational and rotational motion. Sensor placement geometry was parameterized using the DH conversion and implemented using MATLAB. The radar velocity measurements were assumed to be bias-free. Nonetheless, Gaussian measurement noise was included to reflect environmental and hardware variability.

Inertial data was obtained from TM's built-in ideal six DoF IMU model. Since the raw IMU outputs were noise and bias free, these disturbances were applied in post-processing. Gaussian white noise was added independently to each accelerometer and gyroscope channel to model white sensor noise, such as jitter and thermal

noise. In addition, constant bias offsets were applied to simulate static internal calibration and manufacturing errors. These bias values were randomly sampled from set intervals, of $\left[-0.2, 0.2\right]$ m/s^2 for the accelerometers and $\left[-0.1, 0.1\right]$ rad/s for the gyroscopes, at the start of each simulation run and held constant throughout.

This combined simulation approach enabled repeatable scenario testing, consistent MC experimentation, and targeted analysis of estimation performance under known sensor error conditions.

3.4.2 Radar Configuration Evaluation

To assess the minimum number of radar sensors required for body-frame velocity estimation, a series of simulations were conducted with different sensor configurations. Specifically, setups using one, two, three, and four radars were evaluated to determine how sensor count influenced estimation accuracy and robustness.

For each case, a single scenario was used and used on the fusion filter. The only difference between the runs was the number of radars. The evaluation metrics, called the Figures of Merit (FoM), RMSE, MAE, and MaxAE described in (3.45), (3.46), and (3.47), respectively, were used to compare the runs and to evaluate the filter performance.

3.4.3 Observability Analysis via Null-space Accumulation

To characterize the limitations of the filter, an observability analysis was performed using the null-space of the discrete-time observability Gramian. In non-linear and time-varying systems, the Gramian $\mathbf{W}_o$ is often not full rank due to partial system excitation. Consequently, this study focused on analyzing the accumulation of unobservable directions over time.

At each time-step, the system process and measurement linearized Jacobians were computed to form the local observability matrix. A local Gramian was then constructed, and its null-space was extracted to identify state directions that are unobservable. All local null-space directions were summed an normalized to acquire the average directions over time. This was used to identify what states were unobservable and how system changes affected which states were most effected.

To improve the observability of accelerometer biases, particularly along the x and y axes, the IMU was mounted with a deliberate $-45°$ pitch and roll. This orientation was hypothesized tp distribute the gravitational vector and the angular changes across all three accelerometer channels, enhancing the filers ability to estimate the b_{a_x} and b_{a_y} bias states.

3.4.4 Driving Scenario Design

To evaluate estimation performance under diverse driving conditions, a range of test tracks were developed with TM. These tracks were designed to test the filter on varied tracks with varied vehicle maneuvers, including straight-line motion, acceleration events, and different turning maneuvers.

Five tracks were used in total:

- **Steady-State Cruise Track:** A straight, level road on which the vehicle travels at constant velocity with no acceleration or turning. This track isolated drift and sensor bias effects under steady-state conditions.

- **Varied Velocity Track:** A flat road with repeated cycles of braking and acceleration. The vehicle periodically decelerates to low speed before resuming motion. This scenario emphasized the filter's ability to handle abrupt changes in velocity.

- **Uphill Track:** The vehicle accelerates along a flat road, transitions onto a 10° incline, and then continues onto a level segment. This track tested the filter's ability to estimate changes in the vehicles pitch angle without changes in other attitude states.

- **Dynamic J-Turn Track:** A straight segment is followed by a sharp, 5° banked J-turn. The maneuver has strong lateral acceleration and moderate roll. This track tested the filter's ability to estimate changes in the vehicles roll angle with small changes in other attitude states.

- **Infinity with Bank Track:** A figure-eight track with banked curves at each loop, simulating continuous coupled roll and yaw dynamics. This track tested the filter's ability to handle turn transitions and complex motion.

3.4.5 Evaluation Metrics

The FoM used to evaluate the filter's performance were computed for each state variable across all runs and scenarios. These are explained in greater detail in [29].

The FoM that were relevant in this thesis and their corresponding formulas are:

- **Root Mean Square Error (RMSE):** Measures average squared deviation, emphasizing larger errors:

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{k=1}^{N}(\hat{x}_k - x_k)^2}, \tag{3.45}$$

- **Mean Absolute Error (MAE):** Measures average absolute deviation, ro-

bust to outliers:

$$\mathrm{MAE} = \frac{1}{N} \sum_{k=1}^{N} |\hat{x}_k - x_k| \,, \tag{3.46}$$

- **Maximum Absolute Error (MaxAE):** Captures worst-case deviation across the simulation horizon:

$$\mathrm{MaxAE} = \max_k |\hat{x}_k - x_k| \,. \tag{3.47}$$

4
Results and Discussion

This chapter presents and analyzes the key outcomes of the proposed method, linking them to the research questions and design choices outlined earlier. The chapter is divided into five sections, the first two explain foundational design decisions, while the remaining three assess the filter's performance under increasingly realistic and challenging configurations. All results are influenced by stochastic sensor noise.

Section 4.1 explores the impact of the number of radar sensors on estimation quality, validating the use of four radars in terms of observability and robustness. In Section 4.2, the exclusion of ϕ_r^b from the state vector is justified, supported by evidence demonstrating its non-observability and negligible effect on performance.

The filter performance evaluation begins with Section 4.3, establishing a baseline scenario free of sensor or state biases. This serves as a reference to understand the filter's behavior under ideal conditions. Section 4.4 introduces realistic sensor biases and analyzes the filter's ability to estimate and compensate for them across various driving scenarios.

The final section, Section 4.5, investigates a tilted IMU configuration, revealing how the IMU orientation affects bias observability and filter convergence. This section highlights the practical implications of sensor placement on filter performance in real-world conditions.

4.1 Number of Radars

This section evaluates how the number of ground-speed radars influences the observability and accuracy of the state filter, particularly in estimating three dimensional translational and angular velocities. Monte Carlo simulations were conducted using one, two, three, and four radars, and the results were statistically aggregated across runs. Table 4.1 presents the RMSE, MAE, and MaxAE metrics for the translational and angular velocity variables.

The results revealed that estimation performance improved significantly with additional radars, especially for translational velocity. The transition from two to three radars lead to a large error metric reduction, reflecting improved observability and reduced velocity uncertainty. When only a single radar was used, the system became unobservable, as indicated by high RMSE and MaxAE values across all ve-

locity components.

Beyond three radars, the improvements increased marginally. This suggested a saturation effect once the system became sufficiently determined. These findings aligned with prior studies [6], which indicated that three well-placed radars was sufficient to recover three dimensional translational and angular velocity from radar measurements.

It is worth noting that angular velocity estimates remained largely unaffected by the radar count. This was expected since angular rates were primarily derived from the IMU's gyroscopes, which operates independently of the radar configuration. In contrast, accurate translational velocity estimation depended heavily on high-quality measurements from multiple radar units.

Importantly, this study fused radar and IMU measurements in a fusion-based estimation framework. While the IMU contributed robustness, the radar system was critical for achieving low-error velocity estimates.

Table 4.1: FoM for translational and angular velocities across radar configurations.

Radars	Category	Variable	RMSE	MAE	MaxAE
1	Translational Velocities [m/s]	v_x	414.4425	351.2765	730.7039
		v_y	550.3641	458.8107	1008.0860
		v_z	348.9044	299.9092	592.0472
	Rotational Velocities [rad/s]	ω_x	0.0012	0.0009	0.0061
		ω_y	0.0020	0.0013	0.0201
		ω_z	0.0036	0.0014	0.0412
2	Translational Velocities [m/s]	v_x	65.4013	46.0968	151.6317
		v_y	0.0751	0.0584	0.3578
		v_z	106.4270	75.0048	246.7247
	Rotational Velocities [rad/s]	ω_x	0.0015	0.0012	0.0067
		ω_y	0.0021	0.0014	0.0200
		ω_z	0.0038	0.0018	0.0416
3	Translational Velocities [m/s]	v_x	0.0360	0.0281	0.2152
		v_y	0.0720	0.0560	0.3438
		v_z	0.0455	0.0352	0.2822
	Rotational Velocities [rad/s]	ω_x	0.0017	0.0014	0.0069
		ω_y	0.0022	0.0015	0.0202
		ω_z	0.0038	0.0018	0.0414
4	Translational Velocities [m/s]	v_x	0.0324	0.0252	0.1976
		v_y	0.0637	0.0494	0.3205
		v_z	0.0376	0.0289	0.2548
	Rotational Velocities [rad/s]	ω_x	0.0017	0.0014	0.0069
		ω_y	0.0022	0.0015	0.0202
		ω_z	0.0038	0.0018	0.0413

4.2 Exclusion of ϕ_r^b

Evaluation of the system showed that the roll angle relative to the road, ϕ_r^b, was consistently unobservable across all tested configurations. This included scenarios with and without sensor biases, and across various driving tracks. The only exception occurred under conditions of sustained and significant lateral velocity, which are uncommon during nominal operation of large heavy-duty vehicles.

While small amounts of lateral motion slightly improved the certainty of ϕ_r^b, the improvement was insufficient. In realistic driving conditions, lateral excitation is too limited to make ϕ_r^b reliably observable.

This conclusion was supported by observing that the estimated uncertainty of ϕ_r^b consistently increased over time unless lateral velocity was present. The null-space

analysis identified ϕ_r^b as an unobservable mode. These findings are visualized in Figure 4.1 and 4.2.

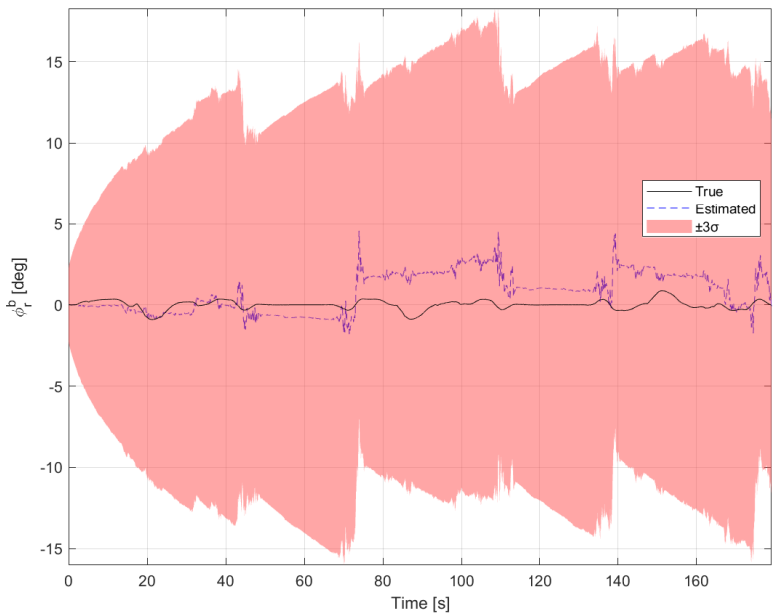

Figure 4.1: Estimated uncertainty of ϕ_r^b over time, showing divergence due to limited lateral velocity.

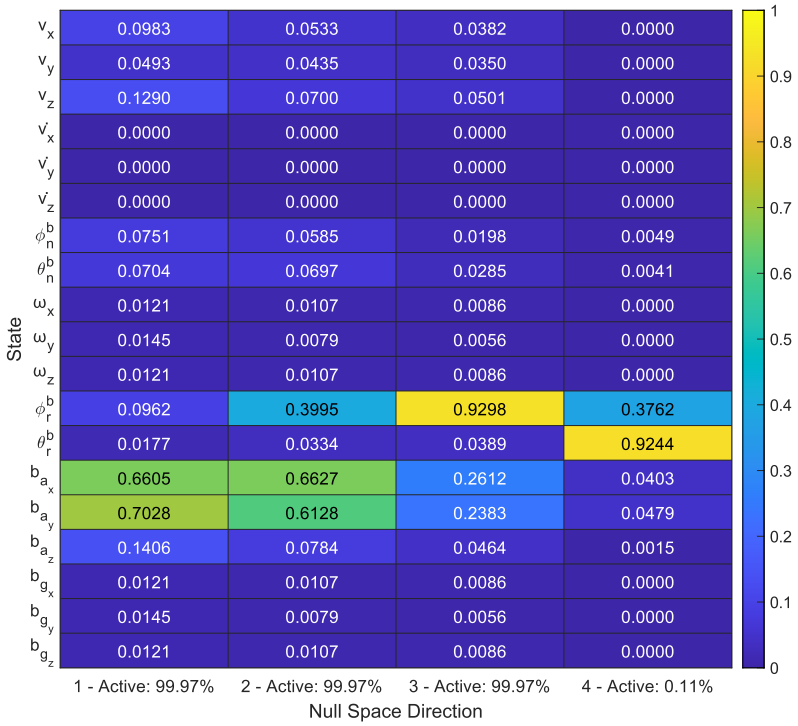

Figure 4.2: Accumulated null-space directions, highlighting that ϕ_r^b lies in an unobservable subspace.

As a result, ϕ_r^b was excluded from the state vector and instead modeled as a Gaussian random variable: $\phi_r^b \sim \mathcal{N}(0, 0.03)$.

4.3 Configuration 1: Unbiased IMU with Conventional Orientation

This configuration was intended to establish a baseline for system performance under idealized conditions. Specifically, beyond the conditions constant over all configurations, it assumed that

- No bias terms were included in the state vector.
- All sensor measurements were free from systematic bias.

The filter's behavior under these conditions was evaluated across five distinct driving scenarios. These tracks are described in more detail in Section 3.4.4. For each scenario, the results from a single realization will be shown. All estimated states as well as the road slope will be presented once and thereafter only the attitudes will be shown. The remaining figures can be found in Appendix A. In addition, for each scenario, the FoM obtained from a 100-run MC simulation are presented in a table.

4.3.1 Steady-State Cruise Track

The *Steady-State Cruise Track* served as an effective baseline for evaluating whether the filter diverged from the true state under near-static conditions. It functioned as the initial validation step to assess whether the filter was both well-structured and well-tuned. In this scenario, the variables that were directly measured, namely, the translational and angular velocities, were expected to be the easiest to estimate. These variables were directly supported by sensor inputs and were thus expected to converge rapidly and accurately.

Simulation results confirmed that the translational and angular velocities were estimated with high precision throughout the duration of the track. The attitude of the body frame with respect to the navigation frame converged to the true value quickly, indicating effective performance of the filter in processing implied variables. However, the estimation of the body frame attitude relative to the road frame showed higher uncertainty.

These finding are backed up by the result in the Figures 4.3 to 4.6, together with the FoM in Table 4.2.

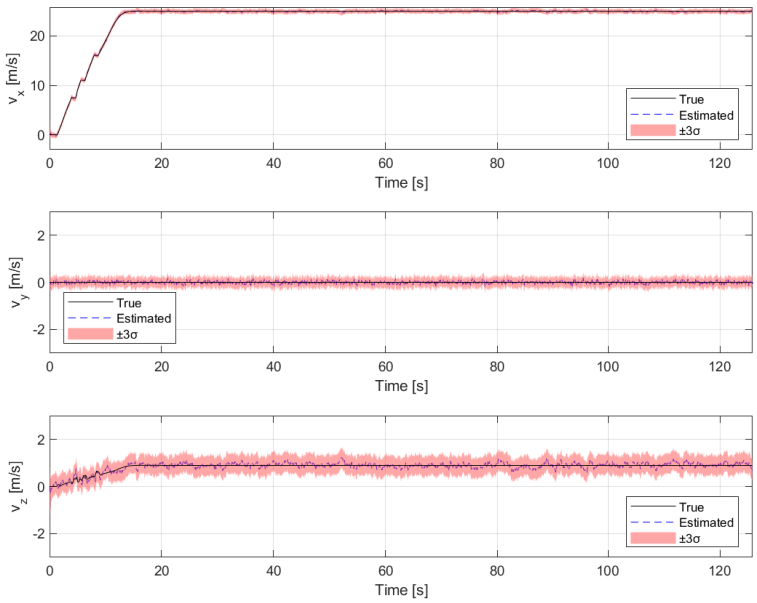

Figure 4.3: Configuration 1: Estimated translational velocities in the *Steady-State Cruise Track*.

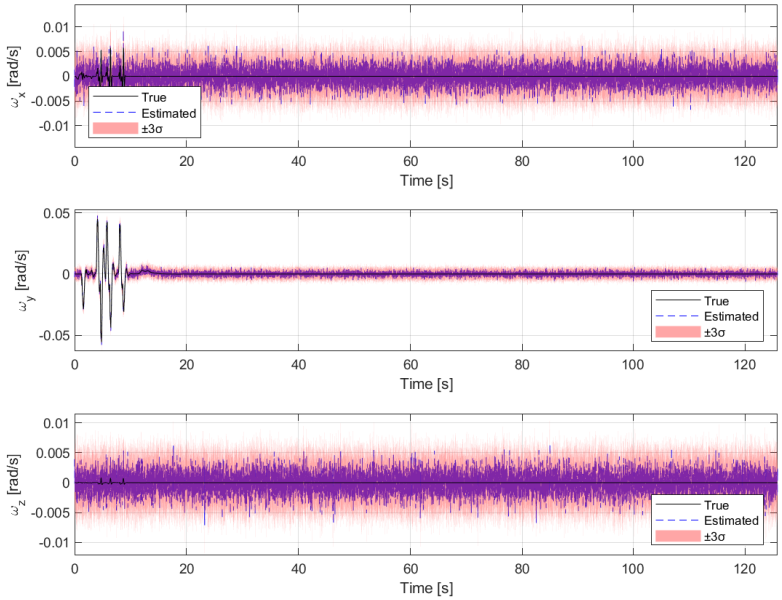

Figure 4.4: Configuration 1: Estimated angular velocities in the *Steady-State Cruise Track*.

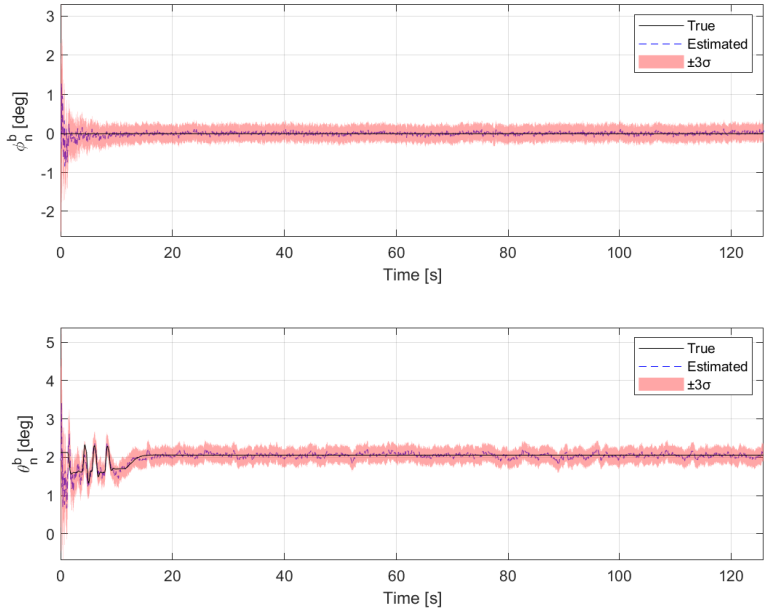

Figure 4.5: Configuration 1: Estimated body attitude w.r.t. navigation frame (ϕ_n^b, θ_n^b) in the *Steady-State Cruise Track*.

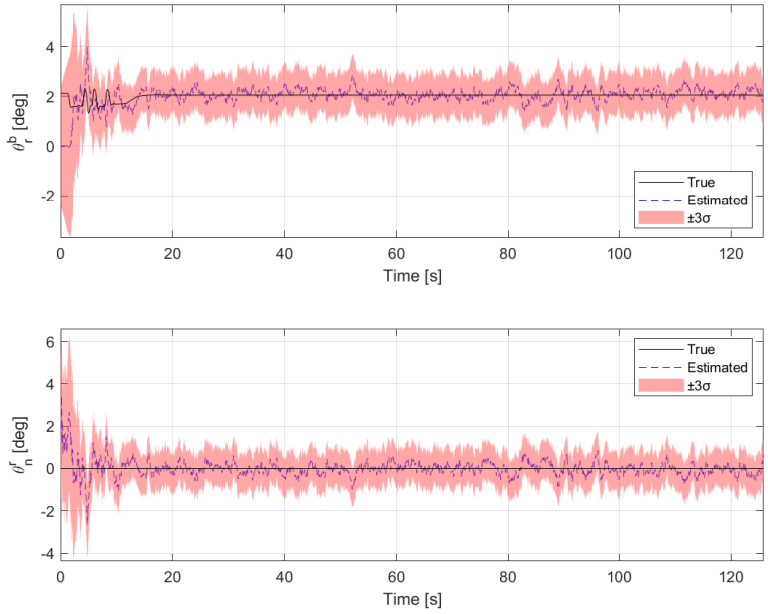

Figure 4.6: Configuration 1: Estimated body pitch w.r.t. road frame (θ_r^b) and estimated road slope w.r.t. navigation frame (θ_n^r) in the *Steady-State Cruise Track*.

Table 4.2: FoM for the *Steady-State Cruise Track* for configuration 1.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1083	0.0863	0.4205
	v_y	0.0549	0.0438	0.2033
	v_z	0.0993	0.0792	0.3857
Rotational Velocities [rad/s]	ω_x	0.0017	0.0014	0.0069
	ω_y	0.0017	0.0013	0.0068
	ω_z	0.0017	0.0013	0.0068
Attitude w.r.t. navigation frame [°]	ϕ_n^b	0.1160	0.0427	2.9773
	θ_n^b	0.1215	0.0643	1.8117
Body Attitude w.r.t. Road [°]	θ_r^b	0.3776	0.2255	2.3310
Road Attitude [°]	θ_n^r	0.4357	0.2725	3.5069

4.3.2 Varied Velocity Track

The second track, the *Varied Velocity Track*, was used to verify whether the filter maintained a similar level of estimation quality as observed in the *Steady-State Cruise Track*. The primary distinction between the two tracks was the longitudinal velocity profile, while the velocity remained constant in the cruise track, it varied in this scenario. This variation enabled a more nuanced evaluation of how the filter handled dynamic changes in the translational velocities in the absence of significant

angular motion.

The results from this track showed that both the translational and angular velocities were estimated accurately throughout the simulation, confirming the robustness of the filter under moderate longitudinal dynamics. The body frame's attitude with respect to the navigation frame also demonstrated reliable convergence to the true values, occurring within a short time window. The estimation of the body frame attitude relative to the road frame remained generally satisfactory, however a moderate level of uncertainty is still present.

These finding are backed up by the result in the Figures 4.7 and 4.8, together with the FoM in Table 4.3.

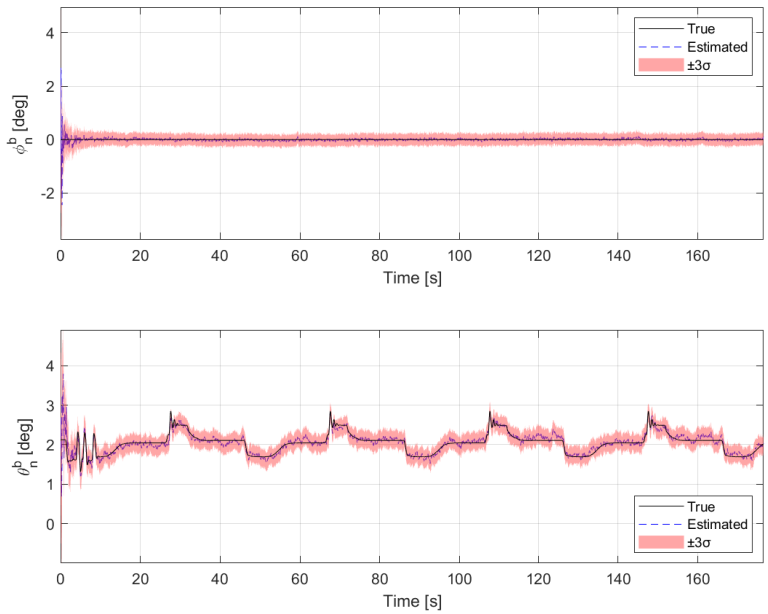

Figure 4.7: Configuration 1: Estimated body attitude w.r.t. navigation frame (ϕ_n^b, θ_n^b) in the *Varied Velocity Track*.

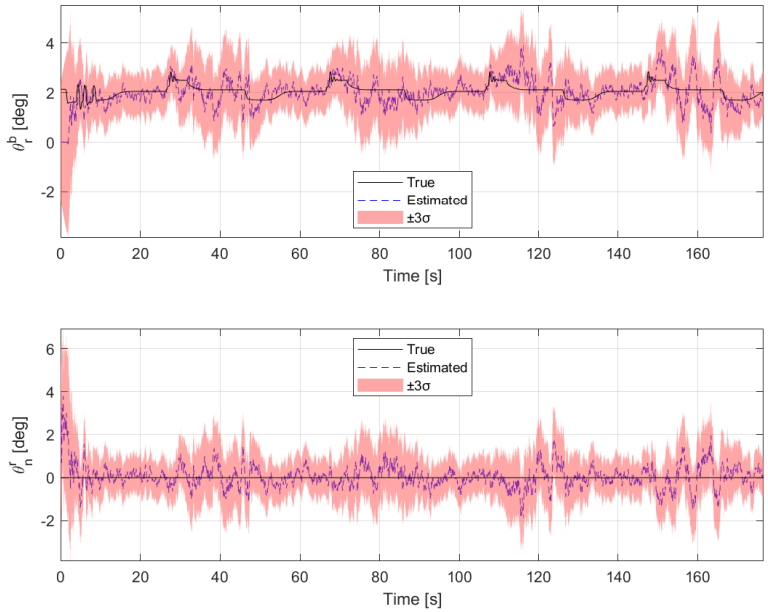

Figure 4.8: Configuration 1: Estimated body pitch w.r.t. road frame (θ_r^b) and estimated road slope w.r.t. navigation frame (θ_n^r) in the *Varied Velocity Track*.

Table 4.3: FoM for the *Varied Velocity Track* for configuration 1.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1101	0.0878	0.4391
	v_y	0.0552	0.0440	0.2077
	v_z	0.1002	0.0799	0.3998
Rotational Velocities [rad/s]	ω_x	0.0017	0.0014	0.0070
	ω_y	0.0017	0.0013	0.0070
	ω_z	0.0017	0.0013	0.0070
Attitude w.r.t. navigation frame [°]	ϕ_n^b	0.1003	0.0399	2.9354
	θ_n^b	0.1111	0.0610	1.8762
Body Attitude w.r.t. Road [°]	θ_r^b	0.4592	0.3167	2.3598
Road Attitude [°]	θ_n^r	0.5180	0.3614	3.6638

4.3.3 Uphill Track

The next natural alteration to the scenario was adding differing road characteristics. This lead to the *Uphill Track*. Accurate estimations in this case further solidified the observation that the measured variables were mostly accurately estimated. Good performance in the estimation of the various attitudes also showed that the filter performed well, despite greater attitude changes.

The simulation results for this track indicated that both the translational and angular velocities were accurately estimated throughout the run, maintaining consistency with previous scenarios. The body frame attitude with respect to the navigation frame also converged well to the true values, with minimal delay, even under changing slope conditions. Estimation of the body frame attitude relative to the road frame was similarly satisfactory as for the previous scenarios.

These finding are backed up by the result in the Figures 4.9 and 4.10, together with the FoM in Table 4.4.

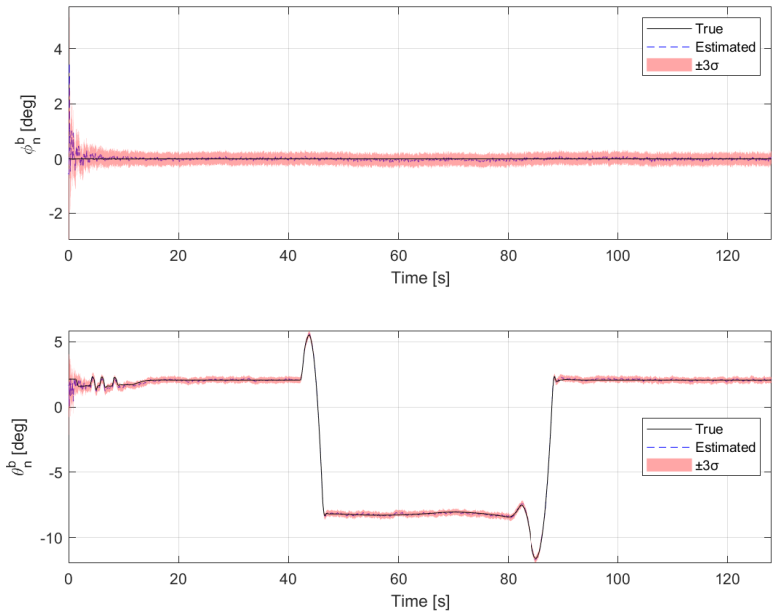

Figure 4.9: Configuration 1: Estimated body attitude w.r.t. navigation frame (ϕ_n^b, θ_n^b) in the *Uphill Track*.

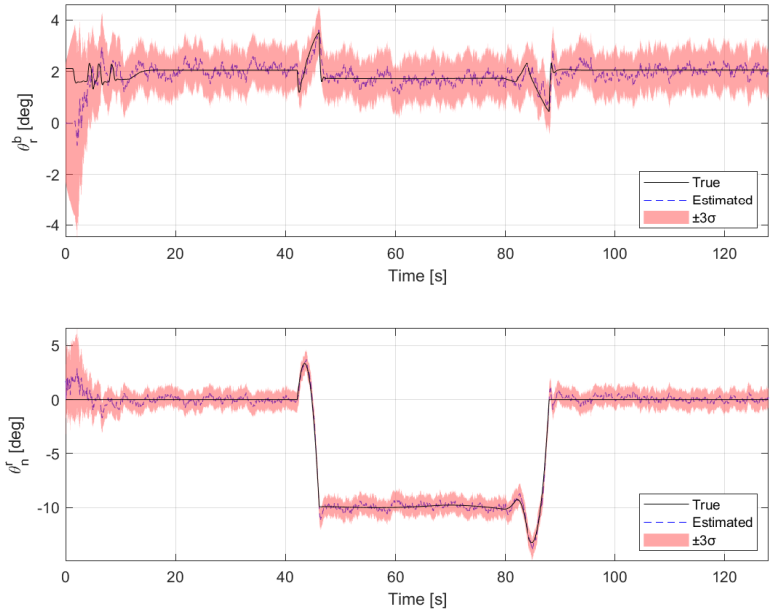

Figure 4.10: Configuration 1: Estimated body pitch w.r.t. road frame (θ_r^b) and estimated road slope w.r.t. navigation frame (θ_n^r) in the *Uphill Track*.

Table 4.4: FoM for the *Uphill Track* for configuration 1.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1097	0.0876	0.4236
	v_y	0.0551	0.0439	0.2029
	v_z	0.1005	0.0803	0.3896
Rotational Velocities [rad/s]	ω_x	0.0017	0.0014	0.0068
	ω_y	0.0017	0.0013	0.0068
	ω_z	0.0017	0.0013	0.0069
Attitude w.r.t. navigation frame [°]	ϕ_n^b	0.1129	0.0425	2.8504
	θ_n^b	0.1208	0.0662	1.7984
Body Attitude w.r.t. Road [°]	θ_r^b	0.3882	0.2389	2.3280
Road Attitude [°]	θ_n^r	0.4466	0.2876	3.4740

4.3.4 Dynamic J-Turn Track

The track *Dynamic J-Turn Track* was used to evaluate the filter performance when the roll angle of the truck was highly dynamic. The filter performed well in all but the body-to-road estimation. This showed that the filter performed well when mainly experiencing roll aligned dynamics.

Throughout the simulation, the translational and angular velocities were estimated

accurately, with low error and consistent convergence. The body attitude with respect to the navigation frame also showed rapid convergence to the true values, demonstrating that the filter could successfully adapt to dynamic roll-dominant motion. However, when evaluating the attitude of the body frame relative to the road frame, the filter performed slightly worse than in previous scenarios. This was due to the body pitch and roll with regards to the road frame being closely coupled, meaning that since the roll was assumed to be approximately zero, a deviation from this value caused less accurate pitch estimation. Due to a small roll deviation, this was not significant, yet had more impact than in previous scenarios.

These finding are backed up by the result in the Figures 4.11 and 4.12, together with the FoM in Table 4.5.

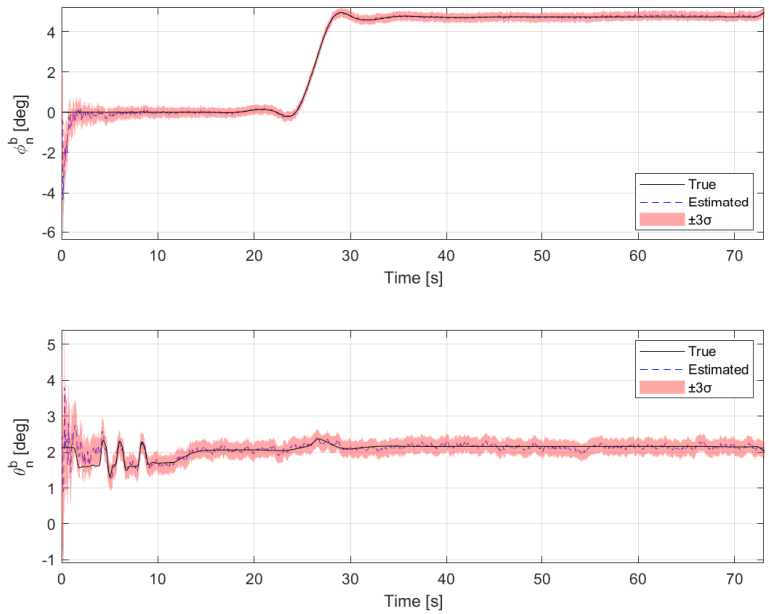

Figure 4.11: Configuration 1: Estimated body attitude w.r.t. navigation frame (ϕ_n^b, θ_n^b) in the *Dynamic J-Turn Track*.

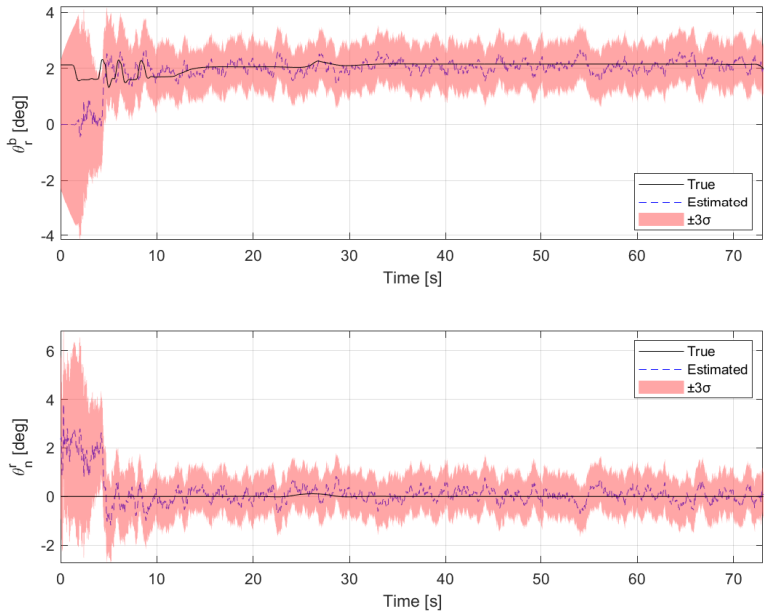

Figure 4.12: Configuration 1: Estimated body pitch w.r.t. road frame (θ_r^b) and estimated road slope w.r.t. navigation frame (θ_n^r) in the *Dynamic J-Turn Track*.

Table 4.5: FoM for the *Dynamic J-Turn Track* for configuration 1.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1111	0.0885	0.4200
	v_y	0.0550	0.0439	0.1975
	v_z	0.1017	0.0811	0.3858
Rotational Velocities [rad/s]	ω_x	0.0017	0.0014	0.0067
	ω_y	0.0017	0.0013	0.0066
	ω_z	0.0017	0.0013	0.0066
Attitude w.r.t. navigation frame [°]	ϕ_n^b	0.1430	0.0514	2.7610
	θ_n^b	0.1500	0.0759	1.7590
Body Attitude w.r.t. Road [°]	θ_r^b	0.4630	0.2696	2.3198
Road Attitude [°]	θ_n^r	0.5278	0.3201	3.4654

4.3.5 Infinity with Bank Track

The final track tested for this configuration was the *Infinity with Bank Track*. It allowed for the evaluation of the filter's ability to, despite a highly dynamic situation, maintain consistently accurate state estimates. The filter performed well in all but the body pitch with regard to the road frame estimation, and thus proved its adaptability and robustness in a wide range of dynamic conditions.

The simulation results confirmed that the translational and angular velocities were accurately estimated throughout the entire track, even under complex motion involving frequent changes in direction and road banking. The body frame attitude with respect to the navigation frame also converged quickly to the true values, demonstrating that the filter could effectively handle high-dynamic scenarios. The issues with the body-to-road pitch estimation seen in the *Dynamic J-Turn Track* were exaggerated in this scenario when the roll angle was significant, that is at the apex of the curves. The estimation of the pitch estimation was still reliable between these curves.

These finding are backed up by the result in the Figures 4.13 and 4.14, together with the FoM in Table 4.6.

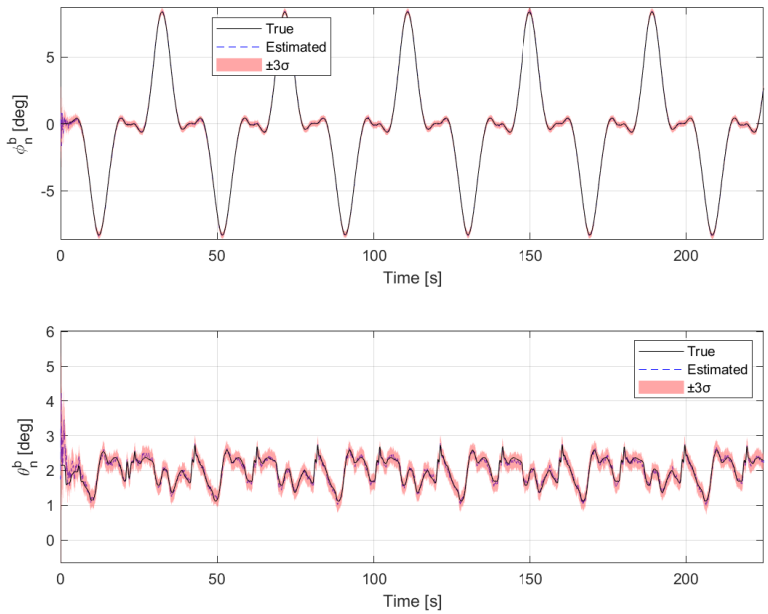

Figure 4.13: Configuration 1: Estimated body attitude w.r.t. navigation frame (ϕ_n^b, θ_n^b) in the *Infinity with Bank Track*.

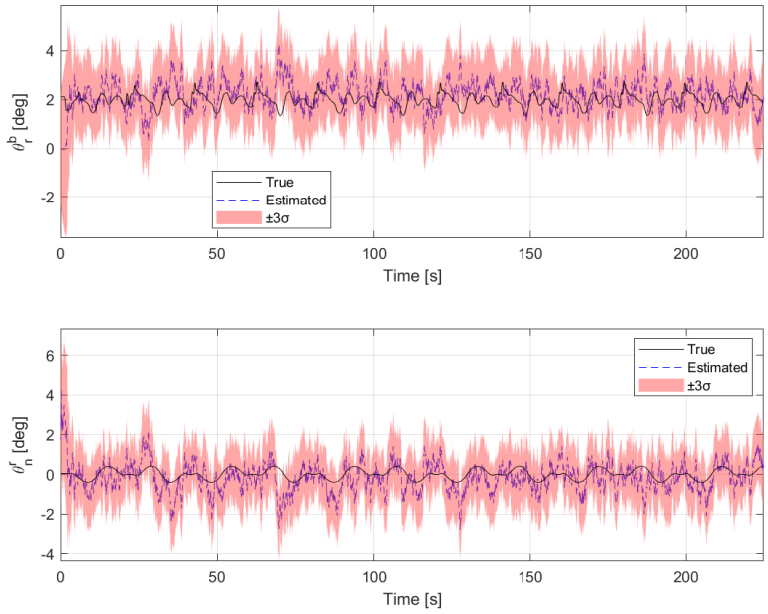

Figure 4.14: Configuration 1: Estimated body pitch w.r.t. road frame (θ_r^b) and estimated road slope w.r.t. navigation frame (θ_n^r) in the *Infinity with Bank Track*.

Table 4.6: FoM for the *Infinity with Bank Track* for configuration 1.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1086	0.0865	0.4486
	v_y	0.0551	0.0440	0.2143
	v_z	0.0995	0.0793	0.4078
Rotational Velocities [rad/s]	ω_x	0.0017	0.0014	0.0071
	ω_y	0.0017	0.0013	0.0070
	ω_z	0.0017	0.0013	0.0071
Attitude w.r.t. navigation frame [°]	ϕ_n^b	0.0998	0.5132	3.0482
	θ_n^b	0.1047	0.5500	1.9346
Body Attitude w.r.t. Road [°]	θ_r^b	0.6581	0.214 77	2.4651
Road Attitude [°]	θ_n^r	0.7082	0.218 61	3.6068

4.3.6 Discussion of the Observed Filter Behavior

The unbiased system analysis is an important first step when attempting to understand how well the base models and tuning works. If the tuning is sufficient and the models are correct, the variables ought to conform to the true values with some variation due to process and measurement noise. This gives an insight into which variables are observable, and ought thus be subject to an observability analysis.

In the case of the non-biased filter, these insights were clear. All the variables that depend directly on the IMU and the radars were observable and conform quickly with the respective true values. However, the body pitch with regard to the road frame, which is dependent on the pseudo measurement, as described in (3.8), behaved worse than the other variables. This is due to the coupling between the body to rad pitch and roll. Since it was established in Section 4.2 that the roll is unobservable, this variable was assumed to be a zero mean Gaussian random variable. As a consequence, when, in reality, there is roll, the pitch estimation will be poor. Nevertheless, the roll is often insignificant in nominal driving scenarios, allowing for this simplification without introducing major errors.

Beyond the body pitch with regard to the road frame, the other variables are also interesting to discuss. The measured variables are, as expected accurately estimated. These barely contain any noise, which has been filtered out, resulting in consistently low error with fast convergence, regardless of the track. Especially noteworthy is the constant and low FoM for the angular velocities. The reason for the FoM behavior is that the filter has more than sufficient high quality measurements to accurately estimate the true angular velocity values. This results in small error variation caused by the noise. This becomes apparent when viewing the FoM with more than four decimal places of precision, however, this level of accuracy is not necessary to view the overall performance of the filter, and therefore not shown. When additional disturbances, such as bias, is added, as in configuration 2 and 3, the differences in the FoM become more apparent.

The implied variables, that is the body attitude w.r.t. the navigation frame, are also accurate. These were more uncertain compared to the measured variables, however, after convergence, they consequently stayed very close to the true values. The post processed road slope angle, is dependent on the performance of the implied and pseudo measured variables. This means that the performance of the slope was coupled to the magnitude of the variables on which it was dependent. This is only an issue when there is a significant amount of body roll with regard to the road frame.

The complete non-biased system has thus unveiled the properties of the system that it was meant to. It showed that the process model and measurement model is correct and the tuning is sufficient. However, due to the uncertain nature of the pseudo measurement, and lack of more information about the variables dependent on this measurement, the performance of these are deteriorated or unusable.

4.4 Configuration 2: Biased IMU with Conventional Orientation

The section configuration was meant to exemplify the performance of the system given, beyond the conditions constant over all configurations, that

- Bias states were present in the state vector.

- Bias was present in the measurements from the IMU sensors.

This configuration's performance was determined by testing in four scenarios. These are described in greater detail in Section 3.4.4. Note that the *Steady-State Cruise Track* was not tested since its performance was nearly identical to that of the *Varied Velocity Track*. For each scenario, the results from a single realization will be shown. Since the performance of the velocities and the body attitude with regard to the road frame estimation were similar to those of configuration 1, only the results of the IMU bias estimation and the body-to-navigation frame angle estimation will be shown. The omitted results can be found in Appendix A. In addition, for each scenario, the FoM obtained from a 100-run MC simulation are presented in a table.

4.4.1 Varied Velocity Track

The *Varied Velocity Track* was interesting in the case where bias was present since it included some dynamic movement in terms of acceleration. It was, however, nearly static in other aspects. This made it a useful case for evaluating the filter's ability to estimate IMU biases when limited lateral and angular velocity was present.

This limited excitation caused the filter to not converge within the simulation window. While the accelerometer z-axis bias converged effectively, due to its alignment with gravity, the x and y axes remained largely unobservable under these low-dynamic conditions. As a result, the body frame attitude with respect to the navigation frame was poorly estimated. Although the estimated orientation followed the correct trend over time, the presence of an uncorrected offset rendered the absolute estimates inaccurate. In contrast, the body frame attitude relative to the road frame were estimated satisfactorily in the pitch direction.

These finding are backed up by the result in the Figures 4.15 and 4.16, together with the FoM in Table 4.7.

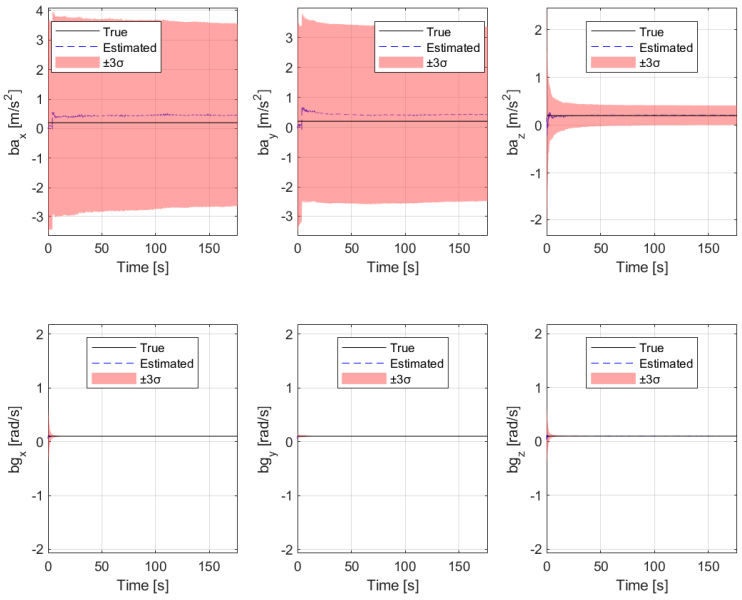

Figure 4.15: Configuration 2: Estimated IMU bias convergence for the *Varied Velocity Track*.

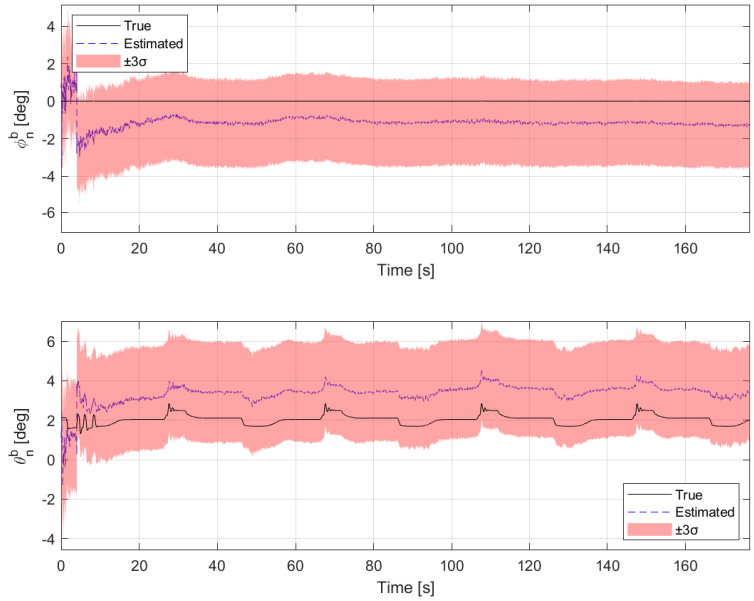

Figure 4.16: Configuration 2: Estimated body attitude w.r.t. navigation frame (ϕ_n^b, θ_n^b) for the *Varied Velocity Track*.

Table 4.7: FoM for the *Varied Velocity Track* for configuration 2.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1099	0.0877	0.4053
	v_y	0.0578	0.0453	0.3949
	v_z	0.1001	0.0797	0.3717
Rotational Velocities [rad/s]	ω_x	0.0029	0.0016	0.0605
	ω_y	0.0018	0.0014	0.0444
	ω_z	0.0033	0.0020	0.0550
Attitude w.r.t. navigation frame [°]	ϕ_n^b	2.7203	2.6837	4.8566
	θ_n^b	1.8441	1.8007	3.4586
Attitude w.r.t. Road [°]	θ_r^b	0.4668	0.3191	2.4186
Road Attitude [°]	θ_n^r	1.9350	1.8447	4.7220
Biases	b_{a_x}	0.3149	0.3088	0.4835
	b_{a_y}	0.4600	0.4549	0.6824
	b_{a_z}	0.0326	0.0249	0.3794
	b_{g_x}	0.0023	0.0004	0.0596
	b_{g_y}	0.0007	0.0002	0.0441
	b_{g_z}	0.0028	0.0013	0.0538

4.4.2 Uphill Track

The *Uphill Track* contributed valuable insights into how the filter reacted when an angular velocity around a single axis and how this affected its ability to estimate IMU biases. It served as a useful case between low-dynamic and highly dynamic scenarios.

In contrast to the straight *Varied Velocity Track*, the accelerometer biases on the x and z axes converged effectively during the simulation, owing to the influence of the slope and associated dynamic changes. This enabled more accurate estimation of orientation states that relied on these axes. However, the y-axis accelerometer bias remained largely unaffected and did not converge, which continued to limit the accuracy of variables dependent on this component. The translational and angular velocities were, as in other scenarios, estimated with high consistency and minimal error. The attitude estimation w.r.t. the navigation frame was significantly improved for the pitch component, particularly once the vehicle encountered the slope.

These finding are backed up by the result in the Figures 4.17 and 4.18, together with the FoM in Table 4.8.

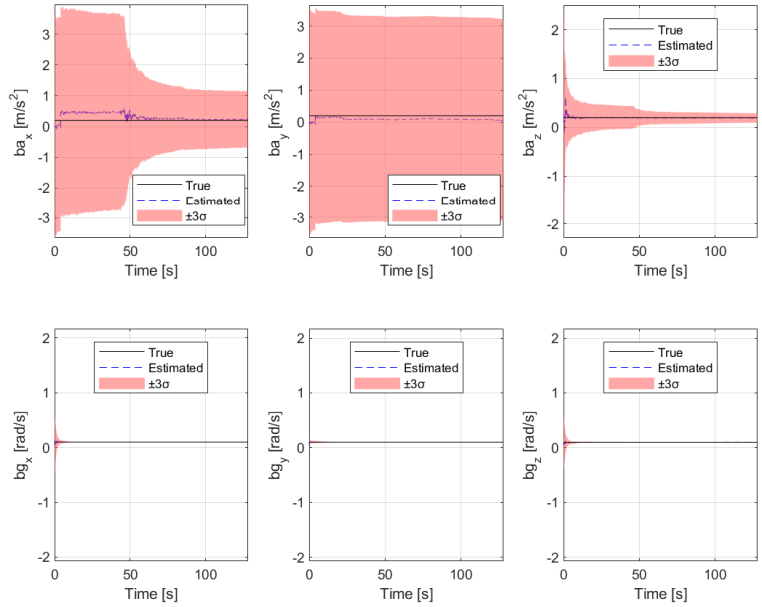

Figure 4.17: Configuration 2: Estimated IMU bias convergence for the *Uphill Track*.

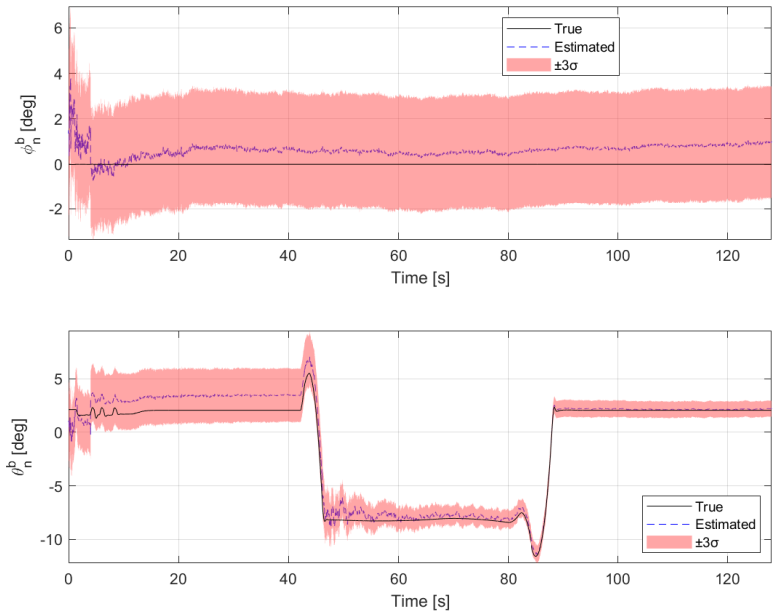

Figure 4.18: Configuration 2: Estimated body attitude w.r.t. navigation frame (ϕ_n^b, θ_n^b) for the *Uphill Track*.

Table 4.8: FoM for the *UphillTrack* for configuration 2.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1099	0.0877	0.3942
	v_y	0.0589	0.0459	0.3846
	v_z	0.1011	0.0806	0.3645
Rotational Velocities [rad/s]	ω_x	0.0032	0.0017	0.0585
	ω_y	0.0019	0.0014	0.0419
	ω_z	0.0037	0.0022	0.0539
Attitude w.r.t. navigation frame [°]	ϕ_n^b	2.9866	2.9308	5.4191
	θ_n^b	1.1565	0.8846	3.7317
Attitude w.r.t. Road [°]	θ_r^b	0.3973	0.2418	2.3798
Road Attitude [°]	θ_n^r	1.2650	0.9614	4.7021
Biases	b_{a_x}	0.1962	0.1528	0.5237
	b_{a_y}	0.5015	0.4935	0.7622
	b_{a_z}	0.0310	0.0199	0.3902
	b_{g_x}	0.0027	0.0005	0.0580
	b_{g_y}	0.0008	0.0003	0.0416
	b_{g_z}	0.0033	0.0016	0.0530

4.4.3 Dynamic J-Turn Track

The *Dynamic J-Turn Track* was an excellent contributor to the analysis of how the filter performed, since it introduced a significant change in the roll angles of the vehicle. It included the angular velocities not present in the *Uphill Track* while minimizing the pitch rate, which was prominent in the *Uphill Track*. This configuration allowed for a focused examination of the system's response to roll-inducing dynamics. It was also slightly more dynamic overall, providing additional insight into how the filter handled higher levels of motion.

In this scenario, the accelerometer biases for the x and y axes began to converge successfully, but only after the vehicle initiated the turn. This delayed convergence is evident in Figure 4.19, and it resulted in poor estimation of both attitude and bias-related variables during the initial straight-line segment. Once the turning maneuver began and the required dynamics were present, the filter was able to correct for the biases more effectively. At this point, the body frame attitude w.r.t. the navigation frame converged to the true values.

These finding are backed up by the result in the Figures 4.19 and 4.20, together with the FoM in Table 4.9.

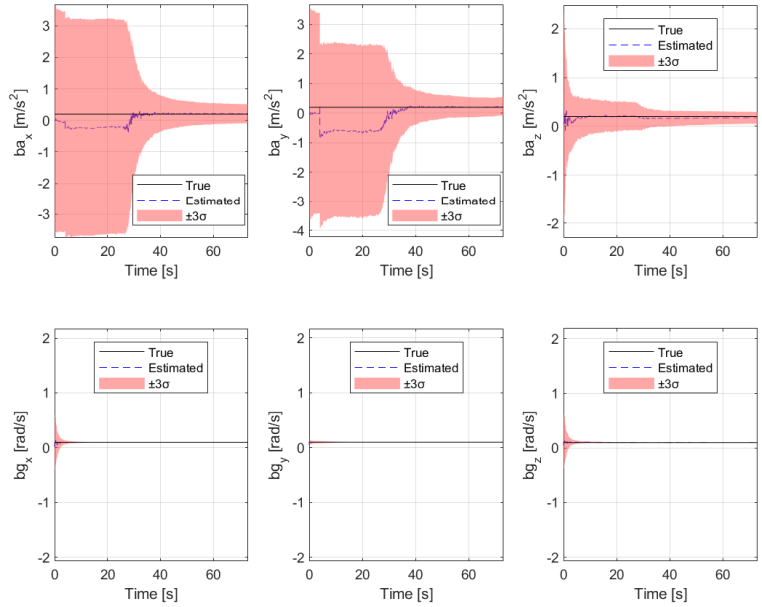

Figure 4.19: Configuration 2: Estimated IMU bias convergence for the *Dynamic J-Turn Track*.

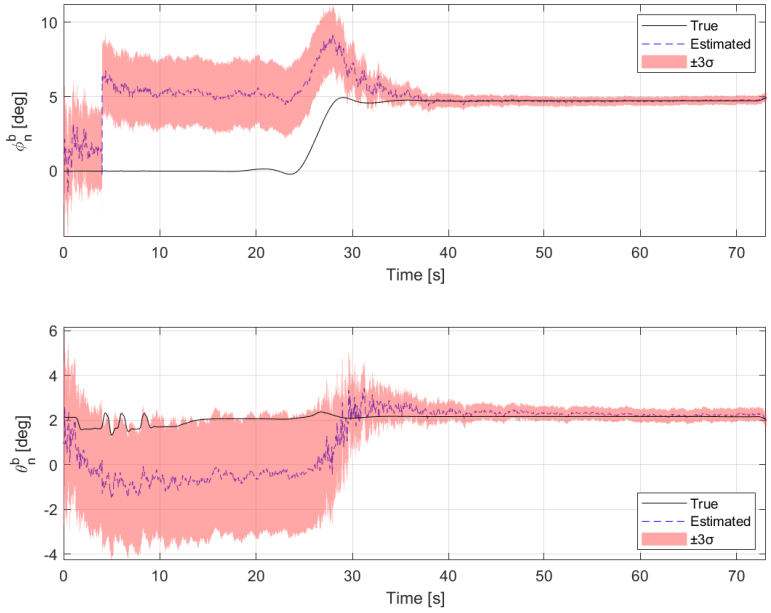

Figure 4.20: Configuration 2: Estimated body attitude w.r.t. navigation frame (ϕ_n^b, θ_n^b) for the *Dynamic J-Turn Track*.

Table 4.9: FoM for the *Dynamic J-Turn Track* for configuration 2.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1126	0.0897	0.3850
	v_y	0.0617	0.0474	0.3956
	v_z	0.1028	0.0819	0.3508
Rotational Velocities [rad/s]	ω_x	0.0040	0.0019	0.0627
	ω_y	0.0020	0.0015	0.0440
	ω_z	0.0045	0.0025	0.0553
Attitude w.r.t. navigation frame [°]	ϕ_n^b	1.6715	1.0961	4.6076
	θ_n^b	1.1800	0.7762	3.6313
Attitude w.r.t. Road [°]	θ_r^b	0.4811	0.2772	2.4217
Road Attitude [°]	θ_n^r	1.3392	0.9264	4.6840
Biases	b_{a_x}	0.1990	0.1337	0.5075
	b_{a_y}	0.2814	0.1933	0.6468
	b_{a_z}	0.0360	0.0237	0.3658
	b_{g_x}	0.0036	0.0008	0.0623
	b_{g_y}	0.0011	0.0004	0.0437
	b_{g_z}	0.0041	0.0019	0.0543

4.4.4 Infinity with Bank Track

The *Infinity with Bank Track* was a highly dynamic track. It was excellent for determining whether the filter would converge to the correct bias values in a short amount of time given enough dynamic movement. In this configuration, it allowed an investigation into whether extreme motion scenarios could potentially challenge or destabilize the filter.

The results in this scenario further confirmed the observations in the *Dynamic J-Turn Track*. The filter's ability to converge toward the correct bias values early in the simulation, even in the presence of noise, was still present and it performed better. Although noise introduces minor fluctuations around the bias estimates, these remain negligible. Due to this the overall filter performance was similar to that of the same scenario under zero-bias conditions, after bias convergence.

These finding are backed up by the result in the Figures 4.21 and 4.22, together with the FoM in Table 4.10.

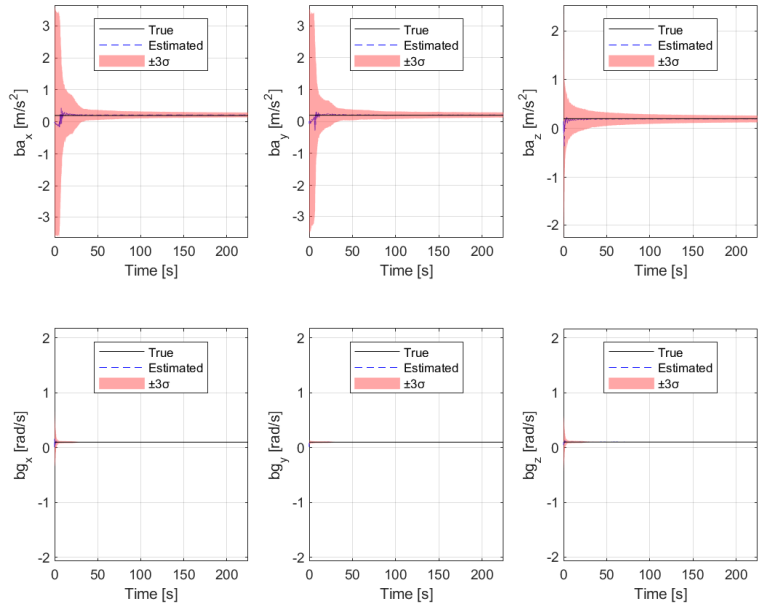

Figure 4.21: Configuration 2: Estimated IMU bias convergence over time for the *Infinity with Bank Track*.

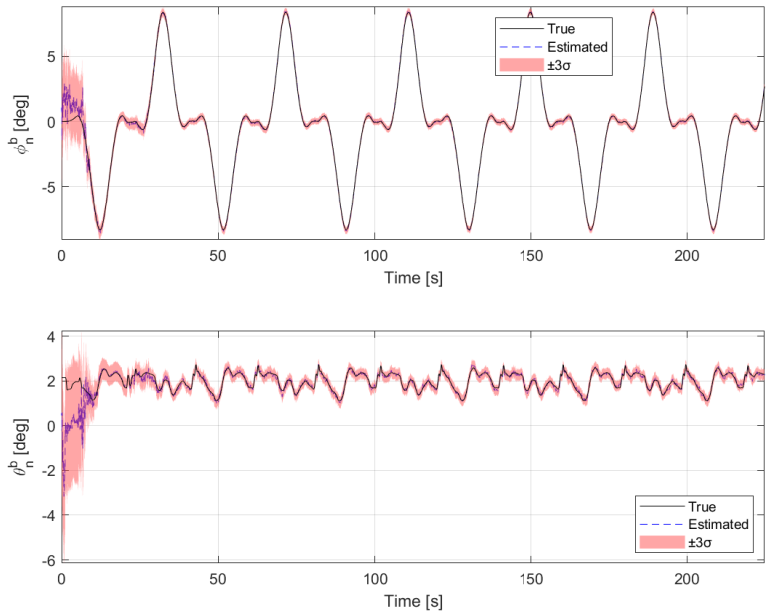

Figure 4.22: Configuration 2: Estimated body attitude w.r.t. the navigation frame (ϕ_n^b, θ_n^b) for the *Infinity with Bank Track*.

Table 4.10: FoM for the *Infinity with Bank Track* for configuration 2.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1112	0.0885	0.4274
	v_y	0.0576	0.0453	0.4015
	v_z	0.1022	0.0813	0.3888
Rotational Velocities [rad/s]	ω_x	0.0027	0.0015	0.0607
	ω_y	0.0019	0.0015	0.0448
	ω_z	0.0031	0.0019	0.0556
Attitude w.r.t. navigation frame [°]	ϕ_n^b	0.2392	0.0997	3.5325
	θ_n^b	0.2268	0.1012	3.0006
Attitude w.r.t. Road [°]	θ_r^b	0.6539	0.5080	2.4884
Road Attitude [°]	θ_n^r	0.7366	0.5609	4.0184
Biases	b_{a_x}	0.0363	0.0245	0.3428
	b_{a_y}	0.0352	0.0170	0.3871
	b_{a_z}	0.0186	0.0077	0.3870
	b_{g_x}	0.0021	0.0003	0.0603
	b_{g_y}	0.0008	0.0006	0.0446
	b_{g_z}	0.0026	0.0012	0.0551

4.4.5 Discussion of the Observed Filter Behavior

The second iteration of the system included IMU biases. This made the system more realistic since, if disturbances attributed to calibration errors are excluded, the biases and the noise is what remains. It is also prudent to note that the filter has been altered to include bias terms for the accelerometers and gyroscopes in the state vector to account for the bias disturbance. The issues introduced by the bias show up in the attitude measures that rely on the IMU.

When studying the results of the various simulations, it is clear that the introduction of IMU biases has no effect on the filter's ability to accurately estimate the translational and angular velocities. These were as accurate as for the filter where biases were not included. The estimation of the pitch angle w.r.t. the road was also unchanged, in its performance, since it was reliant on the estimated translational velocities, which are unaffected.

The difference between the biased and non-biased filters is seen in the estimation of the attitude of the body w.r.t. the navigation frame. These were directly coupled to the bias terms, via (3.31) and (3.32), and did not converge with the true values until biases did. The x and y axes accelerometer biases were passive until rotational dynamics were seen in the system while the z-axis bias converged quickly. This caused the estimations to be poor until enough rotational or gravity derived dynamics had been experienced. The z-axis converged since it is aligned with gravity, allowing the system to acquire an estimate of the bias. In the case of the x and y axes, the-

ses were always orthogonal or nearly orthogonal with gravity, making them reliant on the centripetal acceleration to become observable, see (3.31). The unobservable null-space directions, in figure 4.23, confirm this claim.

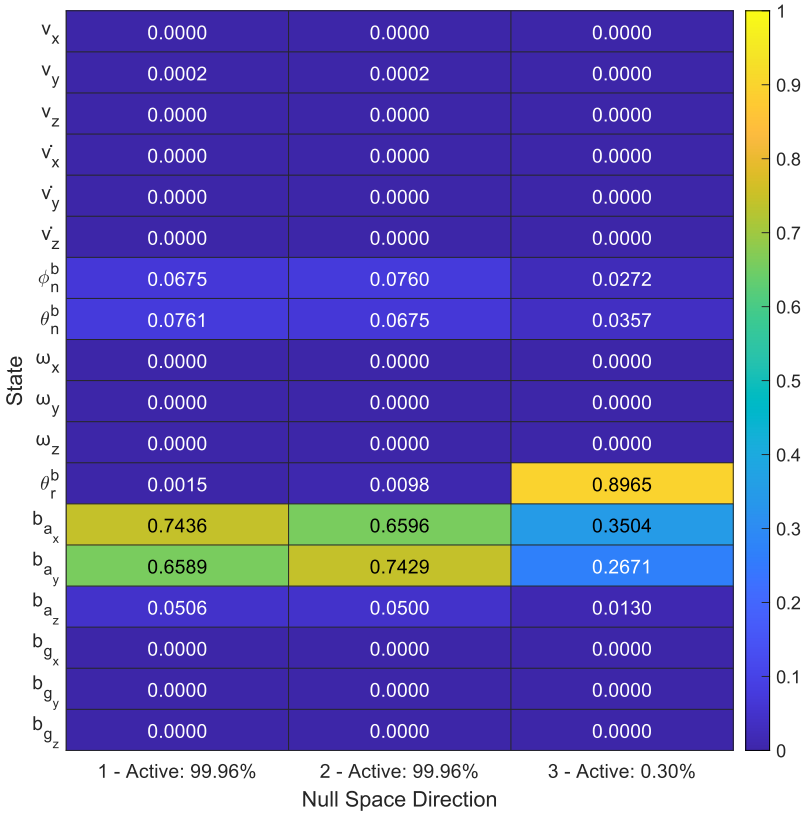

Figure 4.23: The accumulated null-space directions of the biased system.

The filter with bias included is a more realistic solution to real-world scenarios. It accounts for sensor imperfections beyond noise and jitter and can thus be used in less ideal systems. The bias estimation requires angular dynamics in order to make the biases on the x and y axes of the accelerometer observable. Before this has occurred the attitude estimations cannot be trusted however, over time, the filter will converge to the correct bias values and will thus be nearly as accurate as the unbiased version.

4.5 Configuration 3: Biased IMU with Tilted Orientation

Testing a configuration with a tilted IMU was motivated by the discovery that, in certain situations, reorienting the IMU, and thereby changing the system's unobservable direction, improved performance. It was useful when either pitch or roll was exclusively dynamic, before bias convergence. This will be shown through two

of the test tracks. One which is mostly static in its angular changes and one that has the pitch varying.

As stated earlier the results from utilizing this configuration were meant to exemplify the performance of the system given, beyond the conditions constant over all configurations, that:

- Bias states were present in the state vector.
- Bias was present in the measurements from any of the sensors.
- The IMU was oriented in a non-conventional manner where it was pitched up with 45° and then rolled to the left with 45°.

This configuration's performance was determined by testing two scenarios. These are described in greater detail in Subection 3.4.4. For each scenario, the results from a single realization will be shown. Since the performance of this configuration is nearly identical to that of configuration 2 in many aspects, only the results of the IMU bias estimation and the body-to-navigation frame angle estimation will be shown. The omitted results can be found in Appendix A. In addition, for each scenario, the FoM obtained from a 100-run MC simulation are presented in a table.

4.5.1 Varied Velocity Track

The *Varied Velocity Track* was used in this case to exemplify a situation where the IMU orientation did not improve the performance of the filter. As with the conventionally oriented IMU, this configuration failed to produce convergence of the bias variables, leading to inaccurate attitude estimation. Almost all the values that were well estimated in configuration 2 were still well estimated in this realization. This consistency demonstrated the robustness of the filter's velocity estimation performance regardless of IMU orientation. However, the inability of the biases to converge, particularly on all accelerometer axes, negatively impacted the accuracy of the body frame attitude with respect to the navigation frame. As a result, the attitude estimates were degraded. Furthermore, this tilted configuration performed slightly worse in this regard than the conventional IMU orientation, due to additional misalignment between the IMU axes and the gravitational direction.

These finding are backed up by the result in the Figures 4.24 and 4.25, together with the FoM in Table 4.11.

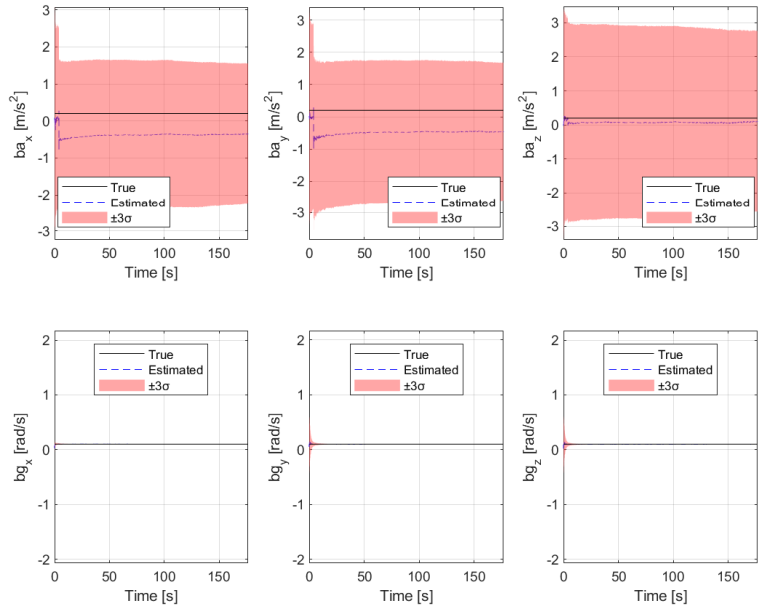

Figure 4.24: Configuration 3: Estimated IMU bias convergence over time for the *Varied Velocity Track*.

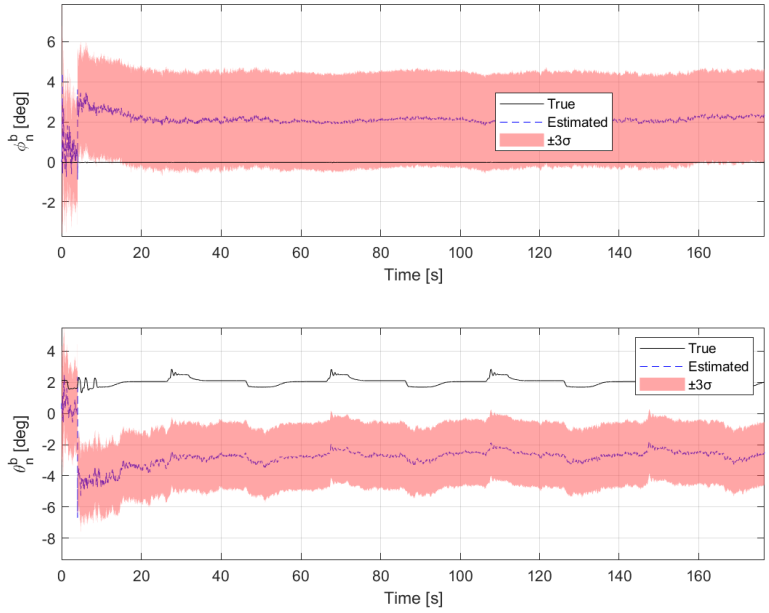

Figure 4.25: Configuration 3: Estimated body attitude w.r.t. the navigation frame (ϕ_n^b, θ_n^b) for the *Varied Velocity Track*.

Table 4.11: FoM for the *Varied Velocity Track* for configuration 3.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1104	0.0880	0.4046
	v_y	0.0580	0.0454	0.3988
	v_z	0.1006	0.0801	0.3674
Rotational Velocities [rad/s]	ω_x	0.0030	0.0016	0.0667
	ω_y	0.0018	0.0014	0.0316
	ω_z	0.0032	0.0019	0.0551
Attitude w.r.t. navigation frame [°]	ϕ_n^b	2.7183	2.6788	4.9328
	θ_n^b	1.9938	0.1012	3.7717
Attitude w.r.t. Road [°]	θ_r^b	0.4665	0.3200	2.3942
Road Attitude [°]	θ_n^r	2.0828	1.9939	4.776
Biases	b_{a_x}	0.2320	0.2280	0.4373
	b_{a_y}	0.3733	0.3684	0.5775
	b_{a_z}	0.3781	0.3729	0.5901
	b_{g_x}	0.0017	0.0009	0.0436
	b_{g_y}	0.0023	0.0007	0.0551
	b_{g_z}	0.0023	0.0007	0.0548

4.5.2 Uphill Track

The *Uphill Track* was used in this case to exemplify a situation where the IMU orientation did improve the performance of the filter. As with the conventionally oriented IMU, the bias variables converged during this scenario. However, in this tilted configuration, the accelerometer y-axis bias, which previously failed to converge, successfully started to converge. Although the convergence of the x and z axes biases was slower compared to the conventional setup, this trade-off enabled more uniform observability across all three axes. It resulted in quicker and more reliable bias convergence under certain motion conditions, without a significant loss in estimation performance.

These finding are backed up by the result in the Figures 4.26 and 4.27, together with the FoM in Table 4.12.

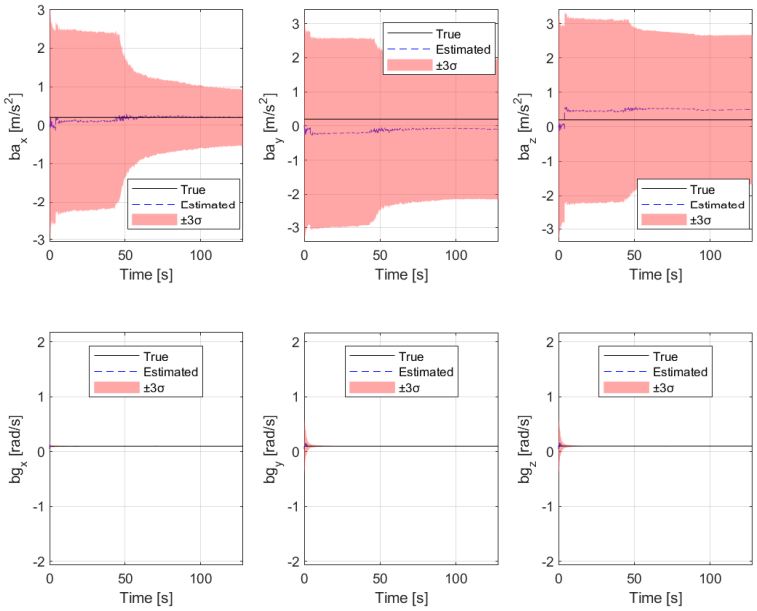

Figure 4.26: Configuration 3: Estimated IMU bias convergence over time for the *Uphill Track*.

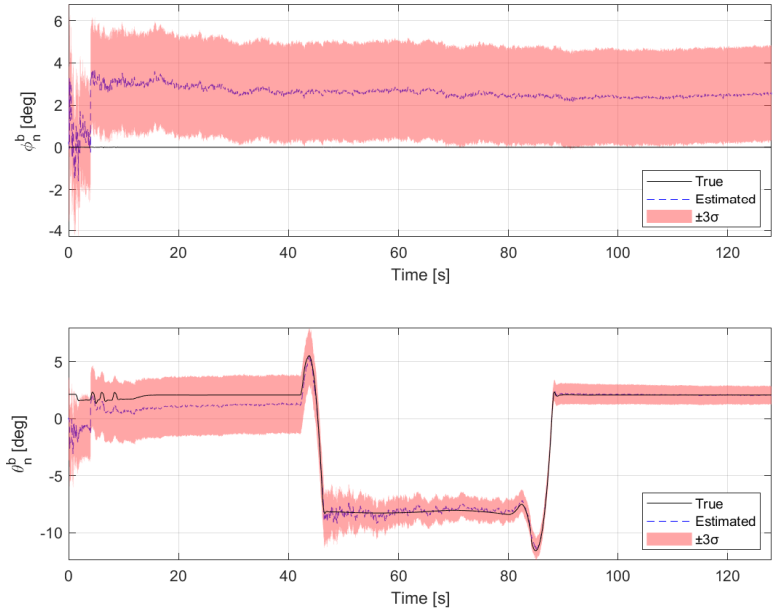

Figure 4.27: Configuration 3: Estimated body attitude w.r.t. the navigation frame (ϕ_n^b, θ_n^b) for the *Uphill Track*.

Table 4.12: FoM for the *Uphill Track* for configuration 3.

Category	Variable	RMSE	MAE	MaxAE
Translational Velocities [m/s]	v_x	0.1093	0.0870	0.3970
	v_y	0.0589	0.0458	0.3944
	v_z	0.1004	0.0799	0.3685
Rotational Velocities [rad/s]	ω_x	0.0033	0.0017	0.0696
	ω_y	0.0018	0.0014	0.0330
	ω_z	0.0035	0.0021	0.0535
Attitude w.r.t. navigation frame [°]	ϕ_n^b	2.8394	2.7886	4.8886
	θ_n^b	1.3915	0.1012	4.1859
Attitude w.r.t. Road [°]	θ_r^b	0.3996	0.2409	2.4252
Road Attitude [°]	θ_n^r	1.4907	1.1329	5.0989
Biases	b_{a_x}	0.1620	0.1262	0.4595
	b_{a_y}	0.3668	0.3505	0.6003
	b_{a_z}	0.3716	0.3556	0.6121
	b_{g_x}	0.0019	0.0011	0.0437
	b_{g_y}	0.0027	0.0009	0.0550
	b_{g_z}	0.0027	0.0008	0.0564

4.5.3 Discussion of the Observed Filter Behavior

The third and final configuration is very similar to the second. The reason for considering this configuration is to illustrate how a difference in orientation of the IMU sensor can change the filter's ability to converge to the true accelerometer bias values. The hypothesis that was set out to be proven was if the angular velocities and the gravity are dispersed onto the different axes, it would allow for slower but robust bias convergence. To show this, the IMU was reoriented, and then tested on two tracks, one that was not expected to converge with the new configuration nor the old and one that was expected to converge with the new configuration but not the old.

Since it is, in all other aspects identical to the filter with bias states, it is equivalently proficient at estimating the states not dependent on the IMU. However, the IMU dependent states, that is the body attitude w.r.t. the navigation frame and the IMU biases, act differently. In the case of non-convergence, the tilted IMU configuration performed, as expected, worse than the conventionally oriented IMU since, in this case, the z-axis accelerometer bias did not converge. This is expected since no angular velocities are present and the gravity is not aligned with any of the axes directly. When an angular velocity was introduced around one of the body frame axes this changes. In the case of the conventionally oriented IMU, the accelerometer bias on the y-axis does not converge since it is neither affected by gravity nor the centripetal acceleration. When the IMU was tilted, all three axes measured some kind of non-bias, non-translational acceleration. This caused all three axes' accelerometer bias

to gain observability. When the dynamics started, the bias estimation started converging. The cost of this was, however, slower convergence since the dynamics were spread over all three axes instead of being concentrated onto a single axis.

This phenomena was caused by the realigning of the unobservable null-space direction affecting the bias terms. The directions became less aligned with any one accelerometer bias term allowing for convergence during dynamic sections of the track. This spread is illustrated in Figure 4.28.

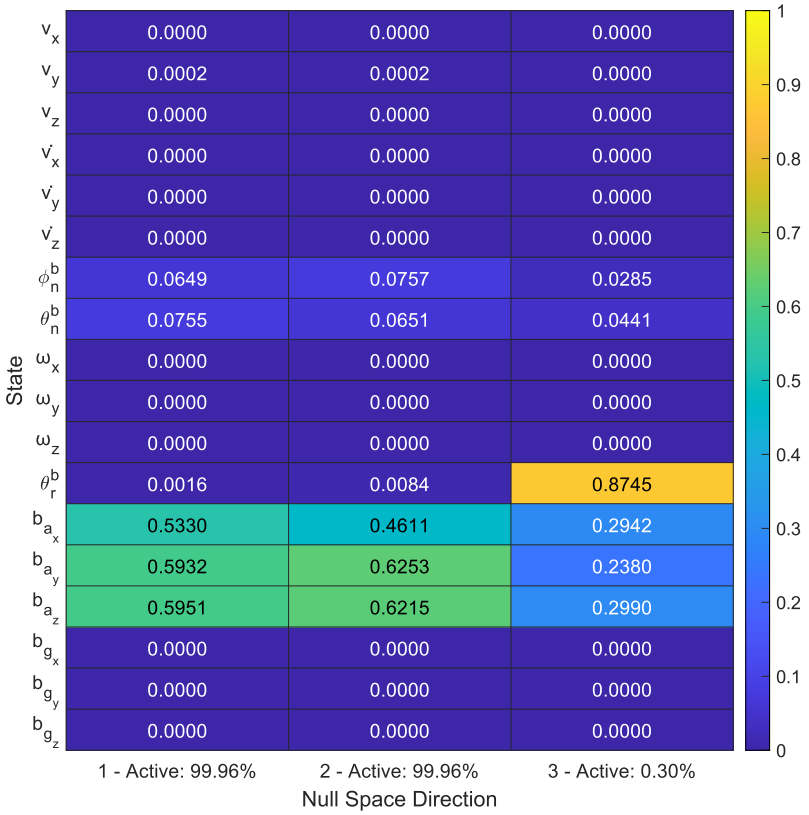

Figure 4.28: The accumulated null-space directions of the biased system with a tilted IMU.

The filter with bias states included and with a tilted IMU is a variant of the original filter that may have applications as a more robust version. It is slower to react to changes in the bias but will adapt, universally, in any situation when some dynamics are present. It is mostly applicable in slow systems where instantaneous convergence is not necessary. A calibration of the exact tilt may further improve the performance given which bias is the most important for the user. It is however an improvement upon the conventionally oriented IMU in long term performance and adaptability.

5
Conclusion

Accurate ego-motion and road attitude estimation is vital for heavy-duty vehicles. This thesis set out to address this issue by developing a system capable of accurately estimating these parameters. The proposed approach achieves an accurate result that accounts for sensor imperfections not caused by calibration errors.

The main contributions of this thesis lie in the fusion of radar and inertial sensors and how these achieve the desired results. The final filter utilizes a forked update step approach that allows for multirate sensors, that is sensors with different update frequencies. This enables seamless fusion of all sensors and produces an output that inherits the frequency of the sensor with the highest update rate. The filter was configured into three separate realizations, one without bias states, one with bias states and conventionally oriented sensors, and a third with bias states and a tilted IMU to address unobservable directions in the system.

The first configuration allowed for identifying flaws in the system. It revealed that using a pseudo measurement to estimate the body attitude with regard to the road frame was insufficient for determining the body roll with regard to the road frame, due to low levels of lateral velocity. As a result, this roll angle was modeled as a Gaussian random variable and excluded from the state vector. The second and third configurations provided insights into how well the system adapted to sensor biases. It was found that estimating these biases depended on the presence of either centripetal acceleration or a significant gravitational component. When both these were insufficient, accelerometer bias estimation was impossible. The main difference between the two biased configurations was the IMU orientation. While the reoriented IMU led to slower convergence across all dynamic scenarios, it eventually allowed for convergence of all three accelerometer biases. In contrast, the conventionally oriented IMU converged more quickly, but in some cases the y-axis bias failed to converge at all.

This thesis has therefore shown that ego-motion of a heavy-duty vehicle can be estimated using the proposed configuration. It has also demonstrated that body attitude with regard to the road frame estimation is possible on at least one axis, however, estimating both requires further research. It has also exemplified the advantages and drawbacks of reorienting the IMU sensor in order to improve bias estimation. However, there are limitations to this study. It only handles disturbances not related to calibration and does not process issues such as radar distortion and scale factor errors. Furthermore, the size of the filter may make it infeasible to di-

rectly incorporate into heavy-duty vehicles, and has therefore only been tested with simulation data.

As noted, there are several areas for future research that could improve the system and enhance understanding of ego-motion and road attitude estimation. Incorporating the range measurements from the radars may allow for accurate estimation of the road attitude. In [6], this was done without utilizing the Doppler velocities, which resulted in an estimate that used a large amount of radars and that has issues with road irregularities. A combination of that approach and the findings of this thesis may lead to a more robust solution. It is also beneficial to continue the investigation of how the filer can handle calibration derived disturbances since these are often unavoidable in reality. The modular design of the system allows for integration of additional sensors, which could improve overall accuracy and speed up bias convergence.

In conclusion, the results of this thesis are promising and indicate that the proposed system has the potential to support the development of robust and accurate estimators. Despite current limitations, the thesis provides a proof of concept that, with further development, could be implemented in real-world heavy-duty vehicles.

Bibliography

[1] A. W. Palmer and N. Nourani-Vatani, "Robust odometry using sensor consensus analysis," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 3167–3173.

[2] D. Kellner, M. Barjenbruch, J. Klappstein, J. Dickmann, and K. Dietmayer, "Instantaneous ego-motion estimation using multiple doppler radars," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 1592–1597.

[3] J. K. J. D. Dominik Kellner, Michael Barjenbruch and K. Dietmayer, "Instantaneous ego-motion estimation using doppler radar," *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*, vol. 16, october 2013. [Online]. Available: 10.1109/ITSC.2013.6728341

[4] A. Kramer, C. Stahoviak, A. Santamaria-Navarro, A.-a. Agha-mohammadi, and C. Heckman, "Radar-inertial ego-velocity estimation for visually degraded environments," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 5739–5746.

[5] Y. Z. Ng, B. Choi, R. Tan, and L. Heng, "Continuous-time radar-inertial odometry for automotive radars," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 323–330.

[6] L. Haglund and J. Runeby, "Ego-motion estimation using radars," Master's thesis, Chalmers University of Technology, Gothenburg, Sweden, 2024, master's Thesis in Systems, Control and Mechatronics. [Online]. Available: https://www.chalmers.se

[7] S. Särkkä and L. Svensson, *Bayesian Filtering and Smoothing*, second edition ed. Cambridge University Press, 2023. [Online]. Available: https://www.cambridge.org/9781108926645

[8] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2009.

[9] M. S. Grewal, L. R. Weill, and A. P. Andrews, *Global Positioning Systems, Inertial Navigation, and Integration*, 2nd ed. John Wiley & Sons, 2013.

[10] "Road vehicles — vehicle dynamics and road-holding ability — vocabulary," International Organization for Standardization, Geneva, CH, Standard, November 2011.

[11] T. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*, 04 2021.

[12] F. Maloberti, *Data Converters*. Springer, 2007.

[13] C. Jekeli, *Inertial Navigation Systems with Geodetic Applications*. Walter de Gruyter, 2001.

[14] R. Aljasami, "Fmwc-radar. signal processing and parameter estimation," Linköping, 2002.

[15] Y. Zhou, L. Liu, H. Zhao, M. López-Benítez, L. Yu, and Y. Yue, "Towards deep radar perception for autonomous driving: Datasets, methods, and challenges," *Sensors*, vol. 22, no. 11, p. 4208, 2022. [Online]. Available: https://www.mdpi.com/1424-8220/22/11/4208

[16] H. D. Young and R. A. Freedman, *Sears and Zemansky's University Physics with Modern Physics*, 15th ed. Pearson, 2019.

[17] H. Goldstein, C. P. Poole, and J. L. Safko, *Classical Mechanics*, 3rd ed. Addison Wesley, 2002.

[18] O. M. Baksalary and G. Trenkler, "The moore–penrose inverse: a hundred years on a frontline of physics research," *EPJ H*, vol. 46, no. 9, 2021.

[19] R. Hermann and A. J. Krener, "Nonlinear controllability and observability," *IEEE Transactions on Automatic Control*, vol. 22, no. 5, pp. 728–740, 1977.

[20] J. a. P. Hespanha, *Linear Systems Theory*. Princeton University Press, 2018.

[21] G. P. Huang, A. I. Mourikis, and S. I. Roumeliotis, "Observability-based rules for designing consistent ekf slam estimators," *The International Journal of Robotics Research*, vol. 29, no. 5, pp. 502–528, 2010.

[22] IPG Automotive, *CarMaker Reference Manual 13.0*, IPG Automotive, 2023, accessed from internal documentation or user manual.

[23] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software*. New York: John Wiley & Sons, 2001.

[24] M. Euston, P. Coote, R. Mahony, J. Kim, and T. Hamel, "A complementary filter for attitude estimation of a fixed-wing uav," in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2008, pp. 340–345.

[25] U. Frese, "A kalman filter for odometry using a wheel mounted inertial sensor," *Unpublished or ResearchGate article*, 2013, accessed: 2025-05-11. [Online]. Available: https://www.researchgate.net/publication/298663789

[26] M. Narasimhappa, A. D. Mahindrakar, V. C. Guizilini, M. H. Terra, and S. L. Sabat, "Mems-based imu drift minimization: Sage husa adaptive robust kalman filtering," *IEEE Sensors Journal*, vol. 20, no. 1, pp. 250–258, 2020.

[27] X. Wei, S. Fan, Y. Zhang, W. Gao, F. Shen, X. Ming, and J. Yang, "A robust adaptive error state kalman filter for mems imu attitude estimation under dynamic acceleration," *Measurement*, vol. 242, p. 116097, 2025.

[28] J. Bergman, H. Ohlsson, and N. Wahlström, "Asynchronous sensor fusion of gps, imu, and can-based odometry for heavy-duty vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 7, no. 1, pp. 120–131, 2022.

[29] M. Z. Naser and A. H. Alavi, "Error metrics and performance fitness indicators for artificial intelligence and machine learning in engineering and sciences," *Architecture, Structures and Construction*, vol. 3, no. 4, pp. 499–517, 2023. [Online]. Available: https://doi.org/10.1007/s44150-021-00015-8

A

Omitted Result Plots

A.1 Configuration 1: Unbiased IMU with Conventional Orientation

A.1.1 Varied Velocity Track

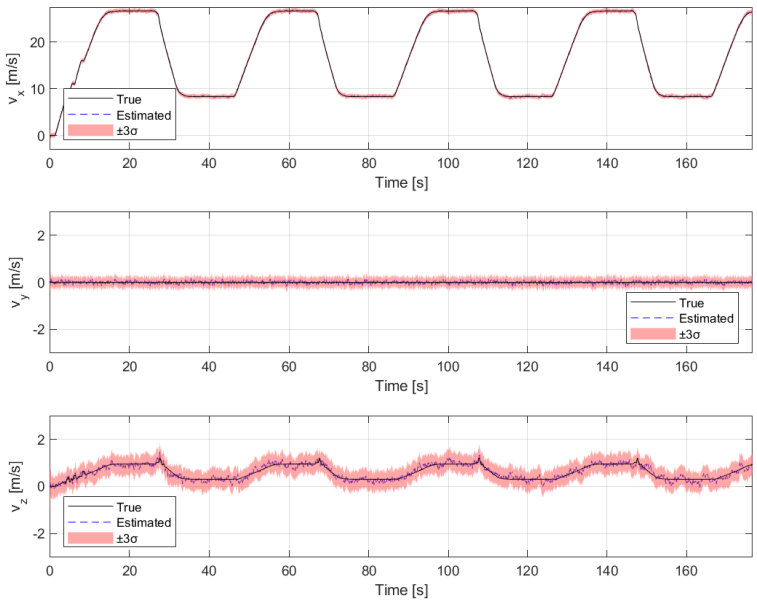

Figure A.1: Configuration 1: Estimated translational velocities in the *Varied Velocity Track*.

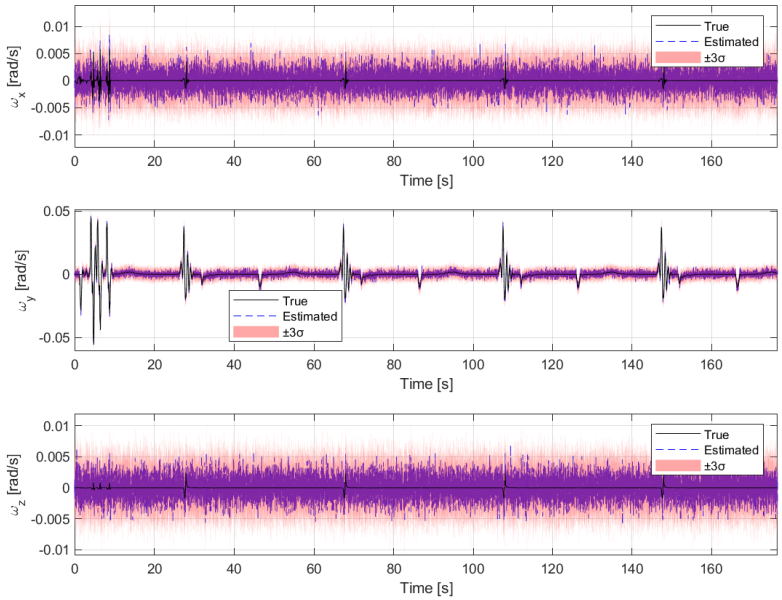

Figure A.2: Configuration 1: Estimated rotational velocities in the *Varied Velocity Track*.

A.1.2 Uphill Track

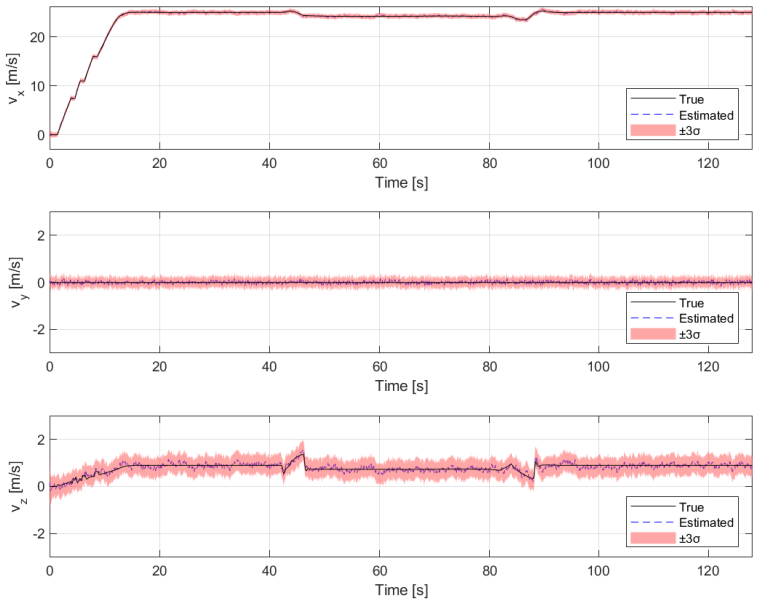

Figure A.3: Configuration 1: Estimated translational velocities in the *Uphill Track*.

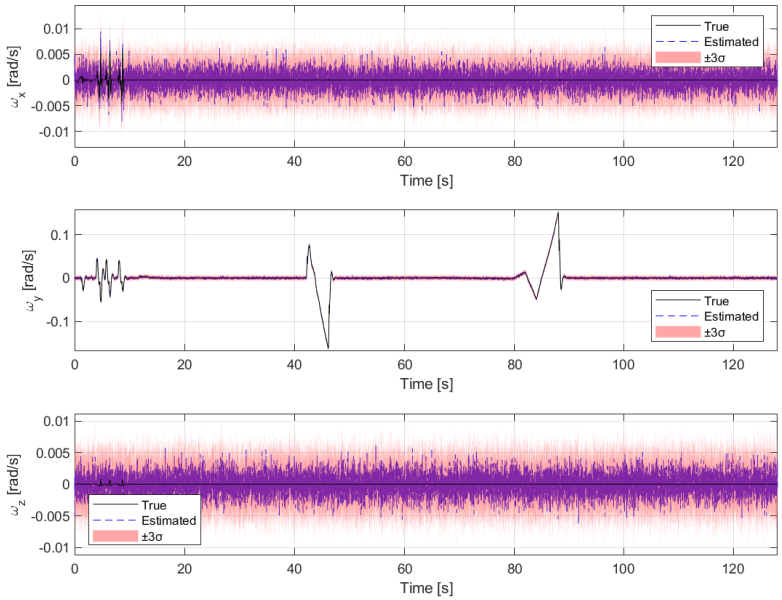

Figure A.4: Configuration 1: Estimated rotational velocities in the *Uphill Track*.

A.1.3 Dynamic J-Turn Track

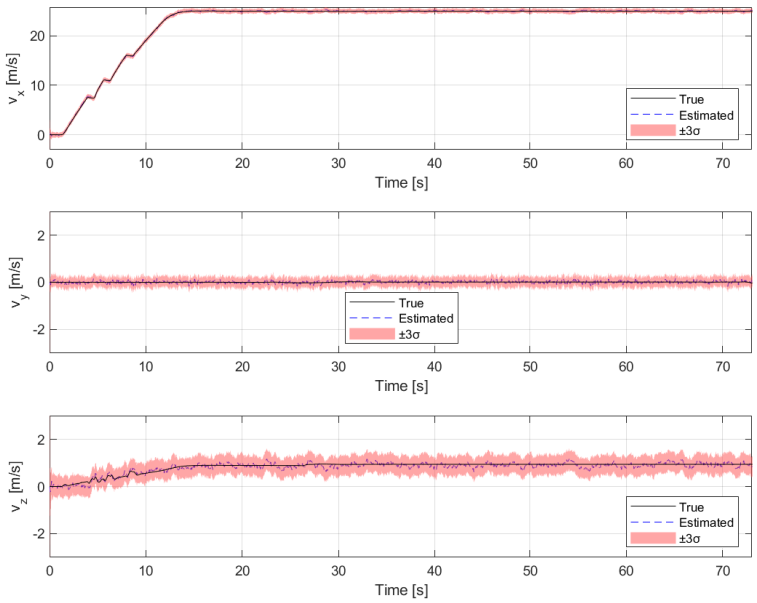

Figure A.5: Configuration 1: Estimated translational velocities in the *Dynamic J-Turn Track*.

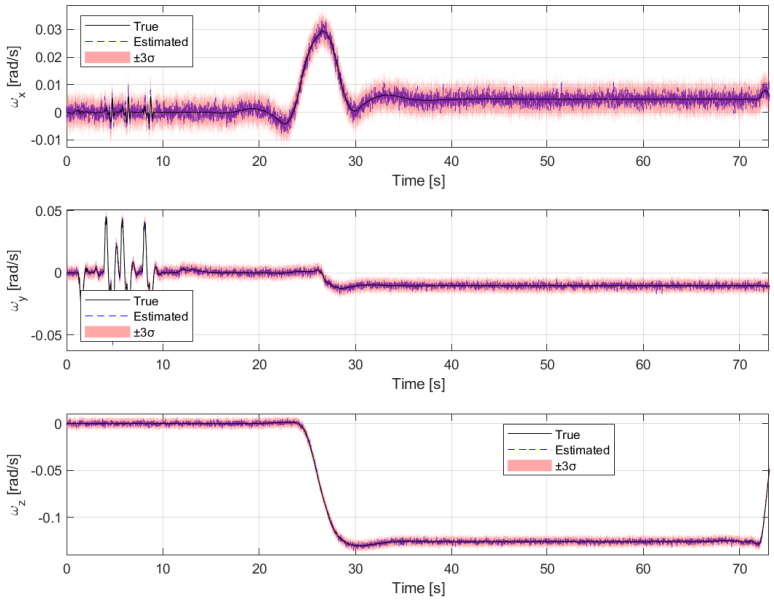

Figure A.6: Configuration 1: Estimated rotational velocities in the *Dynamic J-Turn Track*.

A.1.4 Infinity with Bank Track

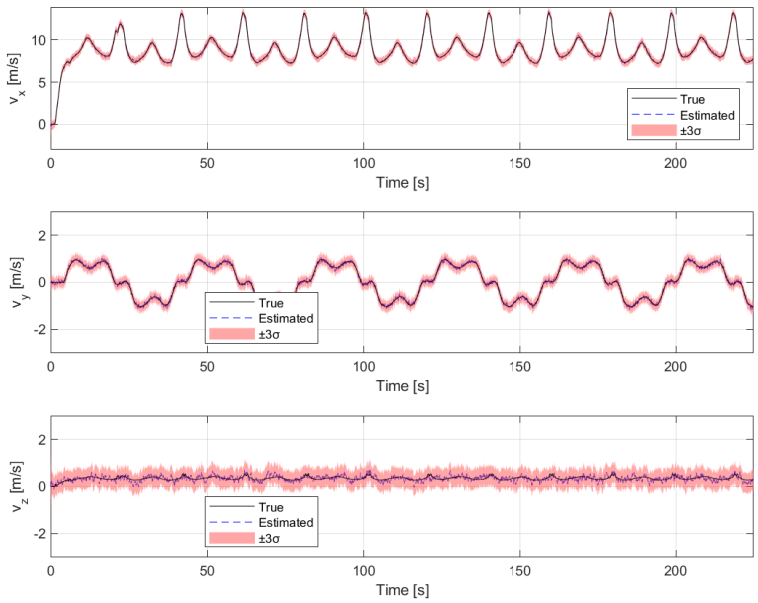

Figure A.7: Configuration 1: Estimated translational velocities in the *Infinity with Bank Track*.

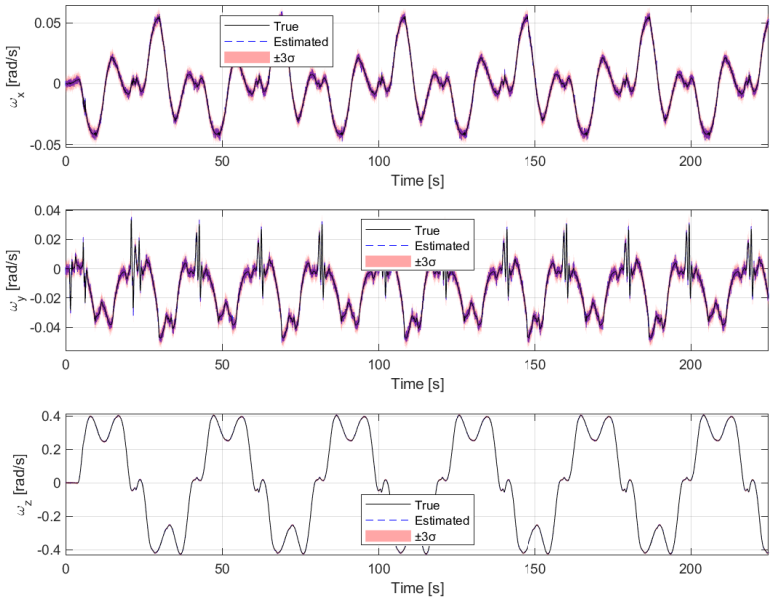

Figure A.8: Configuration 1: Estimated rotational velocities in the *Infinity with Bank Track*.

A.2 Configuration 2: Biased IMU with Conventional Orientation

A.2.1 Varied Velocity Track

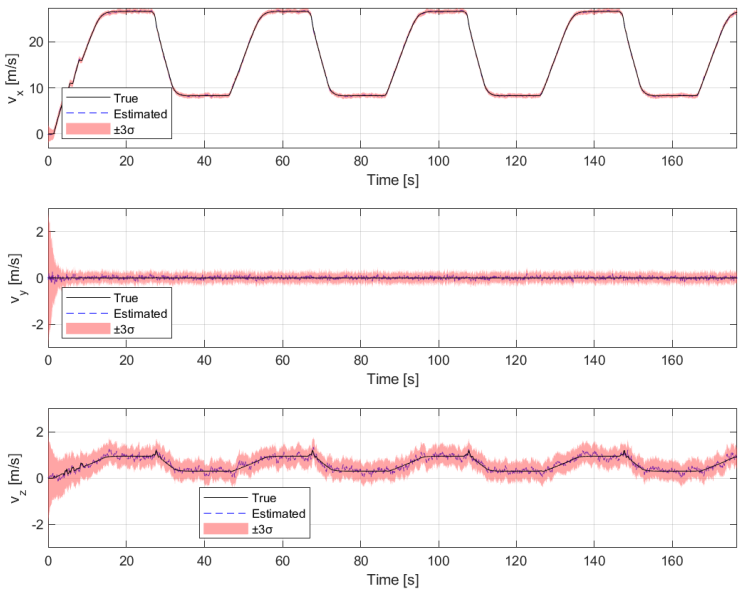

Figure A.9: Configuration 2: Estimated translational velocities in the *Varied Velocity Track*.

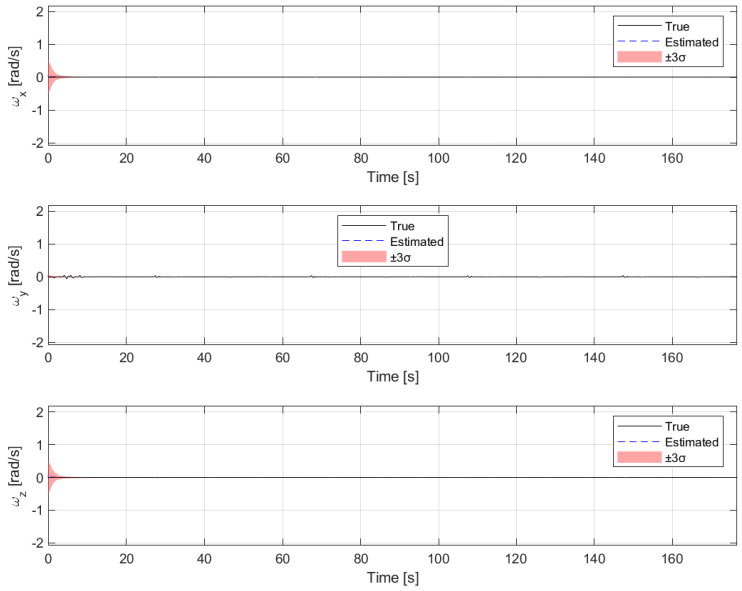

Figure A.10: Configuration 2: Estimated rotational velocities in the *Varied Velocity Track*.

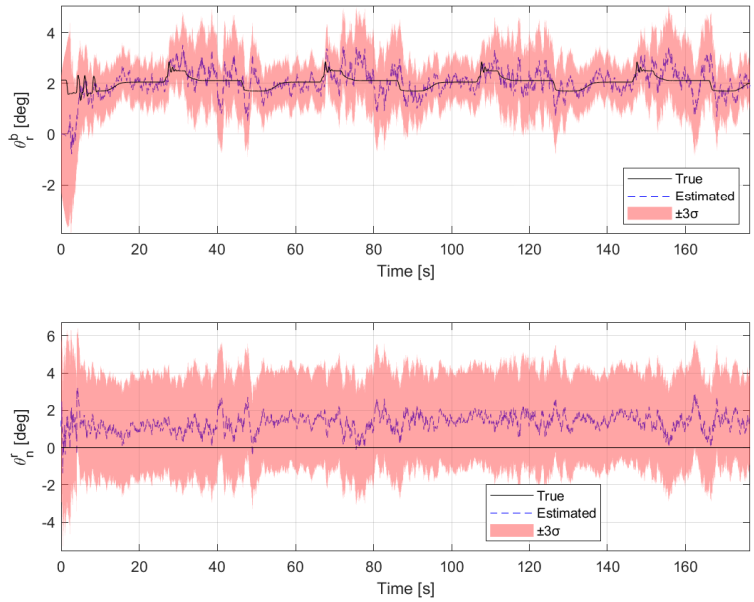

Figure A.11: Configuration 2: Estimated body pitch w.r.t. road frame (θ^b_r) and estimated road slope w.r.t. navigation frame (θ^r_n) in the *Varied Velocity Track*.

A.2.2 Uphill Track

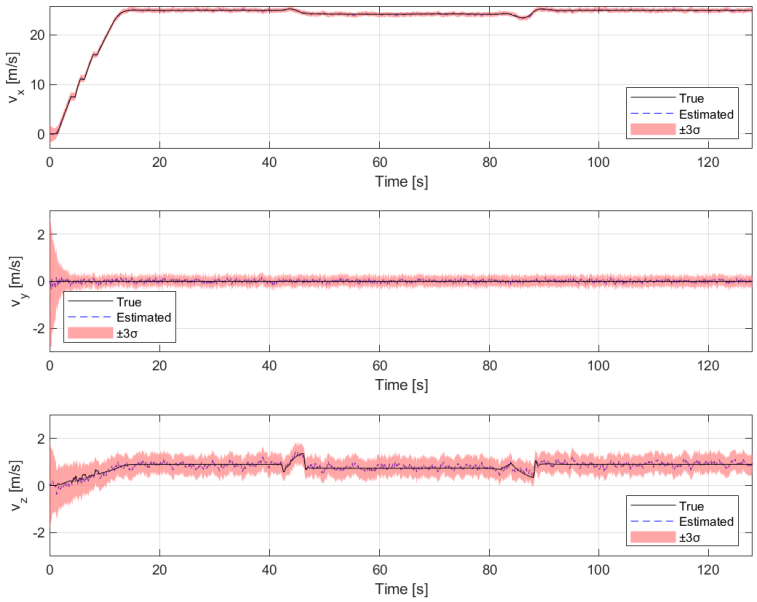

Figure A.12: Configuration 2: Estimated translational velocities in the *Uphill Track*.

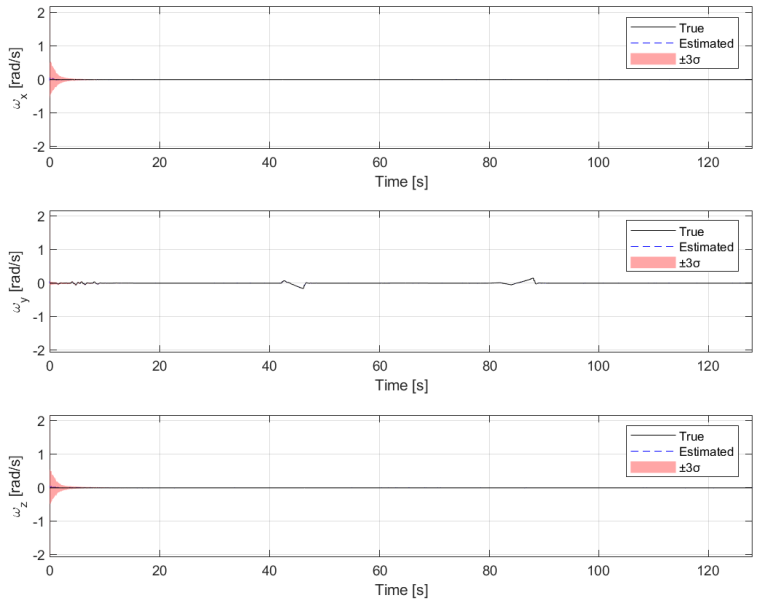

Figure A.13: Configuration 2: Estimated rotational velocities in the *Uphill Track*.

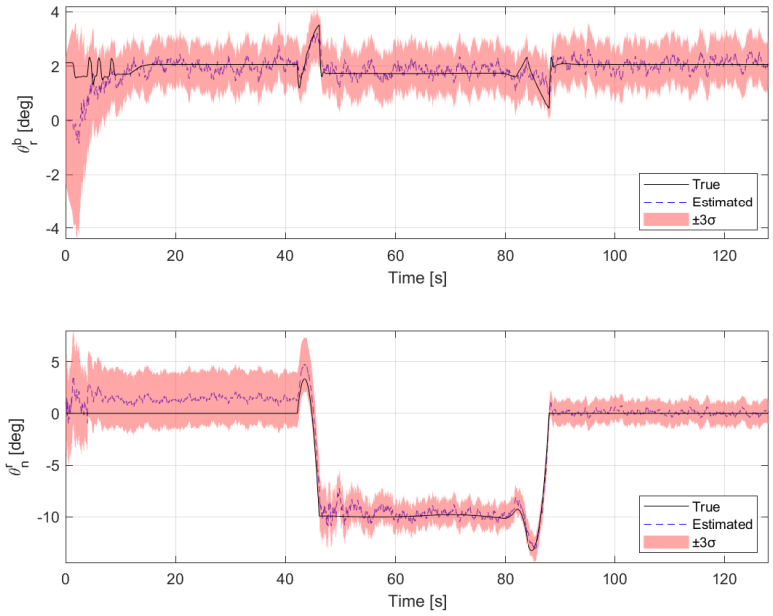

Figure A.14: Configuration 2: Estimated body pitch w.r.t. road frame (θ_r^b) and estimated road slope w.r.t. navigation frame (θ_n^r) in the *Uphill Track*.

A.2.3 Dynamic J-Turn Track

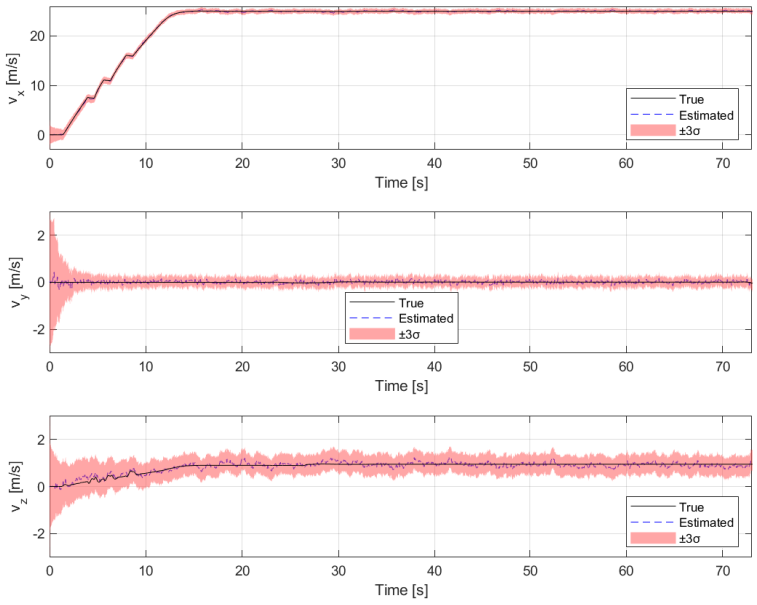

Figure A.15: Configuration 2: Estimated translational velocities in the *Dynamic J-Turn Track*.

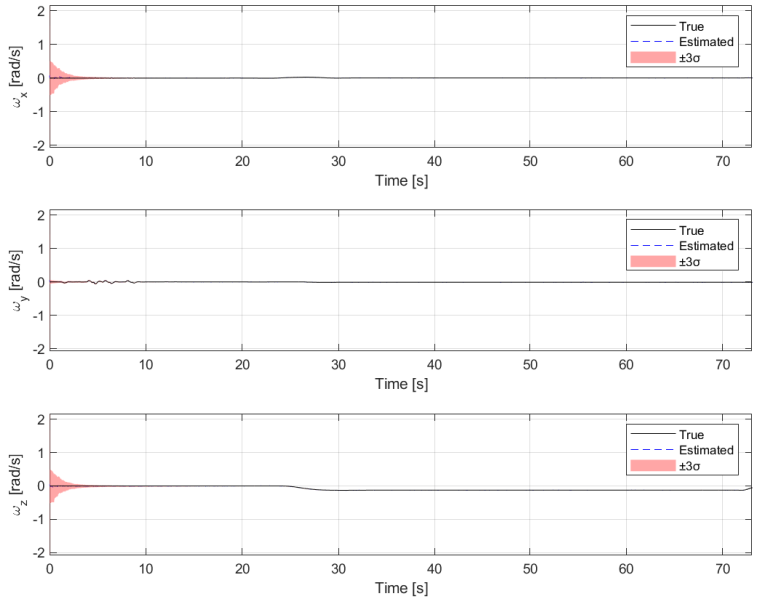

Figure A.16: Configuration 2: Estimated rotational velocities in the *Dynamic J-Turn Track*.

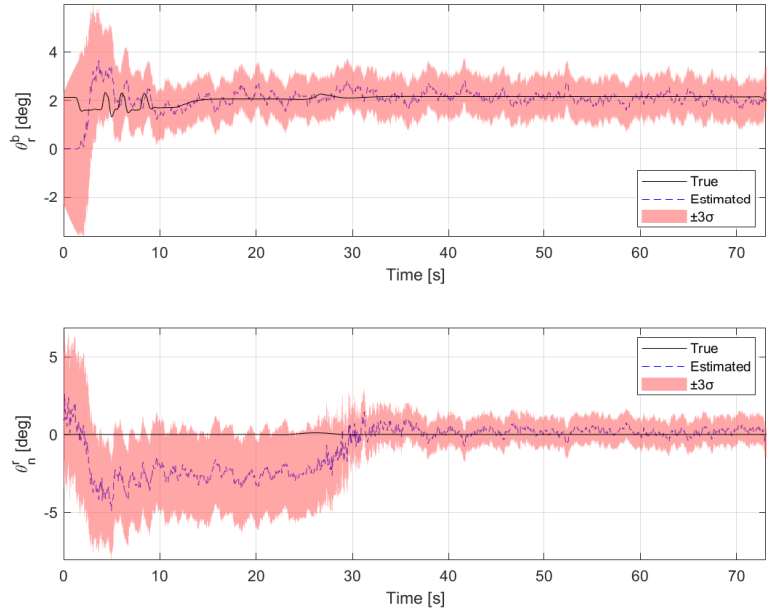

Figure A.17: Configuration 2: Estimated body pitch w.r.t. road frame (θ_r^b) and estimated road slope w.r.t. navigation frame (θ_n^r) in the *Dynamic J-Turn Track*.

A.2.4 Infinity with Bank Track

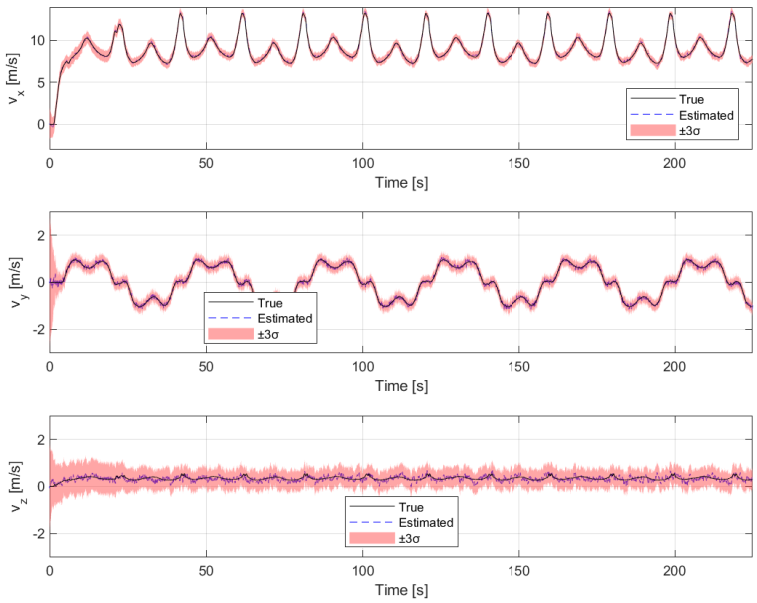

Figure A.18: Configuration 2: Estimated translational velocities in the *Infinity with Bank*.

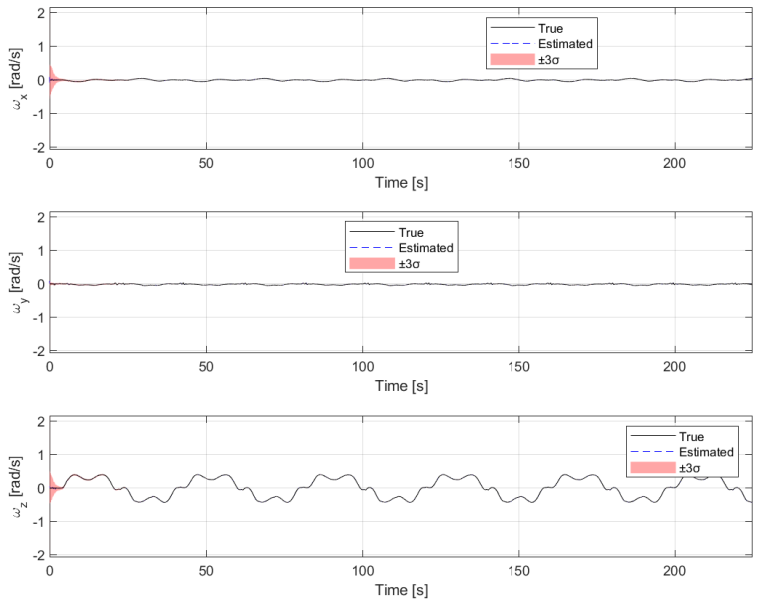

Figure A.19: Configuration 2: Estimated rotational velocities in the *Infinity with Bank*.

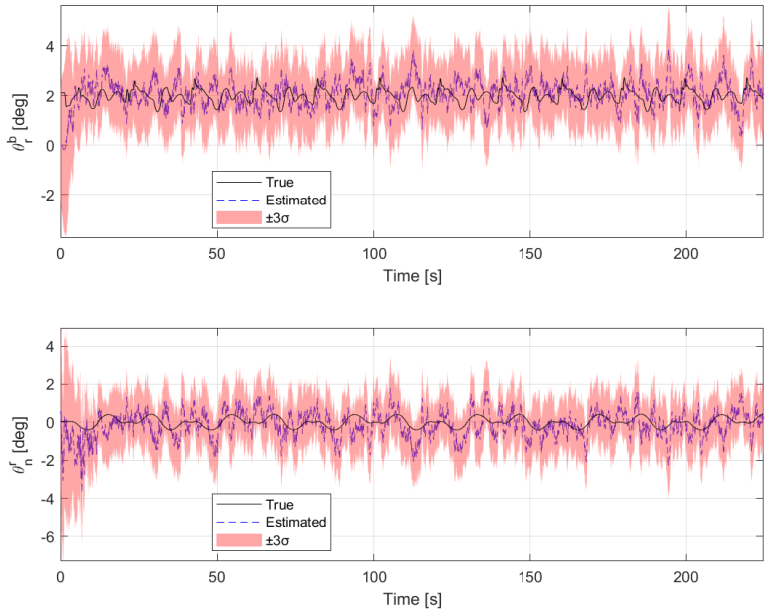

Figure A.20: Configuration 2: Estimated body pitch w.r.t. road frame (θ_r^b) and estimated road slope w.r.t. navigation frame (θ_n^r) in the *Infinity with Bank Track*.

A.3 Configuration 3: Biased IMU with Tilted Orientation

A.3.1 Varied Velocity Track

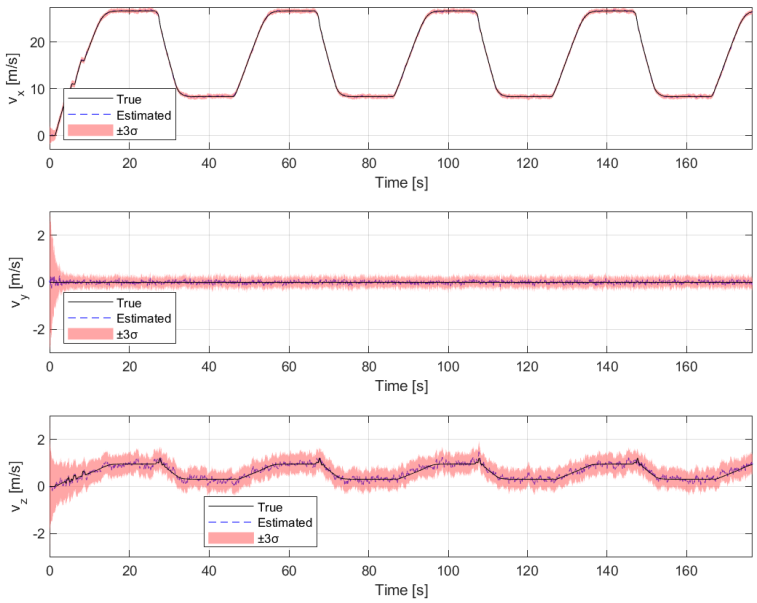

Figure A.21: Configuration 3: Estimated translational velocities in the *Varied Velocity Track*.

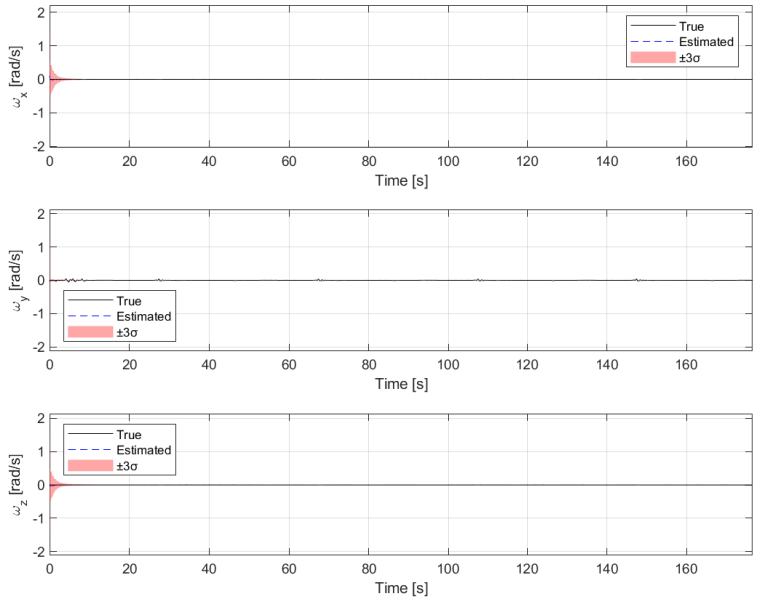

Figure A.22: Configuration 3: Estimated rotational velocities in the *Varied Velocity Track*.

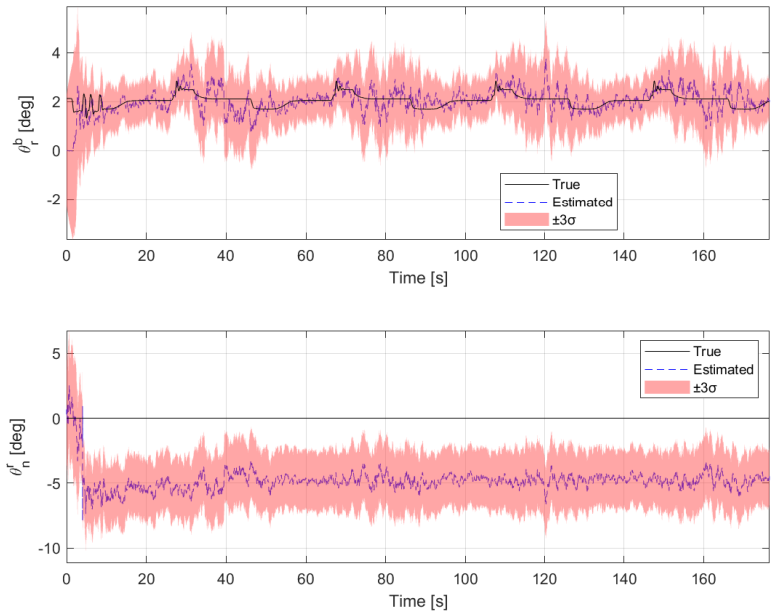

Figure A.23: Configuration 3: Estimated body pitch w.r.t. road frame (θ_r^b) and estimated road slope w.r.t. navigation frame (θ_n^r) in the *Varied Velocity Track*.

A.3.2 Uphill Track

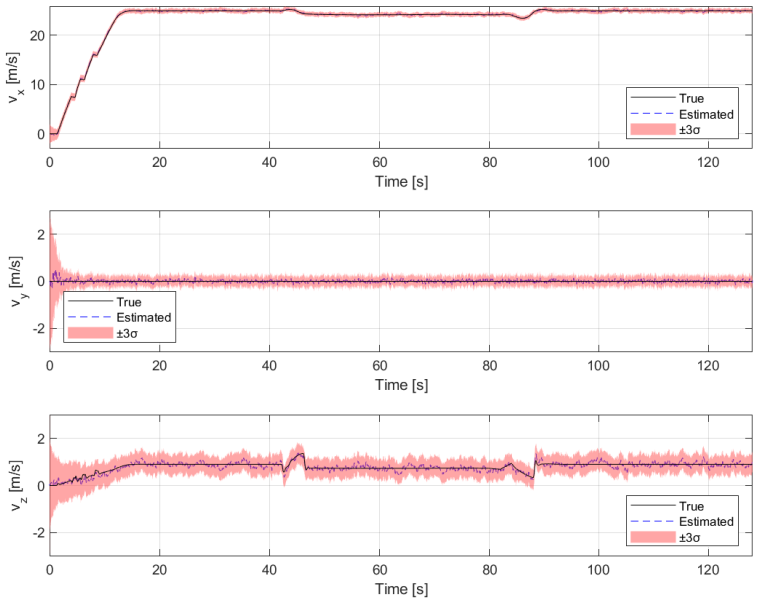

Figure A.24: Configuration 3: Estimated translational velocities in the *Uphill Track*.

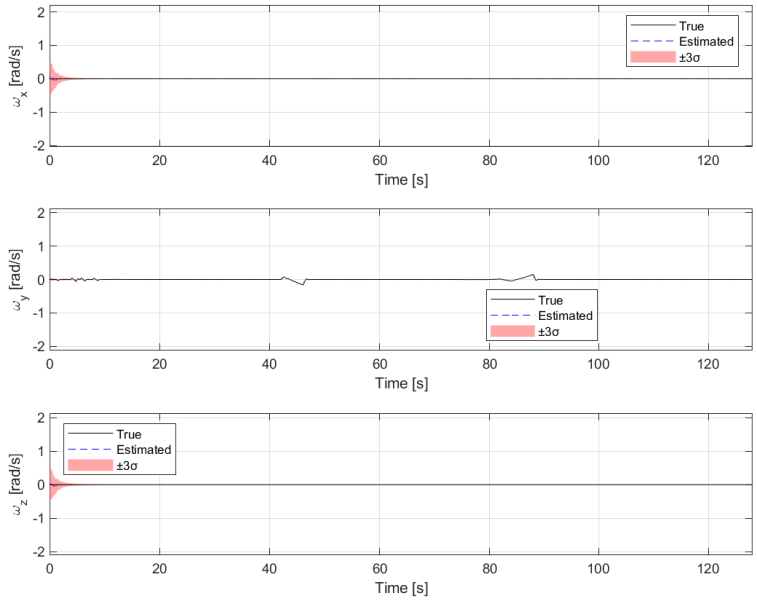

Figure A.25: Configuration 3: Estimated rotational velocities in the *Uphill Track*.

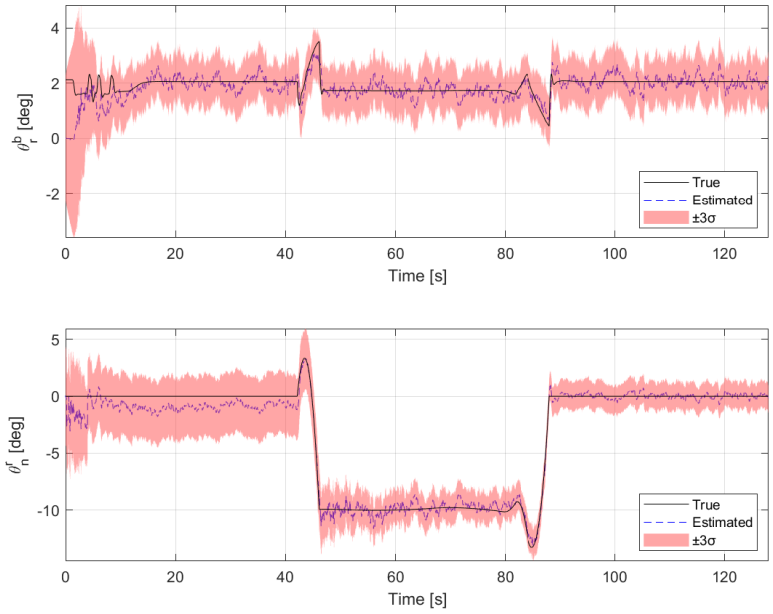

Figure A.26: Configuration 3: Estimated body pitch w.r.t. road frame (θ_r^b) and estimated road slope w.r.t. navigation frame (θ_n^r) in the *Uphill Track*.

B

Noise Covariance Matrix Tuning

$$\Sigma_{\mathbf{x}} = \begin{bmatrix}
10^{-3} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 10^{-3} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 10^{-3} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 10^{-1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 10^{-1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 10^{-1} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 10^{-7} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-7} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-3} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-3} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-3} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-4} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-8} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-8} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-8} & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-12} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-12} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 10^{-12}
\end{bmatrix}$$

$$\Sigma_{\mathbf{y1}} = \begin{bmatrix}
0.04^2 & 0 & 0 & 0 & 0 & 0 \\
0 & 0.04^2 & 0 & 0 & 0 & 0 \\
0 & 0 & 0.04^2 & 0 & 0 & 0 \\
0 & 0 & 0 & 0.0017^2 & 0 & 0 \\
0 & 0 & 0 & 0 & 0.0017^2 & 0 \\
0 & 0 & 0 & 0 & 0 & 0.0017^2
\end{bmatrix}$$

$$\Sigma_{\mathbf{y2}} = \begin{bmatrix}
0.05^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0.05^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0.05^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0.05^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0.05^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0.05^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 3 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.04^2 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.04^2 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.04^2 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0017^2 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0017^2 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0017^2
\end{bmatrix}$$

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY