NONLINEAR FILTERING WITH APPLICATIONS TO ASTRODYNAMICS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Zvonimir Stojanovski

May 2023

© 2023 Zvonimir Stojanovski

ALL RIGHTS RESERVED

NONLINEAR FILTERING WITH APPLICATIONS TO ASTRODYNAMICS

Zvonimir Stojanovski, Ph.D.

Cornell University 2023

Filtering methods are among the basic tools used for state estimation and control in many areas of engineering, particularly in aerospace. For some systems, the problem of state estimation remains challenging for various reasons, including highly nonlinear dynamics and measurements, as well as unusual process and measurement noise distributions. Although a wide variety of filtering methods have been developed over the past few decades, both for general purposes and specific applications, the need for developing new methods remains strong. This thesis is devoted to the development and improvement of nonlinear filtering methods and their applications to challenging problems in astrodynamics.

Most filters currently in use compute the mean and covariance of the state of the system. This is sufficient to describe a Gaussian or nearly Gaussian distribution, but to describe a wider family of possible distributions, higher-order moments are required. In this thesis, we develop the Higher-Order Unscented Estimator, which accounts for skewness and kurtosis in addition to the mean and covariance. We test this filter in simulations of three nonlinear dynamical systems and find that it is more robust than other estimators in the presence of outliers in the process and measurement noise.

Next, we apply the Square Root Sigma Point Filter to the problem of autonomous cross-calibration for Earth-imaging satellites. This is a novel application of nonlinear filtering. We develop a simulation framework for a constellation

of Earth-imaging satellites, which includes detailed models of the satellites' dynamics and cameras, and we demonstrate the application of this method to two satellites in coplanar low-Earth orbits.

Finally, we apply a Gaussian mixture sigma point filter to the problem of exoplanet orbit fitting. To avoid the singularities associated with the classical orbital elements, we introduce a new nonsingular parametrization for Keplerian orbits. With these new elements, the filter works well, and estimates of the classical elements can be obtained by simple transformations.

Throughout this work, we show that nonlinear filtering methods can be successfully applied to new dynamical systems. This is achieved through careful modeling and parametrization of the system state and measurements and the development of new filtering techniques.

BIOGRAPHICAL SKETCH

Zvonimir Stojanovski was born in Split, Croatia, and after several years of wandering, his family settled in Carterville, Illinois. After graduating from Carterville High School in 2014, he came to Cornell University, where he liked it so much that he decided to stay there for as long as he could. As an undergraduate, he did research on the automated design and optimization of satellite constellations in Prof. Daniel Selva's Systems Engineering, Architecture, and Knowledge (SEAK) Lab. During that time, he also acted in several plays, was part of the Midnight Comedy Troupe, and wrote and directed *Three Lost*, a dark musical comedy.

Zvonimir received his Bachelor of Science in Mechanical Engineering in May 2018. That summer, he interned at the Air Force Research Laboratory (AFRL) Space Vehicles Directorate as part of the Phillips Scholars Program. In the fall, he began his PhD program in Aerospace Engineering at Cornell. From Fall 2018 to Spring 2020, he did research on multitarget tracking in Prof. Silvia Ferrari's Laboratory for Intelligent Systems and Controls (LISC). Since Fall 2020, he has been a member of Prof. Dmitry Savransky's Space Imaging and Optical Systems (SIOS) Lab, where his work has focused on nonlinear filtering and its applications to spacecraft state estimation, as well as exoplanet orbit fitting. In Summer 2021 and 2022, he interned at the Johns Hopkins University Applied Physics Laboratory (JHU/APL).

ACKNOWLEDGEMENTS

Learning from Professor Dmitry Savransky was an incredible experience. I am deeply grateful for the opportunity to work with him and for his guidance on my journey from low Earth orbit to distant exoplanets. I am indebted to Professors Mason Peck, Mark Campbell, and Richard Rand for their invaluable advice and support. I would also like to thank my labmates Jacob Shapiro, Dean Keithly, Duan Li, Corey Spohn, Grace Genszler, Jackson Kulik, Rachel Oliver, and Colby Merrill for their friendship and encouragement. Finally, I thank my parents, Dubravka and Goran, and my brother, Ivan, who have always been there for me.

This work was supported in part by the NASA Space Technology Research Grants Early Career Faculty program under NASA grant 80NSSC20K0068.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

INTRODUCTION: ESTIMATION IN AEROSPACE AND ASTRONOMY

In many aerospace applications, it is necessary to obtain estimates of the state of a dynamical system in real time from incomplete and noisy measurements. The Kalman filter, first introduced in 1960, forms the basis of most modern recursive estimators. It alternates between predicting the state of the system and updating the estimate using measurements. The invention of the Kalman filter came just in time for the development of spaceflight—by 1969, it was used on the Apollo navigation computer [25].

Although the original Kalman filter is only for systems with linear dynamics and measurements, it has been extended to nonlinear systems. To adapt to the requirements of a wide range of systems, researchers have developed numerous filtering methods. Some of the most notable methods are listed below.

- *Kalman-Bucy filter:* A continuous-time version of the Kalman filter, with linear dynamics and measurements and white noise.
- *Extended Kalman filter (EKF):* Linearizes the dynamics and measurement functions about the means using Jacobian matrices.
- *Particle Filter (PF):* Approximates the probability density of the state using randomly-generated points.
- *Unscented Kalman filter (UKF):* Approximates the probability density of the state using a finite, deterministic set of points.

In this thesis, we investigate the application of filters to problems in aerospace and astronomy. The foundation of all of these projects is nonlinear Bayesian estimation, and in particular the unscented Kalman filter. The theory and methods of Bayesian estimation are reviewed in Chapter 1.

The first main result in the thesis is a new filter, the Higher-Order Unscented Estimator (HOUSE), presented in Chapter 2. It is a new extension of the unscented Kalman filter with asymmetric sample points and weights chosen to match third- and fourth-order moments in addition to the mean and covariance. Explicit solutions are obtained for sample points and weights, making their evaluation efficient and robust, and rigorous constraints are derived for their applicability. The use of the new filter is demonstrated with three dynamical systems—an aircraft coordinated turn model, a rotating rigid body, and a projectile with drag—and filter performance is compared with that of the conventional unscented Kalman filter and conjugate unscented transform filters. The new filter is found to be more robust in most cases where the initial distribution, process noise, and measurement noise have a high kurtosis, in that it does not generate extreme outliers in the estimation error. Also, execution times for the new filter are found to be only slightly longer than for the conventional unscented Kalman filter and significantly shorter than for the conjugate unscented transform filters.

In Chaper 3, we present a methodology called Autonomous Cross-Calibration for Imaging Satellites (ACCIS). Here, each satellite extracts features from primary mission images and then transmits the features, along with its state estimate, to other satellites. Furthermore, each satellite uses comparisons of the image features, along with conventional state measurements, to estimate its position, attitude, and camera parameters via the Square Root Sigma Point Filter (SRSPF). We demonstrate the feasibility of this method using a simulation of two satellites.

In Chapter 4, we present a new method for exoplanet orbit fitting from direct observations, called the Nonsingular Estimator for Exoplanet Orbits (NEXO),

based on nonlinear filtering and a new set of orbital elements. These elements are nonsingular, and every possible set of values describes an elliptic orbit, making them well-suited to an unconstrained filter; they combine features of the Cohen-Hubbard nonsingular elements, as well as the Thiele-Innes constants. The filtering method combines the SRSPF with Gaussian mixture models. We validate this method using a large sample of simulated exoplanet orbits, and we apply it to real exoplanet data.

CHAPTER 1

BAYESIAN ESTIMATION

In probability theory and statistics, there are two main philosophies: the Bayesian philosophy, in which the probability is interpreted as the degree of belief or certainty, and the frequentist philosophy, in which it is interpreted as a frequency of outcomes in an experiment. In practice, the main difference is that Bayesian techniques start with a prior probability distribution, representing prior knowledge of the system, which is then updated using measurements. This structure is useful for estimation and forms the basis for the Kalman filter.

Since its introduction in 1960 [35], the Kalman filter has become an indispensable tool in control, navigation, and tracking, due to its simplicity and robustness [5]. Although the original Kalman filter assumes that the system dynamics and measurements are linear—one of the very few cases where the filter has a closed form—a wide variety of modifications to the Kalman filter, using various approximation methods, have been developed for handling nonlinear systems with possibly non-Gaussian noise. Fang et al. [19] provide an extensive survey of such filters within a Bayesian framework.

In this chapter, we review the theory of Bayesian estimation, with an emphasis on filtering techniques. In particular, we describe the Kalman filter and some of its extensions to nonlinear systems.

1.1 Notation

Throughout this thesis, vectors are denoted by lowercase boldface letters (e.g., $\mathbf{x}$) and matrices by uppercase boldface letters (e.g., $\mathbf{A}$). An overbar denotes the

mean or expected value of a random variable; e.g., $\bar{\mathbf{x}}$ is the expected value of $\mathbf{x}$. Covariance matrices are denoted by $\mathbf{P}$ with subscripts indicating the variables; e.g., $\mathbf{P}_{xx}$ is the covariance of $\mathbf{x}$, and $\mathbf{P}_{xy}$ is the cross-covariance of $\mathbf{x}$ and $\mathbf{y}$.

1.2 The Nonlinear Filtering Problem

In this section, we present the nonlinear filtering problem and describe the procedures for state prediction and correction.

Suppose that a system with state vector $\mathbf{x}$ evolves in discrete time according to

$$\mathbf{x}(k+1) = \mathbf{f}(\mathbf{x}(k), \mathbf{w}(k), k), \tag{1.1}$$

where $\mathbf{w}$ is the process noise, k is the time step, and $\mathbf{f}$ is a known state function. Furthermore, suppose that a measurement $\mathbf{z}$ is taken at each time step, given by

$$\mathbf{z}(k) = \mathbf{h}(\mathbf{x}(k), \mathbf{n}(k), k), \tag{1.2}$$

where $\mathbf{n}$ is the measurement noise, and $\mathbf{h}$ is a known measurement function. We assume that $\mathbf{w}$ and $\mathbf{n}$ are mutually independent and also independent of $\mathbf{x}$. Furthermore, we assume that $\mathbf{w}(k')$ is independent of $\mathbf{w}(k)$ for $k' \neq k$ and that the same is true for $\mathbf{n}$. Thus, the state of the system is a Markov process, which can be informally stated as, "The future is independent of the past if the present is known" [5]. In addition, we assume that the noise is zero-mean, i.e., that $\bar{\mathbf{w}}(k) = \mathbf{0}$ and $\bar{\mathbf{n}}(k) = \mathbf{0}$.

The operation of the filter can be divided into two alternating steps: *prediction*, or finding the distribution of $\mathbf{x}(k+1)$ given the distributions of $\mathbf{x}(k)$ and $\mathbf{w}(k)$;

and *correction* or *update*, or finding the distribution of $\mathbf{x}(k)$ given a measurement $\mathbf{z}(k)$, the distribution of $\mathbf{n}(k)$, and the prior predicted distribution of $\mathbf{x}(k)$.

Except in some special cases, the true distribution of a variable—i.e., its probability density function (PDF)—cannot be determined analytically. This is especially true if $\mathbf{f}$ and $\mathbf{h}$ are not known explicitly, e.g., if their evaluation requires numerical integration of equations of motion. While the PDF can be approximated very accurately using numerical methods, such as finite difference or Gaussian mixture models, this is too computationally expensive for many applications. Instead, the distribution may be described by a finite set of generalized moments. For any random variable $\mathbf{x} \in \Omega$ with PDF p and function $\phi : \Omega \to \mathbb{R}$, we can define a generalized moment $\bar{\phi} \in \mathbb{R}$ as

$$\bar{\phi} = \mathrm{E}[\phi(\mathbf{x})] = \int_{\Omega} p(\mathbf{x})\phi(\mathbf{x})\mathrm{d}\mathbf{x}. \tag{1.3}$$

Furthermore, for any functions $\mathbf{g} : \Omega \to \Omega'$ and $\phi' : \Omega' \to \mathbb{R}$, we can compute a generalized moment of $\mathbf{y} = \mathbf{g}(\mathbf{x})$

$$\bar{\phi}' = \mathrm{E}[\phi'(\mathbf{y})] = \int_{\Omega} \phi'(\mathbf{g}(\mathbf{x}))p(\mathbf{x})\mathrm{d}\mathbf{x}. \tag{1.4}$$

The simplest generalized moments of a random variable $\mathbf{x}$ are its mean or expected value $\bar{\mathbf{x}}$ and covariance $\mathbf{P}_{xx}$, which is defined as

$$\mathbf{P}_{xx} = \mathrm{E}[(\mathbf{x} - \bar{\mathbf{x}})(\mathbf{x} - \bar{\mathbf{x}})^{\mathrm{T}}]. \tag{1.5}$$

In practice, for many estimators, these are the only generalized moments that are considered. Furthermore, in some cases, it is reasonable to assume that the PDF p is of a specific type (e.g., Gaussian) or belongs to a broader family of distributions (e.g., the Pearson family). Then, p can be fully characterized by a finite set of generalized moments.

Given these definitions, we can now describe in detail the processes of prediction and correction.

1.2.1 Prediction

The most general form of the prediction step is given by the Chapman-Kolmogorov equation

$$p_x^{(k+1)}(\mathbf{x}) = \int p_x^{(k+1|k)}(\mathbf{x}|\mathbf{x}')p_x^{(k)}(\mathbf{x}')\mathrm{d}\mathbf{x}', \tag{1.6}$$

where $p_x^{(k)}$ is the PDF of $\mathbf{x}(k)$, and $p_x^{(k+1|k)}$ is the transition density from $\mathbf{x}(k)$ to $\mathbf{x}(k+1)$, i.e., the conditional PDF of $\mathbf{x}(k+1)$ given $\mathbf{x}(k)$ [5]. From (1.1), we see that the only random effects in the transition from $\mathbf{x}(k)$ to $\mathbf{x}(k+1)$ are due to the process noise $\mathbf{w}(k)$; therefore, we have the conditional probability density for $\mathbf{x}(k+1) = \mathbf{x}$ given $\mathbf{x}(k) = \mathbf{x}'$

$$p_x^{(k+1|k)}(\mathbf{x}|\mathbf{x}') = \int \delta(\mathbf{x} - \mathbf{f}(\mathbf{x}', \mathbf{w}, k))p_w^{(k)}(\mathbf{w})\mathrm{d}\mathbf{w}, \tag{1.7}$$

where $p_w^{(k)}$ is the PDF of $\mathbf{w}(k)$, and δ is the Dirac delta function.

A generalized moment $\bar{\phi}$ of $\mathbf{x}(k+1)$ can be obtained by substituting (1.6) into (1.3), which gives

$$\bar{\phi} = \int \phi(\mathbf{x}) \int p_x^{(k+1|k)}(\mathbf{x}|\mathbf{x}')p_x^{(k)}(\mathbf{x}')\mathrm{d}\mathbf{x}'\mathrm{d}\mathbf{x}, \tag{1.8}$$

and substituting (1.7) gives

$$\bar{\phi} = \int \phi(\mathbf{x}) \iint \delta(\mathbf{x} - \mathbf{f}(\mathbf{x}', \mathbf{w}, k))p_w^{(k)}(\mathbf{w})\mathrm{d}\mathbf{w}p_x^{(k)}(\mathbf{x}')\mathrm{d}\mathbf{x}'\mathrm{d}\mathbf{x}. \tag{1.9}$$

Using the properties of the Dirac delta function, this simplifies to

$$\bar{\phi} = \iint \phi(\mathbf{f}(\mathbf{x}', \mathbf{w}, k))p_w^{(k)}(\mathbf{w})\mathrm{d}\mathbf{w}p_x^{(k)}(\mathbf{x}')\mathrm{d}\mathbf{x}', \tag{1.10}$$

which is the expected value of $\phi(\mathbf{f}(\mathbf{x}, \mathbf{w}, k))$. At this point, in order to simplify the notation in the development of the filter, it is useful to define the augmented state

$$\mathbf{y}_{\mathrm{P}}(k) = \begin{bmatrix} \mathbf{x}(k) \\ \mathbf{w}(k) \end{bmatrix}, \tag{1.11}$$

with PDF $p_{\mathrm{YP}}^{(k)}$, given by

$$p_{\mathrm{YP}}^{(k)}\left(\begin{bmatrix} \mathbf{x} \\ \mathbf{w} \end{bmatrix}\right) = p_{w}^{(k)}(\mathbf{w}) p_{x}^{(k)}(\mathbf{x}). \tag{1.12}$$

Then, the generalized moment is given in the form of (1.3) by

$$\bar{\phi} = \int \phi(\mathbf{y}) p_{\mathrm{YP}}^{(k)}(\mathbf{y}) \mathrm{d}\mathbf{y}, \tag{1.13}$$

where

$$\phi\left(\begin{bmatrix} \mathbf{x} \\ \mathbf{w} \end{bmatrix}\right) = \phi(\mathbf{x}). \tag{1.14}$$

1.2.2 Correction

Due to (1.2), the measurement likelihood function is given by

$$p_{z|x}^{(k)}(\mathbf{z}|\mathbf{x}) = \int \delta(\mathbf{z} - \mathbf{h}(\mathbf{x}(k), \mathbf{n}(k), k)) p_{n}^{(k)}(\mathbf{n}) \mathrm{d}\mathbf{n}, \tag{1.15}$$

where $p_{n}^{(k)}(\mathbf{n})$ is the PDF of $\mathbf{n}(k)$. The posterior, or corrected, PDF for $\mathbf{x}(k)$—i.e., the conditional probability density of $\mathbf{x}(k)$ given $\mathbf{z}(k)$—is given by Bayes' rule

$$p_{x|z}^{(k)}(\mathbf{x}|\mathbf{z}) = \frac{p_{z|x}^{(k)}(\mathbf{z}|\mathbf{x}) p_{x}^{(k)}(\mathbf{x})}{\int p_{z|x}^{(k)}(\mathbf{z}|\mathbf{x}') p_{x}^{(k)}(\mathbf{x}') \mathrm{d}\mathbf{x}'}, \tag{1.16}$$

where $p_{x}^{(k)}$ is the prior PDF of $\mathbf{x}(k)$ and $p_{x|z}^{(k)}$ is the posterior PDF of $\mathbf{x}(k)$.

For computing the generalized moments of $\mathbf{z}(k)$ and $\mathbf{x}(k)$, we use an approach similar to that in the prediction step, with an augmented state

$$\mathbf{y}_{\mathrm{C}}(k) = \begin{bmatrix} \mathbf{x}(k) \\ \mathbf{n}(k) \end{bmatrix}. \tag{1.17}$$

Obtaining generalized moments for the posterior density of $\mathbf{x}$, however, is more difficult. Usually, this is done by defining an estimator $\hat{\mathbf{x}}$ of $\mathbf{x}$ that is a function of $\mathbf{z}$ and the prior distribution of $\mathbf{x}$. (Here, we drop the k suffixes for brevity.) The estimator error is then defined as

$$\boldsymbol{\epsilon} = \mathbf{x} - \hat{\mathbf{x}}. \tag{1.18}$$

It is desirable that an estimator be *unbiased*, which means that $\bar{\boldsymbol{\epsilon}} = \mathbf{0}$, or equivalently, that $\mathrm{E}[\hat{\mathbf{x}}(\mathbf{z})] = \bar{\mathbf{x}}$. Furthermore, to make the estimator error as small as possible, we seek to minimize the mean-square error (MSE), which is defined as $\mathrm{E}[\boldsymbol{\epsilon}^{\mathrm{T}}\boldsymbol{\epsilon}]$.

The optimal estimator—the one that is unbiased and minimizes the MSE—is the mean of the Bayesian posterior distribution

$$\hat{\mathbf{x}} = \frac{\int \mathbf{x} p_{z|x}(\mathbf{z}|\mathbf{x}) p_x(\mathbf{x}) \mathrm{d}\mathbf{x}}{\int p_{z|x}(\mathbf{z}|\mathbf{x}) p_x(\mathbf{x}) \mathrm{d}\mathbf{x}}. \tag{1.19}$$

In some special cases, such as when $\mathbf{h}$ is linear and p_x and p_n are Gaussian, this can be expressed in closed form. In many other cases, however, the optimal estimator does not have a useful analytical form, and the computational cost of approximating the PDFs numerically is too high for real-time applications. Instead, we may use a linear estimator of the form

$$\hat{\mathbf{x}} = \mathbf{A}\mathbf{z} + \mathbf{b}, \tag{1.20}$$

where $\mathbf{A}$ is a constant matrix, and $\mathbf{b}$ is a constant vector. Among unbiased linear estimators, the one that gives the smallest MSE, called the linear minimum

mean-square estimator (LMMSE), is given by [5]

$$\hat{\mathbf{x}} = \bar{\mathbf{x}} + \mathbf{P}_{xz}\mathbf{P}_{zz}^{-1}(\mathbf{z} - \bar{\mathbf{z}}), \tag{1.21}$$

with error covariance

$$\mathbf{P}_{\epsilon\epsilon} = \mathbf{P}_{xx} - \mathbf{P}_{xz}\mathbf{P}_{zz}^{-1}\mathbf{P}_{xz}^{\mathrm{T}}. \tag{1.22}$$

In practice, we often use $\mathbf{P}_{\epsilon\epsilon}$ as an approximation of $\mathbf{P}_{xx|z}$, the posterior covariance of $\mathbf{x}$.

1.3 Linear and Linearized Filtering Methods

This section starts with a description of the original Kalman filter, which gives closed-form solutions to the filtering problem for linear systems. Then, we present the extended Kalman filter, which is a straightforward extension to nonlinear systems based on linear approximations.

1.3.1 The Kalman Filter

The Kalman filter gives an exact solution to the estimation problem with (1.1) and (1.2) in the linear form

$$\mathbf{x}(k+1) = \mathbf{F}(k)\mathbf{x}(k) + \mathbf{G}(k)\mathbf{w}(k) \tag{1.23}$$

and

$$\mathbf{z}(k) = \mathbf{H}(k)\mathbf{x}(k) + \mathbf{C}(k)\mathbf{n}(k). \tag{1.24}$$

For the prediction step, we apply (1.23) and the properties of the mean and covariance to obtain

$$\bar{\mathbf{x}}(k+1|k) = \mathbf{F}(k)\bar{\mathbf{x}}(k) \tag{1.25}$$

and

$$\mathbf{P}_{xx}(k+1|k) = \mathbf{F}(k)\mathbf{P}_{xx}(k)\mathbf{F}^{\mathrm{T}}(k) + \mathbf{G}(k)\mathbf{P}_{ww}(k)\mathbf{G}^{\mathrm{T}}(k). \tag{1.26}$$

For details, see [5]. As for the correction step, we compute the measurement mean

$$\bar{\mathbf{z}}(k) = \mathbf{H}(k)\bar{\mathbf{x}}(k|k-1), \tag{1.27}$$

the measurement covariance

$$\mathbf{P}_{zz}(k) = \mathbf{H}(k)\mathbf{P}_{xx}(k|k-1)\mathbf{H}^{\mathrm{T}}(k) + \mathbf{C}(k)\mathbf{P}_{nn}(k)\mathbf{C}^{\mathrm{T}}(k), \tag{1.28}$$

and the cross-covariance between the state and the measurement

$$\mathbf{P}_{xz}(k) = \mathbf{P}_{xx}(k|k-1)\mathbf{H}^{\mathrm{T}}(k). \tag{1.29}$$

From this, we compute the Kalman gain

$$\mathbf{K}(k) = \mathbf{P}_{xz}(k)\mathbf{P}_{zz}^{-1}(k). \tag{1.30}$$

Then, we use the LMMSE formulae (1.21) and (1.22) to obtain the updated mean

$$\bar{\mathbf{x}}(k) = \bar{\mathbf{x}}(k) + \mathbf{K}(k)(\mathbf{z}(k) - \bar{\mathbf{z}}(k)) \tag{1.31}$$

and covariance

$$\mathbf{P}_{xx}(k) = \mathbf{P}_{xx}(k) - \mathbf{K}(k)\mathbf{P}_{xz}^{\mathrm{T}}(k). \tag{1.32}$$

Contrary to popular belief, the Kalman filter does not require that the distributions of $\mathbf{w}(k)$ and $\mathbf{n}(k)$ be Gaussian. However, that is the only case when the Kalman correction is optimal in the mean-square error sense (and not merely the LMMSE).

1.3.2 The Extended Kalman Filter

Although the Kalman filter only solves the linear estimation problem, its simplicity and robustness have motivated the development of similar but suboptimal

filters for nonlinear systems. The most straightforward is the extended Kalman filter (EKF), which uses a linear approximation of the nonlinear system about the mean, using Jacobian matrices.

In the EKF, we evaluate the predicted state and measurement mean using the full nonlinear model at the prior mean, that is,

$$\bar{\mathbf{x}}(k+1) = \mathbf{f}(\bar{\mathbf{x}}(k), \mathbf{0}, k), \tag{1.33}$$

$$\bar{\mathbf{z}}(k) = \mathbf{h}(\bar{\mathbf{x}}(k), \mathbf{0}, k). \tag{1.34}$$

As for computing the covariances and cross-covariances, we define

$$\mathbf{F}(k) = \frac{\partial \mathbf{f}}{\partial \mathbf{x}}(\bar{\mathbf{x}}(k), \mathbf{0}, k), \tag{1.35}$$

$$\mathbf{G}(k) = \frac{\partial \mathbf{f}}{\partial \mathbf{w}}(\bar{\mathbf{x}}(k), \mathbf{0}, k), \tag{1.36}$$

$$\mathbf{H}(k) = \frac{\partial \mathbf{h}}{\partial \mathbf{x}}(\bar{\mathbf{x}}(k), \mathbf{0}, k), \tag{1.37}$$

$$\mathbf{C}(k) = \frac{\partial \mathbf{h}}{\partial \mathbf{x}}(\bar{\mathbf{x}}(k), \mathbf{0}, k). \tag{1.38}$$

Then, we apply the same procedures as for the original Kalman filter. This method has been widely used for a variety of applications. However, there are many dynamical systems for which the linearization is not a sufficiciently accurate approximation. Also, the evaluation of the Jacobians can be complicated and computationally expensive.

1.4 The Unscented Kalman Filter

As a more efficient alternative to the previously developed EKF and particle filter (PF), Julier and Uhlmann [34] introduced the unscented Kalman filter (UKF),

also known as the sigma point filter (SPF), which evaluates the state and measurement functions at a finite, deterministic set of points, called sigma points, and uses weighted sums to compute the predicted and updated mean and covariance in the filter. This process is referred to as the unscented transform (UT). The points and weights are chosen so that the computed mean and covariance are exact up to the second and first order in the state, respectively. In cases where the state distribution is symmetric, the third moments are preserved as well, due to the symmetry of the sigma points. As in the original Kalman filter [35], no assumptions are made about the type of distribution—in particular, the distribution need not be Gaussian.

1.4.1 The Unscented Transform

In many cases, even when the exact PDF is known, it is impossible to compute the generalized moment using (1.3) analytically. However, the integral in this equation can be approximated with reasonable accuracy and efficiency by evaluating ϕ at a deterministic set of sample points $\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(N)}$ with corresponding weights $w_1, \ldots, w_N$ and replacing the integral with a weighted sum

$$\bar{\phi} \approx \sum_{j=1}^{N} w_j \phi(\mathbf{x}^{(j)}). \tag{1.39}$$

This is the principle behind the unscented transform. It is very closely related to a class of methods for numerical integration; these methods are usually referred to as *quadratures* in the univariate case and *cubatures* in the multivariate case.

Choosing the points $\mathbf{x}^{(j)}$ and weights w_j is equivalent to finding a cubature rule with p as a weighting function. Let $\mathcal{F}$ be the family of functions for which

this cubature is exact; that is

$$\mathcal{F} = \left\{ \phi : \Omega \to \mathbb{R} \middle| \int_{\Omega} \phi(\mathbf{x}) p(\mathbf{x}) \mathrm{d}\mathbf{x} = \sum_{j=1}^{N} w_j \phi(\mathbf{x}^{(j)}) \right\}. \tag{1.40}$$

For example, $\mathcal{F}$ might include all polynomials up to a certain degree. We want to choose the points $\mathbf{x}^{(j)}$ and weights w_j so that all functions ϕ that we use for generalized moments $\bar{\phi}$ can be approximated with sufficient accuracy by functions in $\mathcal{F}$. Specifically, suppose that we want to compute $\bar{\phi}' = \mathrm{E}[\phi'(\mathbf{y})]$, where $\mathbf{y} = \mathbf{g}(\mathbf{x})$, for a polynomial ϕ of degree M and an arbitrary function $\mathbf{g}$. If we want to account for K-th order terms in the Taylor expansion of $\mathbf{g}$, then $\mathcal{F}$ must contain polynomials up to degree KM.

In the original unscented transform [34], for a random variable $\mathbf{x} \in \mathbb{R}^n$ with mean $\bar{\mathbf{x}}$ and covariance $\mathbf{P}_{xx}$, the sample points, referred to as *sigma points*, are given by

$$\mathbf{x}^{(j)} = \begin{cases} \bar{\mathbf{x}} + \sqrt{n+\kappa}\mathbf{c}^{(j)}, & 1 \le j \le n \\ \bar{\mathbf{x}} - \sqrt{n+\kappa}\mathbf{c}^{(j-n)} & n+1 \le j \le 2n \\ \bar{\mathbf{x}} & j = 2n+1 \end{cases}, \tag{1.41}$$

where $\mathbf{c}^{(j)}$ denotes the j-th column of $\sqrt{\mathbf{P}_{xx}}$. The corresponding weights are

$$w_j = \begin{cases} \dfrac{1}{2(n+\kappa)}, & 1 \le j \le 2n \\ \dfrac{\kappa}{n+\kappa}, & 2n+1 \end{cases}, \tag{1.42}$$

where κ is a tuning factor. These points and weights form a cubature rule that is exact for polynomials up to the second order. The square root of the covariance matrix is usually taken to be the lower-triangular Cholesky decomposition, due to its efficiency and numerical stability. We adopt this convention in our work.

1.4.2 Procedures

We now describe how the unscented transform is applied to the prediction and correction steps in the unscented Kalman filter.

Prediction

For the general dynamical model (1.1), we define the augmented state

$$\mathbf{y}_{\mathrm{P}}(k) = \begin{bmatrix} \mathbf{x}(k) \\ \mathbf{w}(k) \end{bmatrix}, \tag{1.43}$$

which has mean

$$\bar{\mathbf{y}}_{\mathrm{P}}(k) = \begin{bmatrix} \bar{\mathbf{x}}(k) \\ \mathbf{0} \end{bmatrix} \tag{1.44}$$

and covariance

$$\mathbf{P}_{yy,\mathrm{P}}(k) = \begin{bmatrix} \mathbf{P}_{xx}(k) & \mathbf{0} \\ \mathbf{0} & \mathbf{P}_{ww}(k) \end{bmatrix}. \tag{1.45}$$

To these we apply the unscented transform to obtain the sample points $\mathbf{y}_{\mathrm{C}}^{(j)}$, which comprise $\mathbf{x}^{(j)}$ and $\mathbf{w}^{(j)}$, with corresponding weights w_j. Then, the predicted states at the sigma points are given by

$$\mathbf{x}_{\mathrm{P}}^{(j)} = \mathbf{f}(\mathbf{x}^{(j)}, \mathbf{w}^{(j)}, k). \tag{1.46}$$

From these points, we compute the predicted state mean

$$\bar{\mathbf{x}}(k+1|k) = \sum_{j=1}^{N} w_j \mathbf{x}_{\mathrm{P}}^{(j)} \tag{1.47}$$

and covariance

$$\mathbf{P}_{xx}(k+1|k) = \sum_{j=1}^{N} w_j (\mathbf{x}_{\mathrm{P}}^{(j)} - \bar{\mathbf{x}}(k+1|k))(\mathbf{x}_{\mathrm{P}}^{(j)} - \bar{\mathbf{x}}(k+1|k))^{\mathrm{T}}. \tag{1.48}$$

In the special (but common) case where the process noise is additive

$$\mathbf{x}(k+1) = \mathbf{f}(\mathbf{x}(k), k) + \mathbf{w}(k), \tag{1.49}$$

we can compute the sigma points and weights only for $\mathbf{x}(k)$ based on $\bar{\mathbf{x}}(k)$ and $\mathbf{P}_{xx}(k)$, from which we evaluate

$$\mathbf{x}_{\mathrm{P}}^{(j)} = \mathbf{f}(\mathbf{x}^{(j)}, k). \tag{1.50}$$

Then, the mean is again given by (1.47), and to the covariance (1.48) we add the noise covariance $\mathbf{P}_{ww}$.

Correction

In the correction step, we use the general measurement model (1.2) to define an augmented state

$$\mathbf{y}_{\mathrm{C}}(k) = \begin{bmatrix} \mathbf{x}(k) \\ \mathbf{n}(k) \end{bmatrix}, \tag{1.51}$$

with mean

$$\bar{\mathbf{y}}_{\mathrm{C}}(k) = \begin{bmatrix} \bar{\mathbf{x}}(k|k-1) \\ \mathbf{0} \end{bmatrix} \tag{1.52}$$

and covariance

$$\mathbf{P}_{yy,\mathrm{C}}(k) = \begin{bmatrix} \mathbf{P}_{xx}(k|k-1) & \mathbf{0} \\ \mathbf{0} & \mathbf{P}_{nn}(k) \end{bmatrix}. \tag{1.53}$$

Then, as in the prediction step, we compute the sample points $\mathbf{x}^{(j)}$ and $\mathbf{n}^{(j)}$ and their corresponding weights w_j, we and evaluate the sample measurements

$$\mathbf{z}^{(j)} = \mathbf{h}(\mathbf{x}^{(j)}, \mathbf{n}^{(j)}, k). \tag{1.54}$$

With these sample points, we compute the measurement mean

$$\bar{\mathbf{z}}(k) = \sum_{j=1}^{N} w_j \mathbf{z}^{(j)}, \tag{1.55}$$

covariance

$$\mathbf{P}_{zz}(k) = \sum_{j=1}^{N} w_j(\mathbf{z}^{(j)} - \bar{\mathbf{z}}(k))(\mathbf{z}^{(j)} - \bar{\mathbf{z}}(k))^{\mathrm{T}}, \tag{1.56}$$

and cross-covariance

$$\mathbf{P}_{xz}(k) = \sum_{j=1}^{N} w_j(\mathbf{x}^{(j)} - \bar{\mathbf{x}}(k))(\mathbf{z}^{(j)} - \bar{\mathbf{z}}(k))^{\mathrm{T}}. \tag{1.57}$$

Then, we apply the LMMSE correction, as in the original Kalman filter, to obtain the updated $\bar{\mathbf{x}}(k)$ and $\mathbf{P}_{xx}(k)$. If the measurement noise is additive,

$$\mathbf{z}(k) = \mathbf{h}(\mathbf{x}(k), k) + \mathbf{n}(k), \tag{1.58}$$

we can apply a simplified approach similar to the one for the prediction step, in this case adding $\mathbf{P}_{nn}$ to $\mathbf{P}_{zz}$ in (1.56).

1.4.3 The Square-Root Sigma Point Filter

The Square Root Sigma Point Filter (SRSPF), introduced by Brunke and Campbell [11], is a variant of the UKF that propagates $\sqrt{\mathbf{P}_{xx}}$ rather than $\mathbf{P}_{xx}$. The main advantage of this approach is that it improves numerical stability. First, $\sqrt{\mathbf{P}_{xx}}$ has a lower condition number than $\mathbf{P}_{xx}$ [5]. Also, due to numerical errors, the LMMSE covariance update (1.22) can give a posterior covariance that is not positive-definite; the SRSPF, on the other hand, avoids this problem.

The SRSPF is based on the following observation. Suppose that, for a random variable $\mathbf{y}$, we have a set of sample points $\mathbf{y}^{(1)}, \ldots, \mathbf{y}^{(N)}$ with corresponding weights $w_1, \ldots, w_N$. We define the centralized sample points as

$$\mathbf{y}_c^{(j)} = \mathbf{y}^{(j)} - \bar{\mathbf{y}}. \tag{1.59}$$

We can arrange the centralized sample points in the matrix

$$\mathbf{Y}_c = \begin{bmatrix} \mathbf{y}_c^{(1)} & \cdots & \mathbf{y}_c^{(N)} \end{bmatrix} \tag{1.60}$$

and the weights in the diagonal matrix

$$\mathbf{W} = \operatorname{diag}(w_1, \ldots, w_N). \tag{1.61}$$

Then, the covariance of $\mathbf{y}$ can be expressed as

$$\mathbf{P}_{yy} = \mathbf{Y}_c \mathbf{W} \mathbf{Y}_c^{\mathrm{T}}. \tag{1.62}$$

Assuming that all of the weights are positive, this is equivalent to

$$\mathbf{P}_{yy} = \mathbf{S}\mathbf{S}^{\mathrm{T}}, \tag{1.63}$$

where

$$\mathbf{S} = \mathbf{Y}_c \sqrt{\mathbf{W}}. \tag{1.64}$$

Now, to obtain $\sqrt{\mathbf{P}_{yy}}$, rather than multiplying $\mathbf{S}\mathbf{S}^{\mathrm{T}}$ and then taking its Cholesky factorization, we can apply the QR factorization directly to $\mathbf{S}^{\mathrm{T}}$. That is, we can compute the factorization

$$\mathbf{S}^{\mathrm{T}} = \mathbf{Q}\mathbf{R}, \tag{1.65}$$

where $\mathbf{R}$ is an upper-triangular matrix, and $\mathbf{Q}$ is a matrix with orthonormal columns. This is equivalent to applying the Gram-Schmidt process to the columns of $\mathbf{S}^{\mathrm{T}}$. Then, using (1.63) and the fact that $\mathbf{Q}^{\mathrm{T}}\mathbf{Q} = \mathbf{I}$, we obtain

$$\sqrt{\mathbf{P}_{yy}} = \mathbf{R}^{\mathrm{T}}. \tag{1.66}$$

In the prediction step of the filter, we can apply this technique with $\mathbf{y}^{(j)} = \mathbf{x}_P^{(j)}$ to obtain $\sqrt{\mathbf{P}_{xx}(k+1)}$. As for the update step, we do the same with the LMMSE errors for sample values of $\mathbf{x}(k)$ and $\mathbf{z}(k)$; that is, we use

$$\mathbf{y}^{(j)} = \mathbf{x}^{(j)} - \mathbf{K}\mathbf{z}^{(j)} \tag{1.67}$$

to obtain the updated $\sqrt{\mathbf{P}_{xx}(k)}$.

1.4.4 Gaussian Mixtures

Like the original Kalman filter and the EKF, the UKF parametrizes probability distributions using only their mean and covariance. This is sufficient for many applications, particularly where the distributions are roughly Gaussian. In some cases, however, such as when the distributions are multimodal, the mean and covariance can be uninformative, or even misleading.

A Gaussian mixture model (GMM) provides a method for handling a very wide range of probability distributions, including multimodal ones, while taking advantage of some of the nice features of Gaussian distributions. For a random variable $\mathbf{x}$, a GMM has a PDF of the form

$$p_x(\mathbf{x}) = \sum_{j=1}^{M} W_j g(\mathbf{x}; \bar{\mathbf{x}}^{(j)}, \mathbf{P}_{xx}^{(j)}), \tag{1.68}$$

where $g(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes a Gaussian PDF with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$; the W_j are positive weights that sum to unity; and $\bar{\mathbf{x}}^{(j)}$ and $\mathbf{P}_{xx}^{(j)}$ are the means and covariances, respectively, of the mixture components.

An extension of the UKF using Gaussian mixtures, called the Sigma Point Gaussian Sum Filter (SPGSF), was developed by Šimandl and Duník [67]. For the prediction step, this filter simply performs the UKF prediction step individually for each component of the mixture, with mean $\bar{\mathbf{x}}^{(j)}$ and covariance $\mathbf{P}_{xx}^{(j)}$, and the weights W_j are unchanged. Similarly, for the update step, the SPGSF performs the UKF update step for the mean and covariance of each mixture component. However, in order to remain consistent with Bayes' rule, the updated weights are given by

$$W_j^+ = \frac{W_j \zeta_j}{\sum_{l=1}^{M} W_l \zeta_l}, \tag{1.69}$$

where

$$\zeta_j = g(\mathbf{z}; \bar{\mathbf{z}}^{(j)}, \mathbf{P}_{zz}^{(j)}), \tag{1.70}$$

in which $\bar{\mathbf{z}}^{(j)}$ and $\mathbf{P}_{zz}^{(j)}$ are the measurement mean and covariance obtained from $\bar{\mathbf{x}}^{(j)}$ and $\mathbf{P}_{xx}^{(j)}$ using the unscented transform. Thus, we can handle very complicated state distributions using the UKF as a building block. This approach can also be applied to variants of the UKF; in particular, Šimandl and Duník apply it to the SRSPF [67].

The main disadvantage of this method, particularly for real-time estimation, is the computational cost, which increases linearly with the number of mixture components. Nevertheless, the SPGSF prediction and update procedures lend themselves well to parallelization.

1.4.5 Credible Intervals

The outputs of the UKF are the mean $\bar{\mathbf{x}}$ and covariance $\mathbf{P}_{xx}$ of the system state $\mathbf{x}$. While these contain a great deal of information about the distribution of $\mathbf{x}$, in some cases it may be difficult to interpret the results, particularly if the components of $\mathbf{x}$ cannot be easily visualized. In this section, we describe methods for summarizing the output of the UKF using Bayesian credbile intervals (not to be confused with frequentist confidence intervals) for non-angular and angular functions of $\mathbf{x}$.

To obtain credible intervals that only depend on the moments of the variables, without making any further assumptions about the forms of the probability distributions, we apply the Chebyshev and Markov inequalities. For a random variable X with mean $\bar{X}$ and standard deviation σ_X, the Chebyshev inequality

states that

$$\Pr[|X - \bar{X}| \geq \varepsilon] \leq \frac{\sigma_X^2}{\varepsilon^2} \tag{1.71}$$

for any $\varepsilon > 0$. The closely related Markov inequality states that, if $X > 0$, then

$$\Pr[X \geq \varepsilon] \leq \frac{\bar{X}}{\varepsilon}. \tag{1.72}$$

These two inequalities hold regardless of the distribution of X [5].

Non-Angular Variables

Let y be a function of $\mathbf{x}$. With the unscented transform, we can approximate this variable's mean

$$\bar{y} = \sum_{j=1}^{N} w_j y(\mathbf{x}^{(j)}) \tag{1.73}$$

and standard deviation

$$\sigma_y = \sqrt{\sum_{j=1}^{N} w_j (y(\mathbf{x}^{(j)}) - \bar{y})^2}. \tag{1.74}$$

Then, by applying the Chebyshev inequality (1.71), we obtain the credible interval

$$C_y = \left[\bar{y} - \frac{\sigma_y}{\sqrt{\alpha}}, \bar{y} + \frac{\sigma_y}{\sqrt{\alpha}} \right]. \tag{1.75}$$

with probability $1 - \alpha$.

Angular Variables

For an angular variable, the conventional mean and standard deviation do not provide a good characterization of its distribution, since the variable "wraps around" at $0, 2\pi$, etc. Therefore, we must take a different approach to computing

credible intervals for angular functions of $\mathbf{x}$. Here, we present a novel formula for angular credible intervals, combining directional statistics with the Markov inequality.

Let θ be an angular function of $\mathbf{x}$. We can approximate the expected values of $\cos\theta$ and $\sin\theta$, denoted by $\bar{c}$ and $\bar{s}$, using the unscented transform:

$$\bar{c} = \sum_{j=1}^{N} w_j \cos\theta(\mathbf{x}^{(j)}), \tag{1.76}$$

$$\bar{s} = \sum_{j=1}^{N} w_j \sin\theta(\mathbf{x}^{(j)}). \tag{1.77}$$

From these values, we can obtain the mean direction $\bar{\theta}$ and and mean resultant length $\bar{R}$ of θ, which are defined as [20]

$$\bar{\theta} = \operatorname{atan2}(\bar{s}, \bar{c}), \tag{1.78}$$

$$\bar{R} = \sqrt{\bar{c}^2 + \bar{s}^2}. \tag{1.79}$$

Now, we want to find a credible interval for θ about $\bar{\theta}$. We define

$$\beta = \frac{1}{2}(\theta - \bar{\theta}), \tag{1.80}$$

and using trigonometric identities, we obtain

$$\sin^2\beta = \frac{1}{2}(1 - \cos\theta\cos\bar{\theta} - \sin\theta\sin\bar{\theta}). \tag{1.81}$$

Taking the expected value of this expression gives

$$\mathrm{E}[\sin^2\beta] = \frac{1}{2}(1 - \bar{c}\cos\bar{\theta} - \bar{s}\sin\bar{\theta}), \tag{1.82}$$

and substituting (1.78–1.79), we have

$$\mathrm{E}[\sin^2\beta] = \frac{1}{2}(1 - \bar{R}). \tag{1.83}$$

Then, we can apply the Markov inequality (1.72) to obtain

$$\Pr[\sin^2\beta \geq \varepsilon] \leq \frac{1-\bar{R}}{2\varepsilon} \tag{1.84}$$

for any $\varepsilon > 0$, and it follows that

$$\Pr\left[\sin^2\beta \geq \frac{1-\bar{R}}{2\alpha}\right] \leq \alpha \tag{1.85}$$

for any $0 < \alpha < 1$. From this, we can obtain a credible interval for θ

$$C_\theta = \left[\bar{\theta} - 2\arcsin\sqrt{\frac{1-\bar{R}}{2\alpha}},\ \bar{\theta} + 2\arcsin\sqrt{\frac{1-\bar{R}}{2\alpha}}\right] \tag{1.86}$$

with probability $1-\alpha$. Note that, if $1-\bar{R} > 2\alpha$, the inverse sines cannot be evaluated, and C_θ covers the entire circle. This means that there is a high probability that the actual value of θ is far from $\bar{\theta}$.

1.5 Conclusion

In this chapter, we presented the theory of Bayesian estimation leading up to the unscented Kalman filter and some of its extensions. In addition, we derived a new formula (1.86) for credible intervals for angular variables based on the Markov inequality. We use this theory in the development and application of estimation methods throughout the rest of this thesis.

CHAPTER 2

THE HIGHER-ORDER UNSCENTED ESTIMATOR

2.1 Introduction

In recent years, a wide range of filters based on the unscented Kalman filter (UKF) have been developed. Some have targeted specific applications, such as the UnScented QUaternion Estimator (USQUE) proposed by Crassidis and Markley for quaternion-based attitude estimation [15]. Others have sought to generalize the UKF, particularly by adding estimates of higher moments and higher-order approximations in order to better characterize the state distribution. One straightforward generalization of the UKF is the higher order unscented filter (HOUF) developed by Tenne and Singh [75]. In HOUF, the sigma points and weights are chosen to match the moments up to a given order; specifically, the points and weights are the solution to a system of polynomial equations with the moments, which is typically obtained using an analytical solver.

The unscented transform is closely related to multivariate quadrature—or cubature—rules, which approximate multivariate integrals using weighted sums. A cubature rule in $\mathbb{R}^n$ is called a product rule if it is an n-fold tensor product of univariate quadrature rules. For filtering purposes, n is the dimension of the augmented state, i.e., the system state combined with measurement or process noise. Such cubature rules are simple to design but also suffer from the "curse of dimensionality" [74]. Several filtering methods have been developed using both product and non-product cubature rules. Ito and Xong [30] use product rules based on Gauss-Hermite quadratures in a filter for multivariate Gaussian distributions; this filter is further generalized to Gaus-

sian mixtures. Furthermore, Gauss-Laguerre quadrature rules, combined with spherical cubature rules, form the basis for the third-order cubature Kalman filter (CKF) proposed by Arasaratnam and Haykin [4] and the higher-degree CKF by Jia, Xin, and Cheng [32]. Another method, intended for cases where the noise is not linearly correlated with the state, is proposed by Grothe [26]. This method uses sets of points called generators, which are used for computing non-product cubature rules (e.g., [37]). With this, Grothe shows that the original sigma points are a special case of a generator and derives explicit formulas for a fifth-order unscented transform with $2n^2 + 1$ points and weights. A more recent filtering method using non-product cubature rules is based on the conjugate unscented transform (CUT), developed by Adurthi, Singla, and Singh [2, 3]. In this method, the sigma points are generated using scaled conjugate axes; the scales and weights for the points are chosen to match fourth, sixth, or eighth order moments for Gaussian or uniform distributions. The number of sigma points in CUT is $O(2^n)$.

While many filtering methods focus on increasing the accuracy of the state estimates, others aim to keep the computational cost of the filters low. The latter consideration is especially important in real-time applications, embedded systems, or wherever computation time and resources are limited. To this end, Julier and Uhlmann propose a reduced UKF with only $n + 2$ sigma points based on a simplex and chosen to match the mean and covariance and minimize the skewness [33]. An extension of this method with $2n + 3$ points, which matches moments up to the fifth order and minimizes the error for the sixth-order moments, is proposed by Lévesque [45].

It is interesting to note that none of the aforementioned filters explicitly

propagate higher-order moments, such as the skewness and kurtosis, of the state distribution. Rather, the higher moments of the standardized state—i.e., the state transformed so that its mean is zero and its covariance is the identity matrix—are taken to be known *a priori*. This is usually done by assuming a particular form for the state distribution, e.g., Gaussian or uniform. Such assumptions, however, are not always accurate. For example, for the problem of cross-calibration in constellations of imaging satellites, Shapiro [66] showed that measurements extracted from images exhibit highly non-Gaussian noise.

A filtering method that directly accounts for non-Gaussian and non-uniform distributions was proposed by Ponomareva, Date, and Wang [57]. This method retains the general structure of the UKF, with $2n+1$ sigma points, but the points and their corresponding weights are scaled using two factors that take into account third and fourth order moments. Specifically, the factors are chosen so that the averages of the third and fourth order marginal moments of the state components are preserved in the unscented transform. This method is simpler and more computationally efficient than several other higher-order filters in that it requires $O(n)$ sigma points and does not require solving large systems of polynomials using an analytical solver. Further results on this method can be found in [56].

Drawing on the method of Ponomareva et al., we propose a new extension of the unscented Kalman filter with $2n+1$ sigma points that propagates the state skewness and kurtosis in n directions, in addition to the mean and covariance. In this filter's variant of the unscented transform, the sigma points and their weights are asymmetrically scaled to match all of the propagated moments. We are able to obtain simple, explicit formulas for the sigma points and weights,

making their computation efficient and robust. This work was presented in [71]. We call this filter the Higher-Order UnScented Estimator (HOUSE).

In Section 2.2, we derive the formulas for the sigma points and their weights and describe constraints under which they are applicable. Finally, in Section 2.3 we demonstrate the use of HOUSE on three dynamical systems—an aircraft coordinated turn model, a rotating rigid body, and a projectile with drag—and compare its performance to that of a conventional UKF implementation as well as the CUT filters under various conditions.

2.2 Methodology

A straightforward generalization of the original unscented transform (1.41–1.42) was developed by Ponomareva, Date, and Wang [57], with sigma points

$$\mathbf{x}^{(j)} = \begin{cases} \bar{\mathbf{x}} + \alpha \mathbf{c}^{(j)}, & 1 \le j \le n \\ \bar{\mathbf{x}} - \beta \mathbf{c}^{(j-n)} & n+1 \le j \le 2n \\ \bar{\mathbf{x}} & j = 2n+1 \end{cases} \tag{2.1}$$

and weights

$$w_j = \begin{cases} \frac{1}{\alpha(\alpha+\beta)n}, & 1 \le j \le n \\ \frac{1}{\beta(\alpha+\beta)n}, & n+1 \le j \le 2n \\ 1 - \sum_{i=1}^{2n} w_i, & j = 2n+1 \end{cases} \tag{2.2}$$

where α and β are coefficients chosen to preserve the marginal third and fourth order moments averaged over the components of $\mathbf{x}$. We propose a further generalization of this method, with separate coefficients α_i and β_i for each column

of $\sqrt{\mathbf{P}_{xx}}$. Then, the sigma points are given by

$$\mathbf{x}^{(j)} = \begin{cases} \bar{\mathbf{x}} + \alpha_j \mathbf{c}^{(j)}, & 1 \leq j \leq n \\ \bar{\mathbf{x}} - \beta_{j-n} \mathbf{c}^{(j-n)} & n+1 \leq j \leq 2n \\ \bar{\mathbf{x}} & j = 2n+1 \end{cases} \tag{2.3}$$

and the corresponding cubature rule is

$$\int_{\Omega} \phi(\mathbf{x}) p(\mathbf{x}) \mathrm{d}\mathbf{x} = w_{2n+1} \phi(\bar{\mathbf{x}}) + \sum_{j=1}^{n} (w_j \phi(\bar{\mathbf{x}} + \alpha_j \mathbf{c}^{(j)}) + w_{n+j} \phi(\bar{\mathbf{x}} - \beta_j \mathbf{c}^{(j)})) \tag{2.4}$$

In this section, we will derive expressions for the coefficients α_i and β_i for $i = 1, \ldots, n$ along with the weights w_j for $j = 1, \ldots, 2n+1$ such that this unscented transform preserves the marginal skewness and kurtosis for each component of $\tilde{\mathbf{x}} = (\sqrt{\mathbf{P}_{xx}})^{-1}(\mathbf{x} - \bar{\mathbf{x}})$. Using these results, we will describe a new generalization of the unscented Kalman filter, which we call the Higher-Order UnScented Estimator (HOUSE). Furthermore, we will derive conditions under which HOUSE can reliably be applied.

2.2.1 Standardization

For any $\mathbf{x}$, we can define a random variable

$$\tilde{\mathbf{x}} = (\sqrt{\mathbf{P}_{xx}})^{-1}(\mathbf{x} - \bar{\mathbf{x}}) \tag{2.5}$$

with domain $\tilde{\Omega}$. Then, $\tilde{\mathbf{x}}$ has a mean of zero and covariance $\mathbf{I}$, the $n \times n$ identity matrix. Using a change of variables, we have

$$\int_{\Omega} \phi(\mathbf{x}) p(\mathbf{x}) \mathrm{d}\mathbf{x} = \int_{\tilde{\Omega}} \phi\left((\sqrt{\mathbf{P}_{xx}})\tilde{\mathbf{x}} + \bar{\mathbf{x}}\right) \tilde{p}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}} \tag{2.6}$$

where $\tilde{p}$ is the probability density of $\tilde{x}$. The corresponding cubature points for $\tilde{\mathbf{x}}$ are

$$\tilde{\mathbf{x}}^{(j)} = \begin{cases} \alpha_j \mathbf{e}^{(j)}, & 1 \leq j \leq n \\ -\beta_{j-n} \mathbf{e}^{(j-n)} & n+1 \leq j \leq 2n \\ \mathbf{0} & j = 2n+1 \end{cases} \tag{2.7}$$

where the $\mathbf{e}^{(j)}$ are the standard basis vectors, and the cubature rule is

$$\int_{\tilde{\Omega}} \psi(\tilde{\mathbf{x}}) \tilde{p}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}} = w_{2n+1} \psi(\mathbf{0}) + \sum_{j=1}^{n} (w_j \psi(\alpha_j \mathbf{e}^{(j)}) + w_{n+j} \psi(-\beta_j \mathbf{e}^{(j)})) \tag{2.8}$$

For $\phi : \Omega \to \mathbb{R}$, we define

$$\tilde{\phi}(\tilde{\mathbf{x}}) = \phi\left((\sqrt{\mathbf{P}_{xx}}) \tilde{\mathbf{x}} + \bar{\mathbf{x}} \right) \tag{2.9}$$

The cubature rule (2.4) for ϕ holds if and only if (2.8) holds for $\tilde{\phi}$. This can be shown by supposing that the cubature rule (2.8) holds for $\tilde{\phi}$. Then

$$\begin{aligned} \int_{\Omega} \phi(\mathbf{x}) p(\mathbf{x}) \mathrm{d}\mathbf{x} &= \int_{\tilde{\Omega}} \phi\left((\sqrt{\mathbf{P}_{xx}}) \tilde{\mathbf{x}} + \bar{\mathbf{x}} \right) \tilde{p}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}} = \int_{\tilde{\Omega}} \tilde{\phi}(\tilde{\mathbf{x}}) \tilde{p}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}} \\ &= w_{2n+1} \tilde{\phi}(\mathbf{0}) + \sum_{j=1}^{n} (w_j \tilde{\phi}(\alpha_j \mathbf{e}^{(j)}) + w_{n+j} \tilde{\phi}(-\beta_j \mathbf{e}^{(j)})) \\ &= w_{2n+1} \phi(\bar{\mathbf{x}}) + \sum_{j=1}^{n} (w_j \phi(\bar{\mathbf{x}} + \alpha_j \mathbf{c}^{(j)}) + w_{n+j} \phi(\bar{\mathbf{x}} - \beta_j \mathbf{c}^{(j)})) \end{aligned} \tag{2.10}$$

using $\sqrt{\mathbf{P}_{xx}} \mathbf{e}^{(j)} = \mathbf{c}^{(j)}$. Conversely, if (2.4) holds for ϕ, we have

$$\begin{aligned} \int_{\tilde{\Omega}} \tilde{\phi}(\tilde{\mathbf{x}}) \tilde{p}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}} &= \int_{\tilde{\Omega}} \phi\left((\sqrt{\mathbf{P}_{xx}}) \tilde{\mathbf{x}} + \bar{\mathbf{x}} \right) \tilde{p}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}} = \int_{\Omega} \phi(\mathbf{x}) p(\mathbf{x}) \mathrm{d}\mathbf{x} \\ &= w_{2n+1} \phi(\bar{\mathbf{x}}) + \sum_{j=1}^{n} (w_j \phi(\bar{\mathbf{x}} + \alpha_j \mathbf{c}^{(j)}) + w_{n+j} \phi(\bar{\mathbf{x}} - \beta_j \mathbf{c}^{(j)})) \\ &= w_{2n+1} \tilde{\phi}(\mathbf{0}) + \sum_{j=1}^{n} (w_j \tilde{\phi}(\alpha_j \mathbf{e}^{(j)}) + w_{n+j} \tilde{\phi}(-\beta_j \mathbf{e}^{(j)})) \end{aligned} \tag{2.11}$$

demonstrating the bi-directionality of the realtionship between the two cubature rules. Therefore, our problem reduces to finding the coefficients α_j, β_j and the weights w_j for a random, zero-mean vector with covariance $\mathbf{I}$.

2.2.2 Determination of Coefficients and Weights

Suppose that we want the cubature rule (2.8) to hold exactly for $\psi(\tilde{\mathbf{x}}) = 1$ and $\psi(\tilde{\mathbf{x}}) = x_i^k$ with $i = 1, \ldots, n$ and $k = 1, 2, 3, 4$. If $\tilde{\mathbf{x}}_i$ has skewness γ_i and kurtosis κ_i, then we have the system of equations

$$\sum_{j=1}^{2n+1} w_j = 1 \tag{2.12}$$

and

$$w_i\alpha_i - w_{n+i}\beta_i = \quad 0 \tag{2.13}$$

$$w_i\alpha_i^2 + w_{n+i}\beta_i^2 = \quad 1 \tag{2.14}$$

$$w_i\alpha_i^3 - w_{n+i}\beta_i^3 = \quad \gamma_i \tag{2.15}$$

$$w_i\alpha_i^4 + w_{n+i}\beta_i^4 = \quad \kappa_i \tag{2.16}$$

in unknowns α_i, β_i, w_i, w_{n+i}, and w_{2n+1} for $i = 1, \ldots, n$. This system can be solved as follows. First, from (2.13), we have

$$w_{n+i} = \frac{\alpha_i}{\beta_i} w_i \tag{2.17}$$

Substituting this into (2.14) gives

$$w_i = \frac{1}{\alpha_i^2 + \alpha_i\beta_i} \tag{2.18}$$

and applying (2.17) gives

$$w_{n+i} = \frac{1}{\beta_i^2 + \alpha_i\beta_i} \tag{2.19}$$

If we substitute these two results into the higher-moment equations (2.15) and (2.16), we have

$$\frac{\alpha_i^3}{\alpha_i^2 + \alpha_i\beta_i} - \frac{\beta_i^3}{\beta_i^2 + \alpha_i\beta_i} = \gamma_i \tag{2.20}$$

and

$$\frac{\alpha_i^4}{\alpha_i^2 + \alpha_i\beta_i} + \frac{\beta_i^4}{\beta_i^2 + \alpha_i\beta_i} = \kappa_i \tag{2.21}$$

which simplify to

$$\alpha_i - \beta_i = \gamma_i \tag{2.22}$$

and

$$\alpha_i^2 - \alpha_i\beta_i - \beta_i^2 = \kappa_i \tag{2.23}$$

respectively. Combining the last two expressions yields the second order polynomial

$$\alpha_i^2 - \gamma_i\alpha_i + \gamma_i^2 - \kappa_i = 0 \tag{2.24}$$

which has the solution

$$\alpha_i = \frac{\gamma_i \pm \sqrt{4\kappa_i - 3\gamma_i^2}}{2} \tag{2.25}$$

Since the skewness and kurtosis always satisfy the inequality

$$\kappa_i \geq \gamma_i^2 + 1 \tag{2.26}$$

we have

$$4\kappa_i - 3\gamma_i^2 \geq \gamma_i^2 + 4 \tag{2.27}$$

which guarantees that there is a real, positive solution for α_i, namely

$$\alpha_i = \frac{\gamma_i + \sqrt{4\kappa_i - 3\gamma_i^2}}{2} \tag{2.28}$$

Then, using (2.22), we have

$$\beta_i = \frac{-\gamma_i + \sqrt{4\kappa_i - 3\gamma_i^2}}{2} \tag{2.29}$$

which is also real and positive. Furthermore, we can evaluate $w_1, \ldots, w_{2n}$ using (2.18) and (2.19), and we find that these weights are positive as well. Finally, from (2.12), we have

$$w_{2n+1} = 1 - \sum_{j=1}^{2n} w_j \tag{2.30}$$

2.2.3 Tuning the Sigma Points and Weights

From (2.18) and (2.19), we have

$$w_i + w_{n+i} = \frac{1}{\alpha_i \beta_i} \tag{2.31}$$

and substituting this into (2.30) gives

$$w_{2n+1} = 1 - \sum_{i=1}^{n} \frac{1}{\alpha_i \beta_i} \tag{2.32}$$

Furthermore, substituting (2.28) and (2.29) gives

$$w_{2n+1} = 1 - \sum_{i=1}^{n} \frac{1}{\kappa_i - \gamma_i^2} \tag{2.33}$$

which holds regardless of the sign chosen in the solution for α_i. From the inequality (2.26), we have

$$0 < \frac{1}{\kappa_i - \gamma_i^2} \leq 1 \tag{2.34}$$

However, this does not guarantee that w_{2n+1} will be positive, especially if n is large and the κ_i and γ_i are small. In some cases, the fact that w_{2n+1} can be negative can be problematic. For example, approximating the covariance of a random variable $\mathbf{y} = \mathbf{g}(\mathbf{x})$ by

$$\mathbf{P}_{yy} = \sum_{j=1}^{2n+1} w_j (\mathbf{y}^{(j)} - \bar{\mathbf{y}})(\mathbf{y}^{(j)} - \bar{\mathbf{y}})^{\mathrm{T}} \tag{2.35}$$

could give a covariance matrix estimate that is not positive-definite. On the other hand, if all of the weights are non-negative, then the approximated covariance matrix must be at least positive-semidefinite.

One way to ensure that w_{2n+1} stay non-negative is to increase the smaller values of kurtosis, so that $\kappa_i - \gamma_i^2$ is at least $\kappa_{\min}$ in each direction

$$\kappa_i' = \begin{cases} \kappa_{\min} + \gamma_i^2, & \kappa_i - \gamma_i^2 \leq \kappa_{\min} \\ \kappa_i, & \text{otherwise} \end{cases} \tag{2.36}$$

We choose $\kappa_{\min}$ to be the value of kurtosis that gives $w_{2n+1} = \delta$ for a symmetric distribution ($\delta = 0$ in the minimal case); that is

$$\kappa_{\min} = \frac{n}{1-\delta} \tag{2.37}$$

The trade-off of this approach, of course, is that the higher moments are somewhat distorted. Specifically, since the kurtosis increases, the modified distribution has heavier tails and is more non-Gaussian.

2.2.4 Mixed Moments

Let Q be the cubature operator from (2.8)

$$Q[\psi] = w_{2n+1}\psi(\mathbf{0}) + \sum_{j=1}^{n}(w_j\psi(\alpha_j\mathbf{e}^{(j)}) + w_{n+j}\psi(-\beta_j\mathbf{e}^{(j)})) \tag{2.38}$$

If

$$\psi(\tilde{\mathbf{x}}) = \prod_{i=1}^{n}\tilde{x}_i^{k_i} \tag{2.39}$$

where two or more of the k_i are nonzero, then $Q[\psi] = 0$. Therefore, using the coefficients and weights found in Section 2.2.2, the cubature rule (2.8) is exact for $\psi(\tilde{\mathbf{x}}) = \tilde{x}_i\tilde{x}_j$, since the components of $\tilde{\mathbf{x}}$ are uncorrelated. Effectively, when computing $Q[\psi]$ for an arbitrary function ψ, the cubature "drops" the third and fourth order mixed terms in the Taylor expansion of ψ. We do not expect that this significantly impacts the accuracy of the filter.

In the special case where the components of $\tilde{x}$ are not only uncorrelated, but independent, the quadrature rule is also exact for all third and fourth order monomials except those of the form $\tilde{x}_i^2\tilde{x}_j^2$. This assumption of independence is reasonable if the system state is described using a minimal set of coordinates.

2.2.5 Constraints

The main limitation of the HOUSE filter, particularly if the correction (2.36) is applied, is that the sigma points could grow without bound for large values of kurtosis. One way to ensure that this does not occur is to impose the constraint

$$\|\mathbf{x}^{(i)} - \bar{\mathbf{x}}\| < R \tag{2.40}$$

for some radius R. Using (2.3), we see that this constraint is equivalent to

$$\max_{i \in \{1,\dots,n\}} \left(\|\mathbf{c}^{(i)}\| \max(\alpha_i, \beta_i) \right) < R \tag{2.41}$$

Furthermore, from (2.28) and (2.29), we have

$$\max(\alpha_i, \beta_i) = \frac{|\gamma_i| + \sqrt{4\kappa_i - 3\gamma_i^2}}{2} \tag{2.42}$$

so the constraint (2.40) can be expressed as

$$\max_{i \in \{1,\dots,n\}} \left(\frac{|\gamma_i| + \sqrt{4\kappa_i - 3\gamma_i^2}}{2} \|\mathbf{c}^{(i)}\| \right) < R \tag{2.43}$$

The other constraint on the moments comes from (2.33): if all weights are to be non-negative, then the skewness and kurtosis values must satisfy

$$\sum_{i=1}^{n} \frac{1}{\kappa_i - \gamma_i^2} \leq 1 \tag{2.44}$$

To obtain simpler inequalities, we assume that $\mathbf{x}$ has a radially symmetric distribution about $\bar{\mathbf{x}}$. This is effectively assuming that we have the same uncertainty in each component of $\mathbf{x}$; while this is clearly not the case in the filter, we believe that this is reasonable for an order-of-magnitude analysis of the filter's tractability. Then, we have $\kappa_i = \kappa$, $\gamma_i = 0$, and $\|\mathbf{c}^{(i)}\| = \sigma$ for $i = 1, \dots, 2n + 1$; in that case, the inequality (2.43) reduces to

$$\sigma \sqrt{\kappa} < R \tag{2.45}$$

and (2.44) reduces to

$$\frac{n}{\kappa} \leq 1 \tag{2.46}$$

These two simple inequalities provide an approximate range for κ where the filter can operate

$$n \leq \kappa < \frac{R^2}{\sigma^2} \tag{2.47}$$

We expect that the filter will perform better if the range for κ is wider; this is satisfied when

$$n\sigma^2 \ll R^2 \tag{2.48}$$

Furthermore, because of this inequality, we expect that this filter is better suited for lower-dimensional systems.

2.2.6 Computational Complexity

The computational complexity of HOUSE is only slightly greater than that of the conventional UKF. In both filters, the most expensive part of the computation tends to be the evaluation of the nonlinear state and measurement functions, and in both filters these are evaluated $2n + 1$ times. The added complexity of HOUSE is due to the evaluation of the sigma point coefficients and weights, which requires $O(n)$ operations, and the estimation of the skewness and kurtosis, which requires $O(n^2)$ operations. However, these computations are very simple and involve no transcendental functions. The increased memory requirement for HOUSE due to the stored skewness and kurtosis is also $O(n)$.

2.2.7 Procedures

Here, we summarize the HOUSE prediction and correction procedures. These are illustrated in Figures 2.1 and 2.2. A C++ implementation of these procedures can be found in [72].

Prediction

1. For the augmented state

$$\mathbf{y}_{\mathrm{P}}(k) = \begin{bmatrix} \mathbf{x}(k) \\ \mathbf{w}(k) \end{bmatrix} \tag{2.49}$$

generate the modified sigma points

$$\mathbf{y}_{\mathrm{P}}^{(j)}(k) = \begin{bmatrix} \mathbf{x}^{(j)}(k) \\ \mathbf{w}^{(j)}(k) \end{bmatrix} \tag{2.50}$$

and weights w_j, as described in Section 2.2.2.

2. Propagate the state for each sigma point:

$$\mathbf{x}^{(j)}(k+1|k) = \mathbf{f}(\mathbf{x}^{(j)}(k), \mathbf{w}^{(j)}(k), k) \tag{2.51}$$

3. Compute the predicted mean and covariance:

$$\bar{\mathbf{x}}(k+1|k) = \sum_{j=1}^{N} w_j \mathbf{x}^{(j)}(k+1|k) \tag{2.52}$$

$$\boldsymbol{\eta}^{(j)}(k+1|k) = \mathbf{x}^{(j)}(k+1|k) - \bar{\mathbf{x}}(k+1|k) \tag{2.53}$$

$$\mathbf{P}_{xx}(k+1|k) = \sum_{j=1}^{N} w_j \boldsymbol{\eta}^{(j)}(k+1|k) \boldsymbol{\eta}^{(j)}(k+1|k)^{\mathrm{T}} \tag{2.54}$$

4. Compute the standardized states at the sigma points:

$$\tilde{\mathbf{x}}^{(j)}(k+1|k) = \left(\sqrt{\mathbf{P}_{xx}(k+1|k)} \right)^{-1} \boldsymbol{\eta}^{(j)}(k+1|k) \tag{2.55}$$

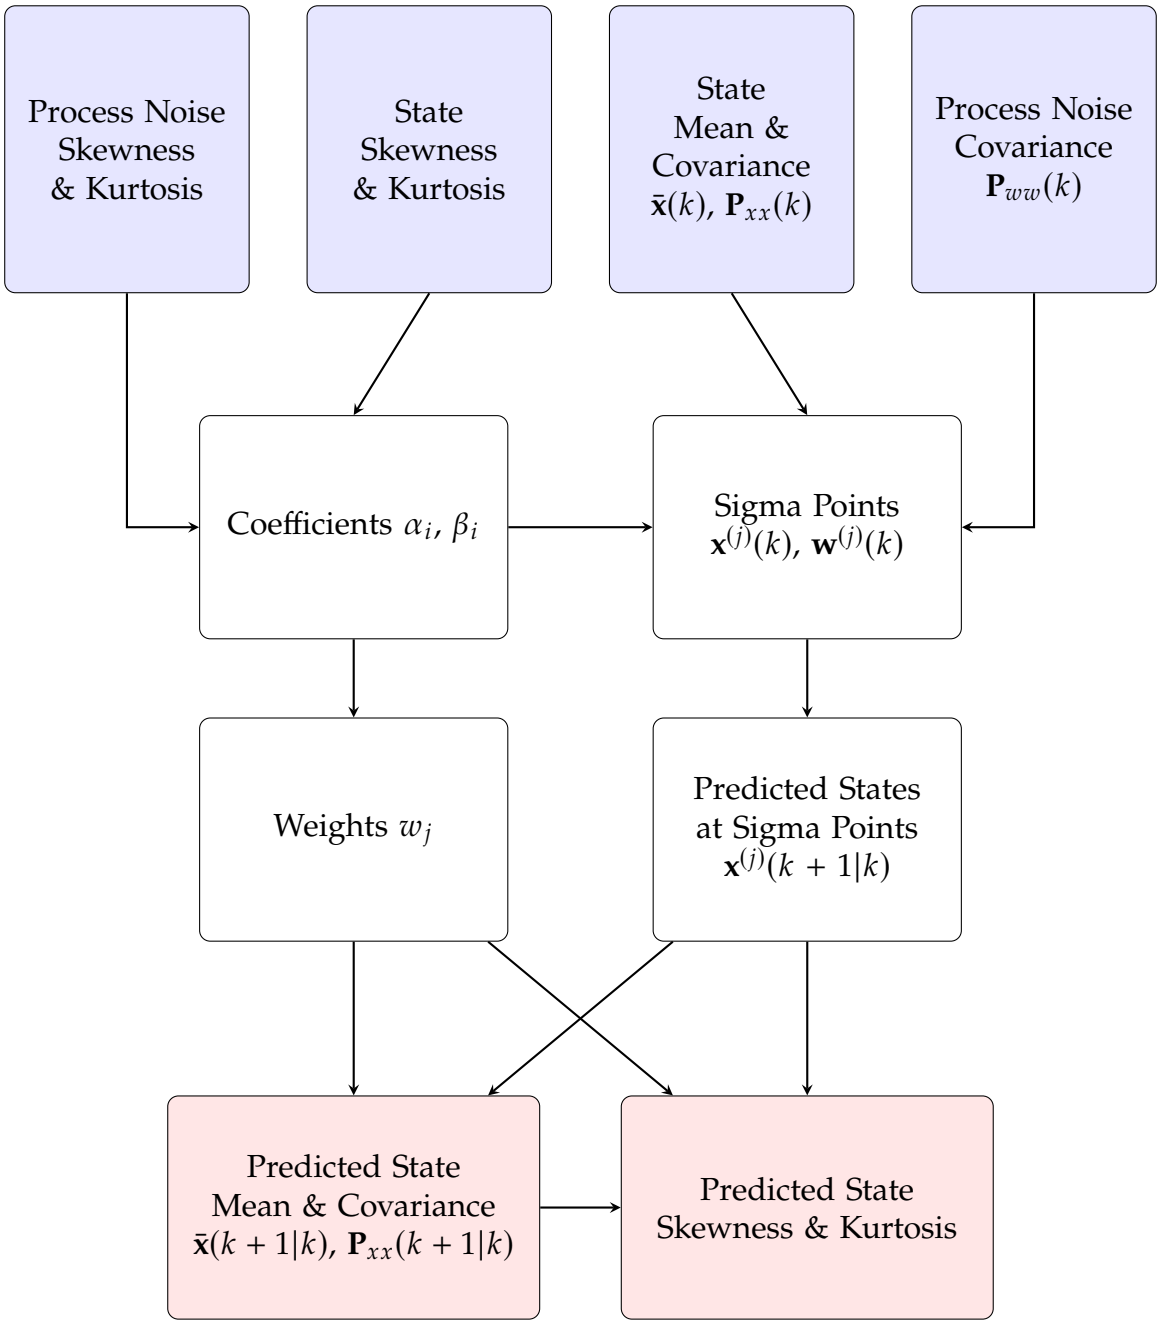

Figure 2.1: Diagram of HOUSE prediction.

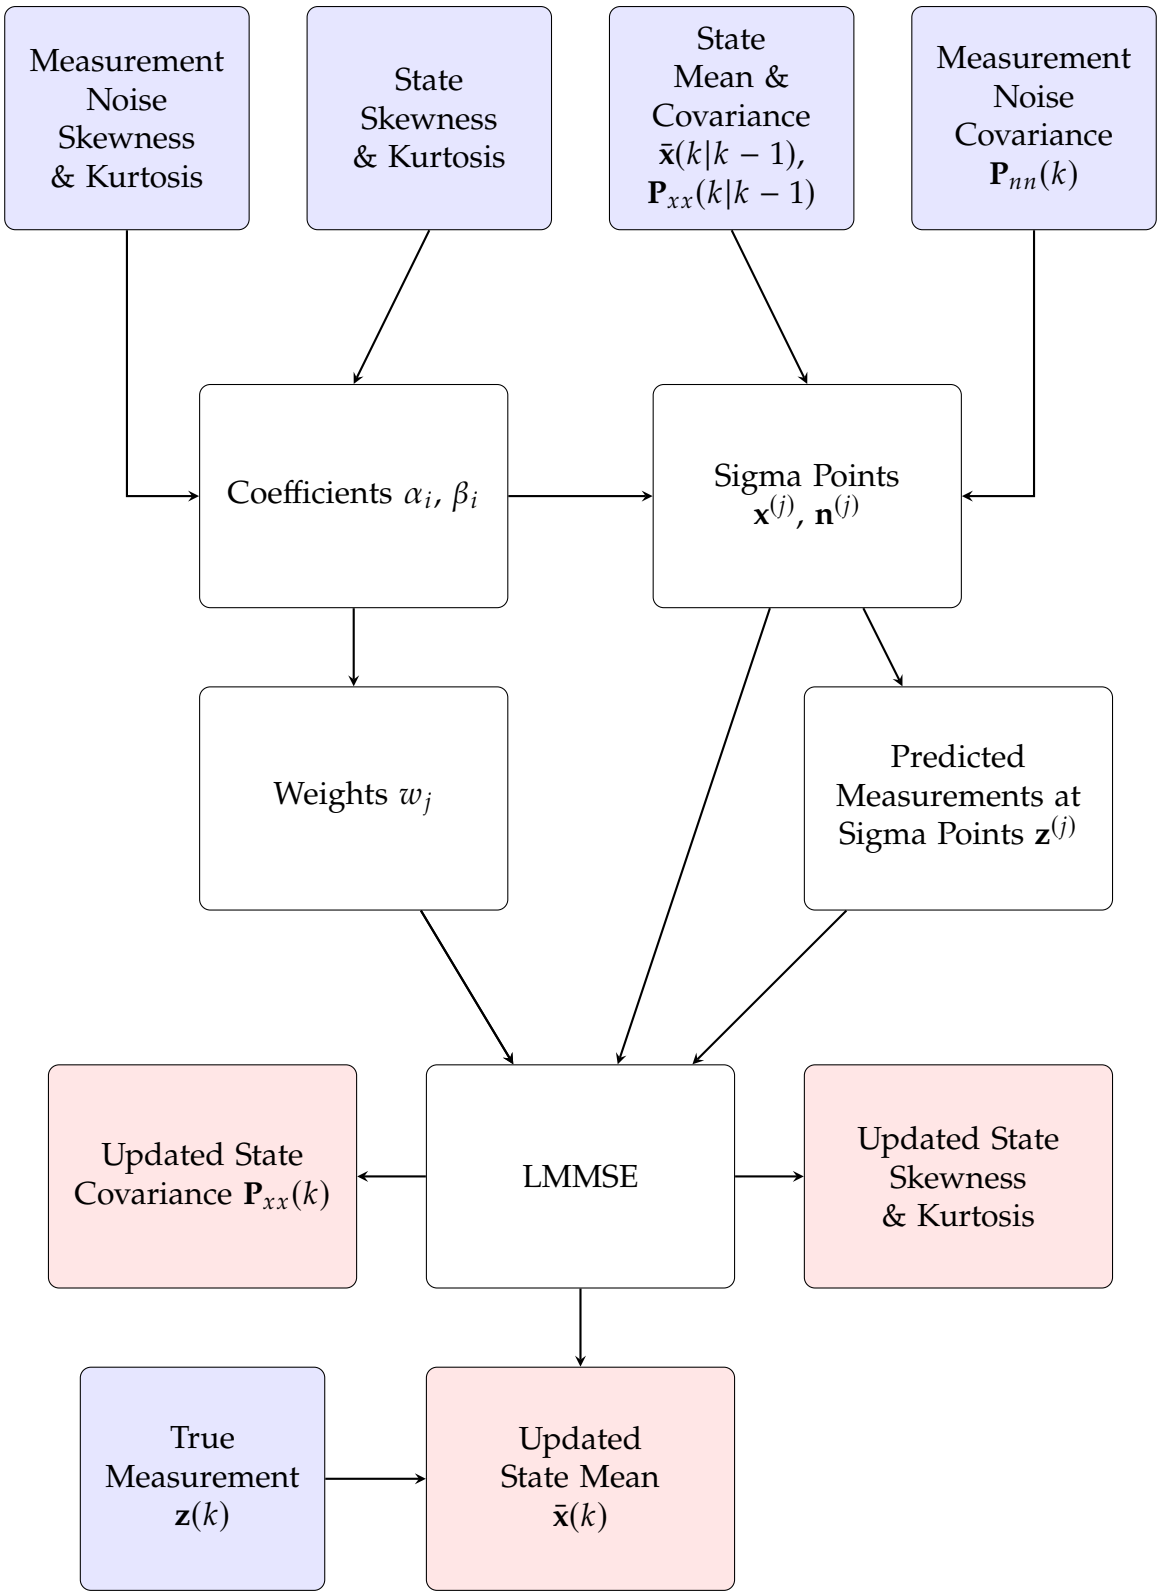

Figure 2.2: Diagram of HOUSE update.

5. Compute the skewness and kurtosis of the standardized state:

$$\gamma_i(k+1|k) = \sum_{i=1}^{N} w_j \tilde{x}_i^{(j)}(k+1|k)^3 \tag{2.56}$$

$$\kappa_i(k+1|k) = \sum_{i=1}^{N} w_j \tilde{x}_i^{(j)}(k+1|k)^4 \tag{2.57}$$

Correction

Here, $\mathbf{z}(k)$ denotes the true measurement.

1. For the augmented state

$$\mathbf{y}_{\mathrm{C}}(k) = \begin{bmatrix} \mathbf{x}(k) \\ \mathbf{n}(k) \end{bmatrix} \tag{2.58}$$

generate the modified sigma points

$$\mathbf{y}_{\mathrm{C}}^{(j)}(k) = \begin{bmatrix} \mathbf{x}^{(j)}(k) \\ \mathbf{n}^{(j)}(k) \end{bmatrix} \tag{2.59}$$

and weights w_j, as described in Section 2.2.2.

2. Compute the measurement for each sigma point:

$$\mathbf{z}^{(j)}(k) = \mathbf{h}(\mathbf{x}^{(j)}(k), \mathbf{w}^{(j)}(k), k) \tag{2.60}$$

3. Compute the measurement mean and covariance:

$$\bar{\mathbf{z}}(k) = \sum_{j=1}^{N} w_j \mathbf{z}^{(j)}(k) \tag{2.61}$$

$$\mathbf{P}_{zz}(k) = \sum_{j=1}^{N} w_j (\mathbf{z}^{(j)}(k) - \bar{\mathbf{z}}(k))(\mathbf{z}^{(j)}(k) - \bar{\mathbf{z}}(k))^{\mathrm{T}} \tag{2.62}$$

$$\mathbf{P}_{xz}(k) = \sum_{j=1}^{N} w_j (\mathbf{x}^{(j)}(k|k-1) - \bar{\mathbf{x}}(k|k-1))(\mathbf{z}^{(j)}(k) - \bar{\mathbf{z}}(k))^{\mathrm{T}} \tag{2.63}$$

4. Compute the LMMSE-updated mean and covariance for the state:

$$\bar{\mathbf{x}}(k) = \bar{\mathbf{x}}(k|k-1) + \mathbf{P}_{xz}(k)\mathbf{P}_{zz}(k)^{-1}(\mathbf{z}(k) - \bar{\mathbf{z}}(k)) \tag{2.64}$$

$$\mathbf{P}_{xx}(k) = \mathbf{P}_{xx}(k|k-1) - \mathbf{P}_{xz}(k)\mathbf{P}_{zz}(k)^{-1}\mathbf{P}_{zx}(k)^{\mathrm{T}} \tag{2.65}$$

5. Compute the LMMSE error and standardized state at the sigma points:

$$\boldsymbol{\epsilon}^{(j)}(k) = \mathbf{x}^{(j)}(k) - \bar{\mathbf{x}}(k|k-1) - \mathbf{P}_{xz}(k)\mathbf{P}_{zz}(k)^{-1}(\mathbf{z}^{(j)}(k) - \bar{\mathbf{z}}(k)) \tag{2.66}$$

$$\tilde{\mathbf{x}}^{(j)} = \left(\sqrt{\mathbf{P}_{xx}(k)}\right)^{-1} \boldsymbol{\epsilon}^{(j)}(k) \tag{2.67}$$

6. Compute the skewness and kurtosis of the standardized state:

$$\gamma_i(k) = \sum_{i=1}^{N} w_j \tilde{x}_i^{(j)}(k)^3 \tag{2.68}$$

$$\kappa_i(k) = \sum_{i=1}^{N} w_j \tilde{x}_i^{(j)}(k)^4 \tag{2.69}$$

2.3 Examples of Application

In this section, we demonstrate the use of HOUSE with simulations of three dynamical systems: an aircraft performing coordinated turns, a rotating rigid body, and a projectile with drag. In each case, we test the performance of HOUSE against the conventional UKF and the CUT filters with Gaussian and non-Gaussian noise. Potential advantages of CUT include guaranteed positive weights without moment distortion and accounting for higher-order mixed moments [3].

In the non-Gaussian case, the noise is sampled from a Pearson type IV distribution, which has a PDF of the form

$$p(x) = K\left(1 + \left(\frac{x-\lambda}{a}\right)^2\right)^{-m} \exp\left(-\nu \arctan\left(\frac{x-\lambda}{a}\right)\right) \tag{2.70}$$

defined on the entire real line, where $a > 0$, $m > 1/2$, λ, and ν are free parameters, and K is a normalizing factor. Given values of the mean, standard deviation, skewness, and kurtosis, the parameters of the Pearson type IV distribution can be determined uniquely, provided that [18]

$$0 < \frac{\gamma^2(\kappa+3)^2}{4(4\kappa-3\gamma^2)(2\kappa-3\gamma^2-6)} < 1 \tag{2.71}$$

Qualitatively, the Pearson type IV distribution has an asymmetric bell shape with heavy tails. It has been used for modeling various random processes, including wind shear fluctuations [61], fluctuating pressure on aircraft skin panels [68], and solar wind intensity [41]. The random noise in our simulations is generated using the procedures described by McGrath and Irving [50].

We implemented HOUSE, the conventional UKF, and the CUT filters in C++, using similar program structures to ensure a fair comparison of run times. The CUT sigma points for standardized Gaussian distributions were pre-computed using code by Adurthi et al. [1]. We used the Eigen 3 library [27] for matrix operations and Burkardt's implementation [12] of the Shampine-Gordon solver [65] for ordinary differential equations. The complete filtering and simulation code can be found in [72].

2.3.1 Aircraft Coordinated Turn

This example, representing a simple air traffic tracking scenario, is used by Adurthi et al. [3] to test the performance of the CUT filters. In this scenario, an aircraft executes the following maneuvers at a constant speed of 120 m/s: heads westward for 125 s, turns southward 90° with a turn rate of 1°/s, heads southward for 125 s, turns westward 90° at −3°/s, and then heads westward for

125 s.

The aircraft's motion is described by the coordinated turn (CT) model. The state vector is $\mathbf{x} = \begin{bmatrix} \xi & \dot{\xi} & \eta & \dot{\eta} & \Omega \end{bmatrix}^{\mathrm{T}}$, where (ξ, η) is the aircraft's position and Ω is its turn rate. The state dynamics are modeled in discrete time as

$$\mathbf{x}(k+1) = \begin{bmatrix} 1 & \frac{\sin(\Omega T)}{\Omega} & 0 & -\frac{1-\cos(\Omega T)}{\Omega} & 0 \\ 0 & \cos(\Omega T) & 0 & -\sin(\Omega T) & 0 \\ 0 & \frac{1-\cos(\Omega T)}{\Omega} & 1 & \frac{\sin(\Omega T)}{\Omega} & 0 \\ 0 & \sin(\Omega T) & 0 & \cos(\Omega T) & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \mathbf{x}(k) + \mathbf{w}(k) \tag{2.72}$$

where T is the time between steps. The covariance matrix of the noise $\mathbf{w}$ is given by

$$\mathbf{P}_{ww} = L_1 \begin{bmatrix} \frac{T^3}{3} & \frac{T^2}{2} & 0 & 0 & 0 \\ \frac{T^2}{2} & T & 0 & 0 & 0 \\ 0 & 0 & \frac{T^3}{3} & \frac{T^2}{2} & 0 \\ 0 & 0 & \frac{T^2}{2} & T & 0 \\ 0 & 0 & 0 & 0 & \frac{L_2}{L_1} T \end{bmatrix} \tag{2.73}$$

where L_1 and L_2 are constants. A radar takes takes measurements of range

$$r = \sqrt{\xi^2 + \eta^2} + n_r \tag{2.74}$$

and bearing

$$\theta = \operatorname{atan2}(\eta, \xi) + n_\theta, \tag{2.75}$$

where n_r and n_θ represent measurement noise.

In the simulations, the process noise constants are $L_1 = 0.16$ and $L_2 = 0.01$; the standard deviations of the process noise n_r and n_θ are 100 m and 1°, respectively. The initial (prior) distribution of the state has mean $\bar{\mathbf{x}}(0) =$

$\begin{bmatrix} 25000\,\text{m} & -120\,\text{m} & 10000\,\text{m} & 0\,\text{m/s} & 0.000001\,\text{rad/s} \end{bmatrix}^{\mathrm{T}}$, with standard deviation 1000 m for the position, 10 m/s for the velocity, and 1°/s for the turn rate.

We consider two cases: one in which the measurement noise is Gaussian and another in which it has a Pearson Type IV distribution with skewness $\gamma = -1$ and kurtosis $\kappa = 20$. In addition, in the Pearson case, the initial state distribution and process noise are assumed to have kurtosis $\kappa = 10$. For each case, we test the performance of HOUSE, the UKF (with $\kappa = 1$), CUT-4, CUT-6, and CUT-8 over 100 trials.

Figure 2.3 shows the root-mean-square error (RMSE) for each filter and length of time between measurements. The RMSE values for the UKF and CUT filters with Gaussian noise are nearly identical to those in [3]. In the case with Gaussian noise, the RMSE is of the same order of magnitude for HOUSE and UKF but lower by an order of magnitude or more for CUT-4, CUT-6, and CUT-8. This is expected, because the CUT filters are designed specifically to match the higher-order moments of Gaussian distributions.

For the case with Pearson Type IV noise, Figure 2.3 shows that the RMSE for HOUSE is lower than for any of the other filters for $T = 1$ s. On the other hand, for larger values of T, the estimation error is lower for CUT-4, CUT-6, and CUT-8. This is due to one of the limitations of HOUSE discussed in Section 2.2.5: the distance of the sigma points from the mean can rapidly increase with the standard deviation and kurtosis of the state and noise. In this case, we know from (2.73) that the standard deviation of the process noise, and therefore the distance of the sigma points from the mean, increases with T; this can lead to unrealistic values for the state components, particularly the turn rate Ω, for large values of T. Modeling the process noise with a smaller kurtosis value

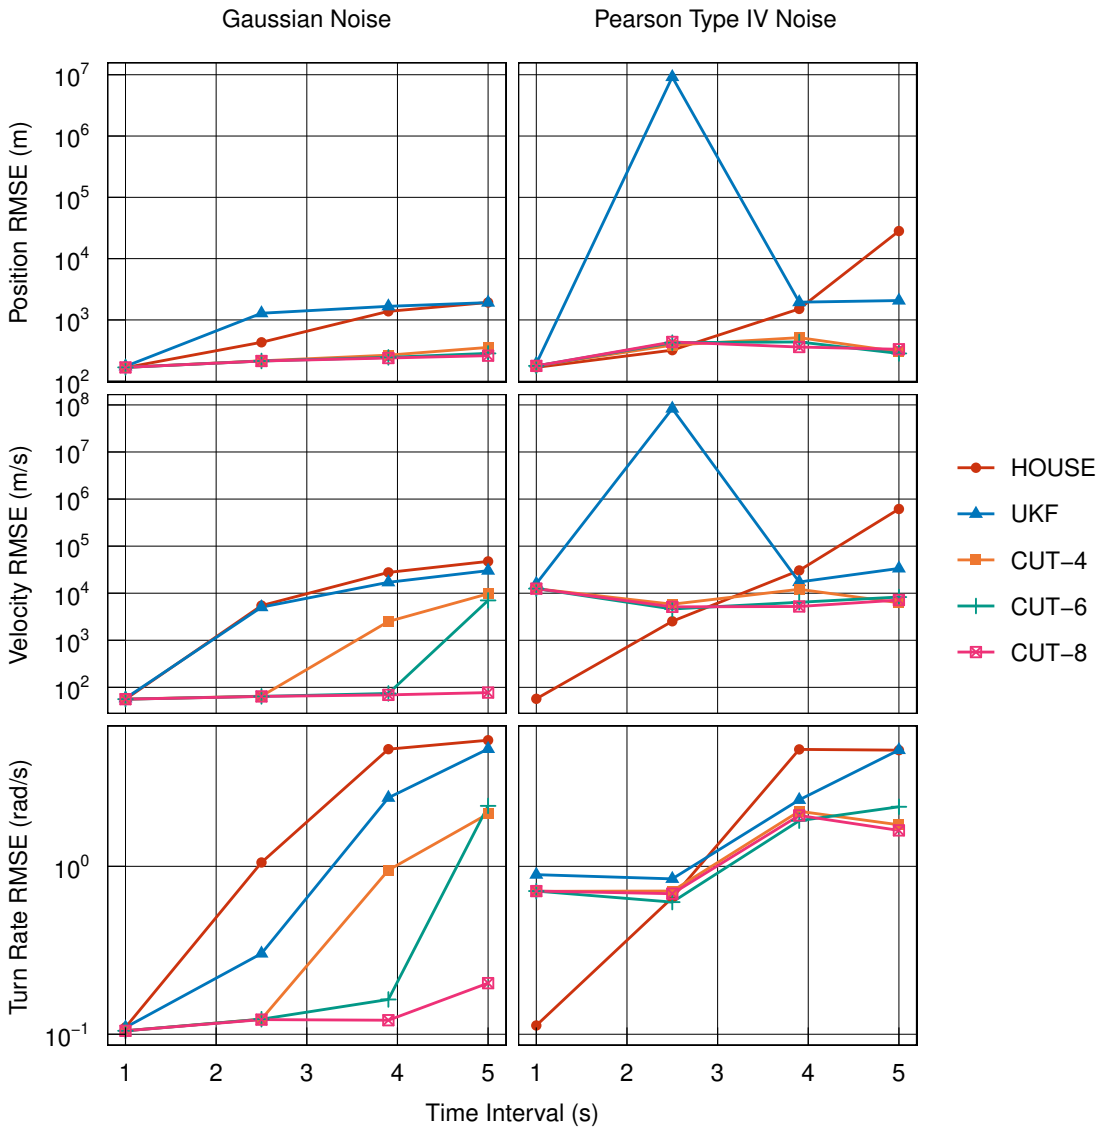

Figure 2.3: Root-mean-square error for aircraft CT example.

could mitigate this problem, but it could also produce negative weights: since the dimension of the augmented state is $n = 10$, the kurtosis value used here ($\kappa = 10$) is already the smallest value that guarantees non-negative weights, given by (2.37).

Figures 2.4 and 2.5 show the error distribution for $T = 1$ s and $T = 5$ s, respectively, for each of the filters. In Figure 2.4, we see that the RMSE for

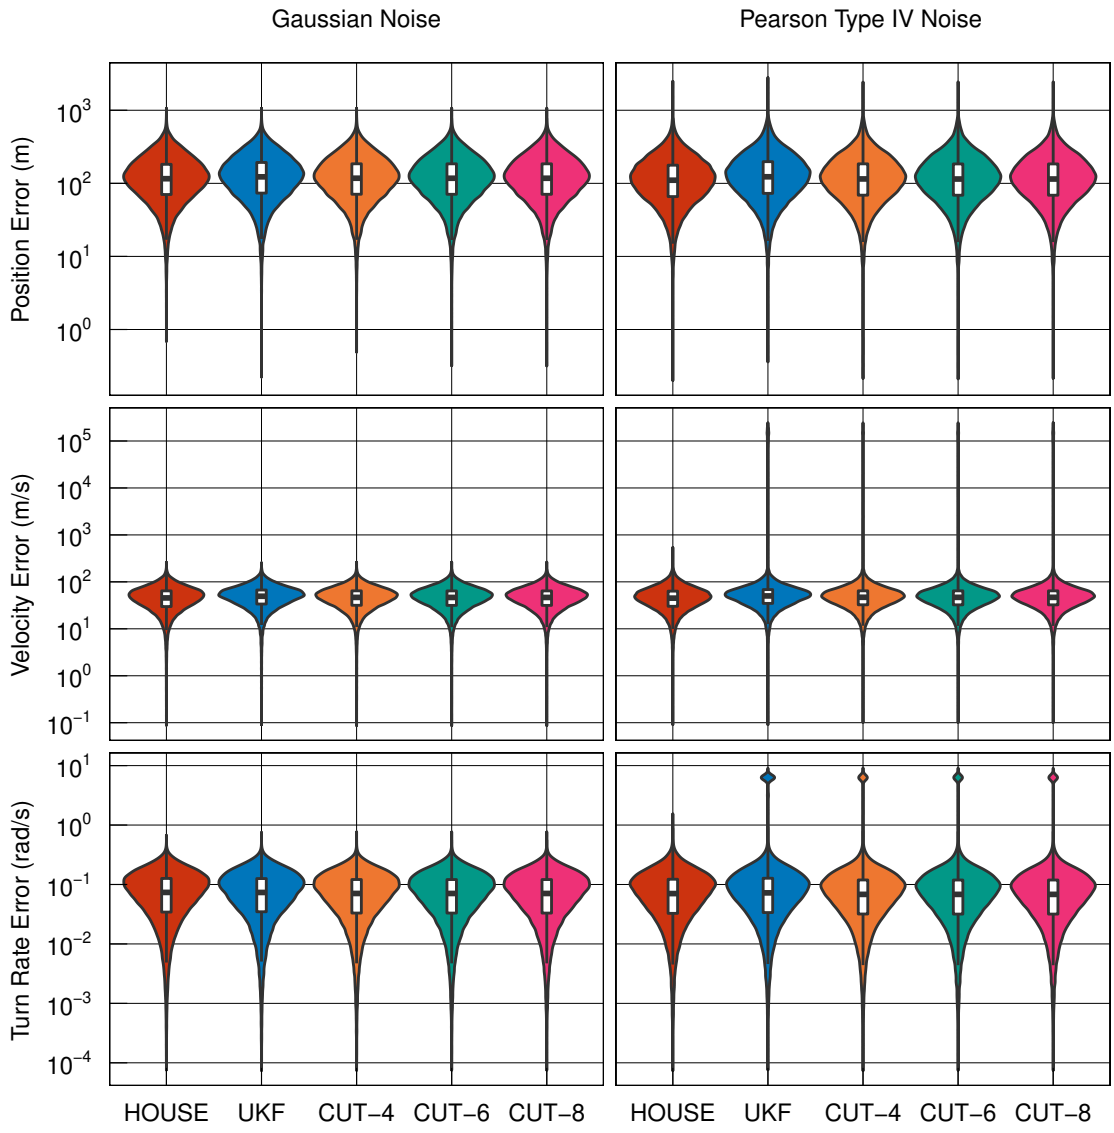

Figure 2.4: Distribution of estimation error for coordinated turn example with sampling time interval 1 s.

Pearson type IV noise is driven by outliers for all of the filters except HOUSE. This effect is closely related to the kurtosis of the noise distributions and is an important characteristic of HOUSE, as discussed in the next example. However, from Figure 2.5, we see how the performance of HOUSE degrades for a larger value of T; in that case, HOUSE produces even more outliers than the other filters.

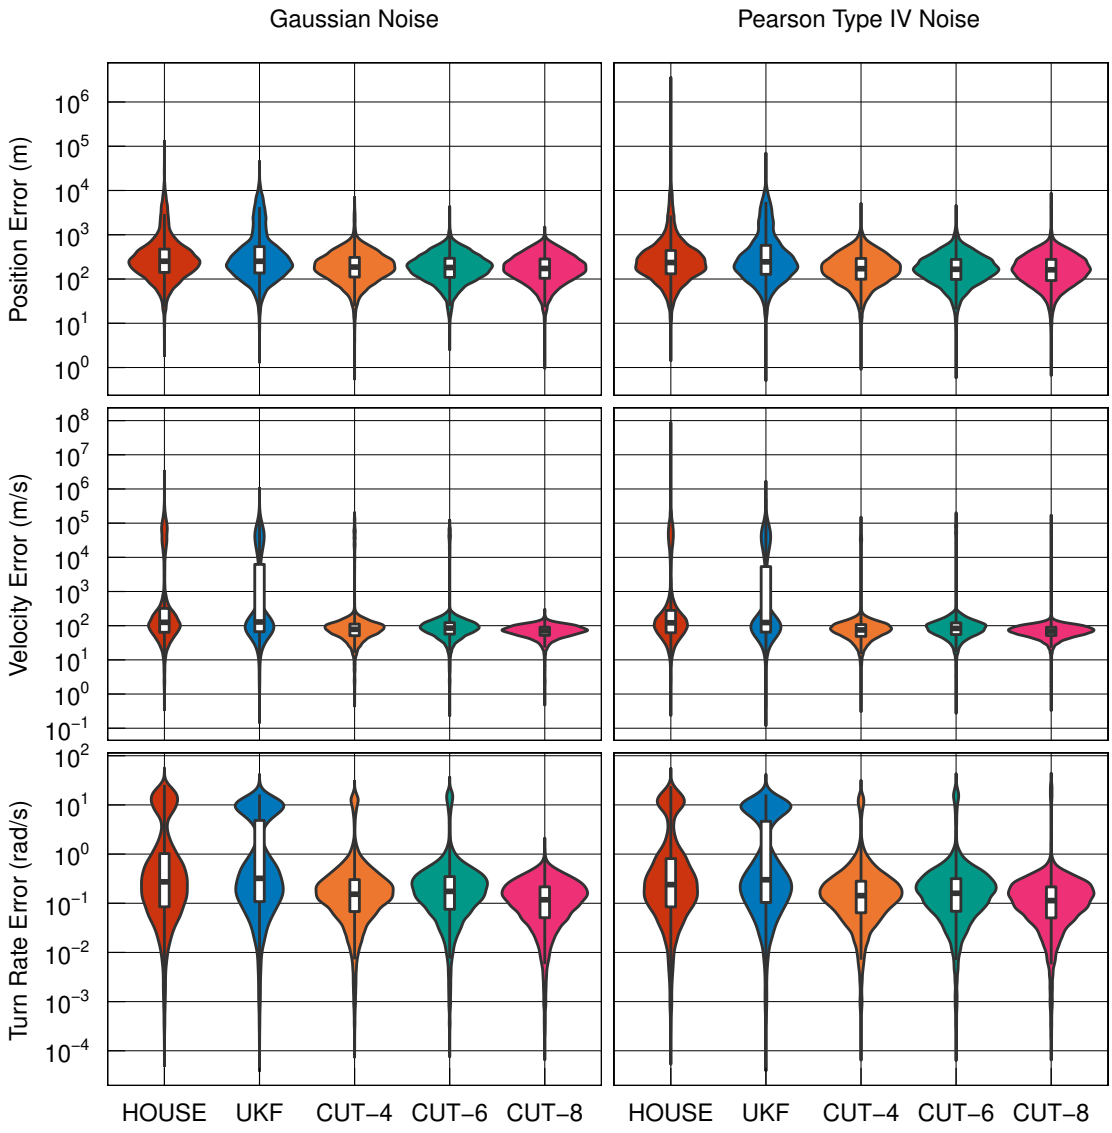

Figure 2.5: Distribution of estimation error for coordinated turn example with sampling time interval 5 s.

2.3.2 Projectile

The projectile's state consists of the position (x, y, z) and velocity $(\dot{x}, \dot{y}, \dot{z})$, and its equations of motion are

$$\ddot{x} = -bv\dot{x} + f_x \tag{2.76}$$

$$\ddot{y} = -bv\dot{y} + f_y \tag{2.77}$$

$$\ddot{z} = -bv\dot{z} + f_z - g \tag{2.78}$$

where b is a constant, (f_x, f_y, f_z) are the specific disturbance forces, g is the acceleration due to gravity ($9.80665\,\mathrm{m/s^2}$), and v is the speed

$$v = \sqrt{\dot{x}^2 + \dot{y}^2 + \dot{z}^2} \tag{2.79}$$

The constant b is given by

$$b = \frac{AC_D\rho}{2m} \tag{2.80}$$

where A is the projectile's surface area, C_D is its drag coefficient, m is its mass, and ρ is the atmospheric density. In this example, we take $b = 0.001\,\mathrm{m^{-1}}$. The observer is located at the origin and takes measurements of azimuth

$$\alpha = \mathrm{atan2}(y, -x) + n_\alpha \tag{2.81}$$

and elevation

$$\epsilon = \mathrm{atan2}\left(z, \sqrt{x^2 + y^2}\right) + n_\epsilon \tag{2.82}$$

where n_α and n_ϵ represent measurement noise.

The disturbance forces are taken to be independent with mean zero and standard deviation $0.01\,\mathrm{m/s^2}$. The two components of the measurement noise n_α and n_ϵ are assumed to be independent with mean zero and a standard deviation of one arcminute. The initial (prior) state distribution has mean $(1000, 1000, 0)$ m

and standard deviation 250 m for the position and mean $(500, 0, 500)$ m/s and standard deviation 100 m/s for the velocity. All components of the initial state are assumed to be independent as well. We consider two cases: one in which the distributions of the initial state, process noise, and measurement noise are Gaussian, and one in which they are Pearson type IV. In the latter case, all of the distributions have kurtosis 30; the initial state and process noise have skewness 1, and the measurement noise has skewness -1. The measurements are sampled at a rate of 5 Hz.

For both the Gaussian case and the Pearson type IV case, we compare HOUSE to the UKF (with tuning factor $\kappa = 1$), CUT-4, and CUT-6 over 100 trials. (We do not test CUT-8 in this example, because the augmented state in the prediction step has dimension $n = 9$, while solutions for CUT-8 sigma points are known for $n \leq 6$ [3].) In each trial, the projectile's trajectory is simulated with randomly generated initial conditions and process noise, with a terminal condition of $z = 0$ (i.e., when the projectile hits the ground). Trajectories lasting less than one second are rejected. Based on each trajectory, a sequence of azimuth-elevation measurements is generated with random measurement noise.

Figure 2.6 shows RMSE throughout the projectile's flight for each of the filters, averaged over 100 trials. Since the total time-of-flight varies between trials, we consider the percentage of the total time-of-flight, rather than the actual time. In the case with Gaussian distributions, the RMSE is nearly equal for HOUSE and the conventional UKF, but it is an order of magnitude lower for CUT-4 and CUT-6. Again, this is expected, because of the design of the CUT filters. However, in the case with Pearson type IV distributions, the RMSE is between one and three orders of magnitude lower for HOUSE than for the UKF, CUT-4, or CUT-6.

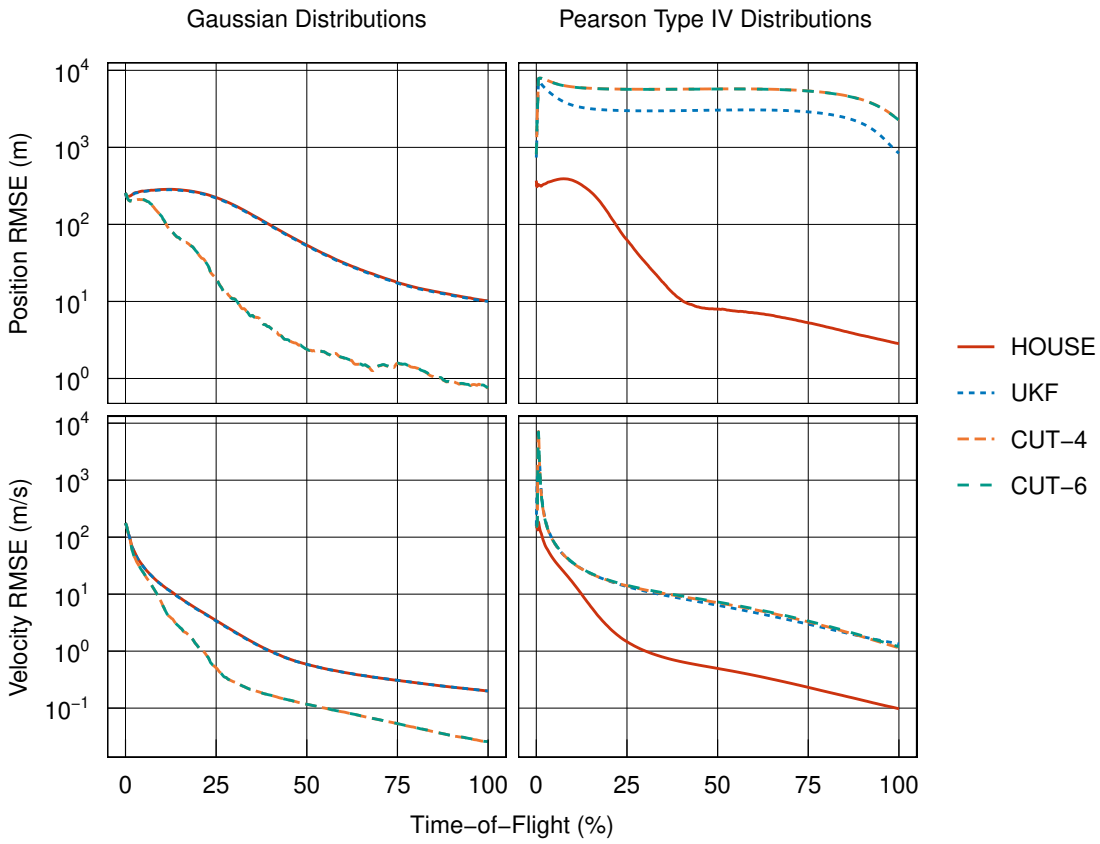

Figure 2.6: Root-mean-square estimation error for projectile example.

Figure 2.7 shows the overall distribution of the estimation error for each of the filters for time $t > 1$ s after launch. (We do not consider the first second of flight, since the high initial errors are not representative of the filters' performance.) It is clear that, in the case with Pearson type IV distributions, the RMSE for the UKF, CUT-4, and CUT-6 is dominated by outliers—those cases in which the position error is 10 km or greater. Except for these outliers, the error distributions for the UKF, CUT-4, and CUT-6 are very similar to that of HOUSE. In the case with Gaussian distributions, on the other hand, there are no such outliers. This is a direct effect of the kurtosis of the distributions, which is interpreted as their "propensity to produce outliers" [78]. Specifically, since a Pearson type IV distribution has a higher kurtois than a Gaussian distribution, we expect to

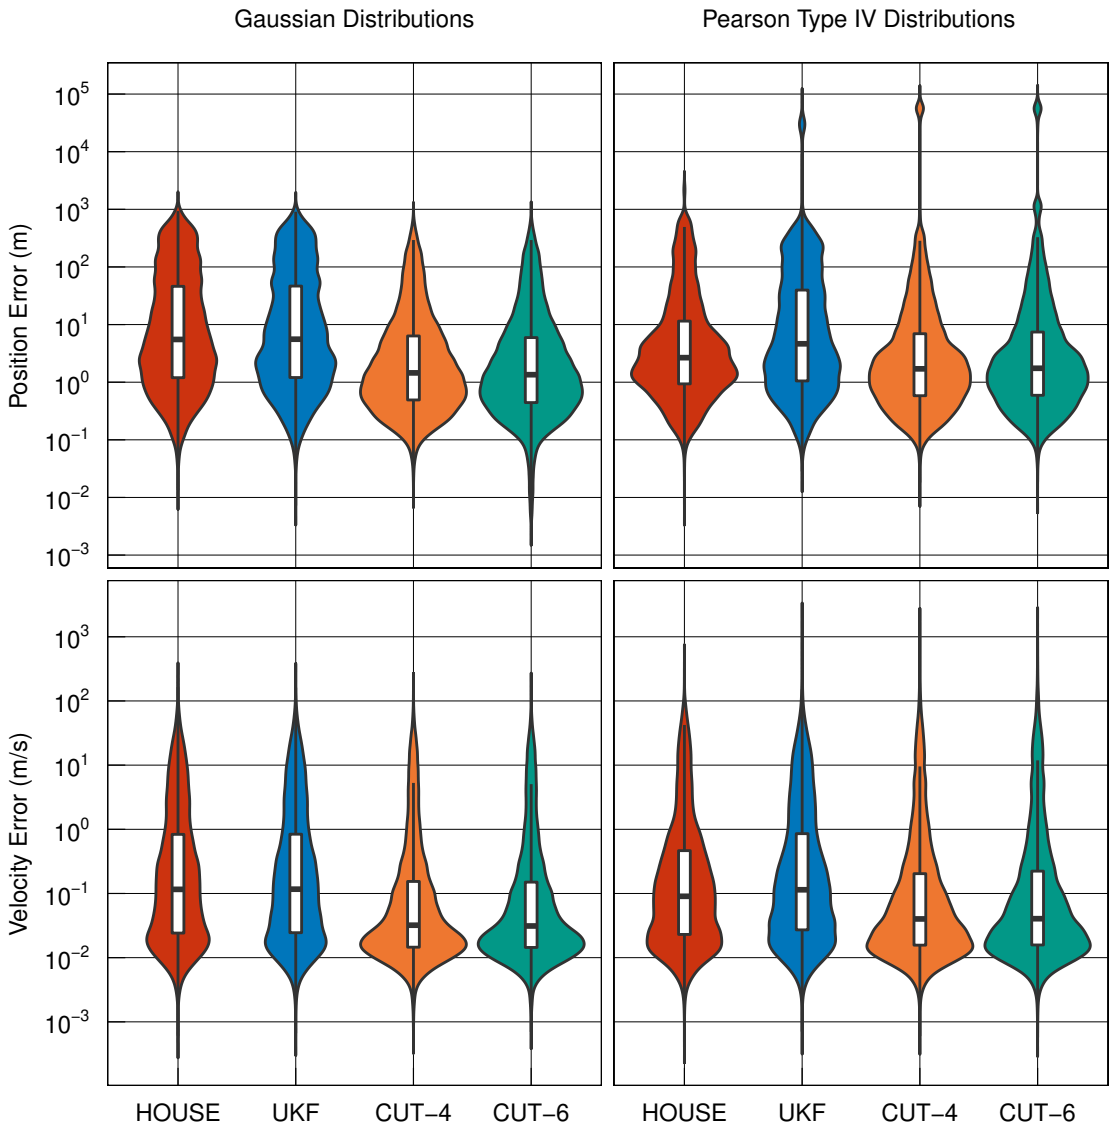

Figure 2.7: Distribution of estimation error for projectile example.

see more outliers in the Pearson type IV case. Based on Figures 2.6 and 2.7, it appears that HOUSE, which directly accounts for the higher kurtosis, is more robust in the presence of outliers than CUT-4, CUT-6, or the UKF.

Figure 2.8 shows the average run times for each of the filters. While HOUSE is slightly slower than the UKF, it is an order of magnitude faster than the CUT-4 or CUT-6 filters, since HOUSE requires many fewer sigma points.

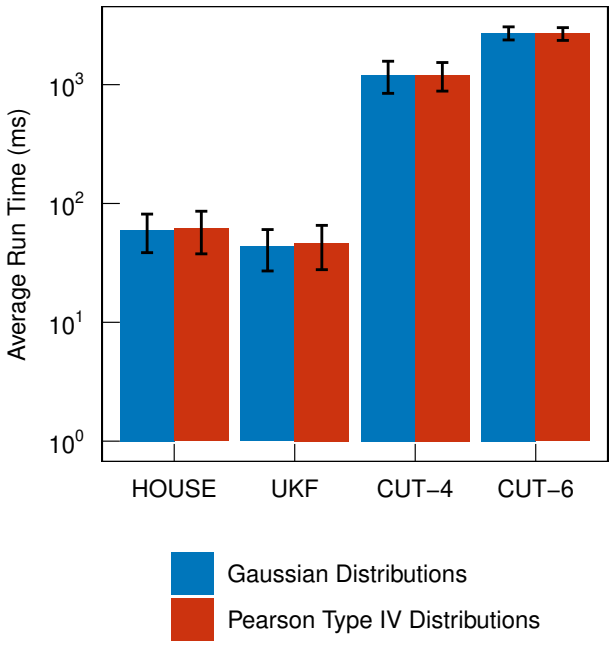

Figure 2.8: Filter run times for projectile example.

2.3.3 Rigid Body

We consider a rigid body with principal moments of inertia about the center of mass I_1, I_2, I_3 and angular velocity components $(\omega_1, \omega_2, \omega_3)$ in the directions of the principal axes, respectively. The time evolution of the angular velocities is described by the Euler equations

$$I_1\dot{\omega}_1 = (I_2 - I_3)\omega_2\omega_3 + \tau_1 \tag{2.83}$$

$$I_2\dot{\omega}_2 = (I_3 - I_1)\omega_3\omega_1 + \tau_2 \tag{2.84}$$

$$I_3\dot{\omega}_3 = (I_1 - I_2)\omega_1\omega_2 + \tau_3 \tag{2.85}$$

where τ_1, τ_2, τ_3 are external disturbance torques about the center of mass in the directions of the principal axes, respectively. We assume that only ω_1 is measured directly.

In our simulations, we consider an asymmetric rigid body with I_1 =

$900\,\mathrm{kg}\cdot\mathrm{m}^2$, $I_2 = 800\,\mathrm{kg}\cdot\mathrm{m}^2$, and $I_3 = 700\,\mathrm{kg}\cdot\mathrm{m}^2$. The disturbance torques are taken to be zero-mean and independent, with standard deviation $0.001\ \mathrm{N}\cdot\mathrm{m}$. The initial angular velocities sampled from a distribution with mean zero and standard deviation 0.01 rad/s in each direction. The measurement noise has mean zero and standard deviation 0.001 rad/s, and the measurements are sampled at 10 Hz. As in the previous example, we consider Gaussian and Pearson Type IV distributions for the initial state, process noise, and measurement noise; specifically, we consider Pearson distributions with skewness $\gamma = -1$ and kurtosis $\kappa = 30$. For both distribution types, we test HOUSE, the UKF (with tuning factor $\kappa = 1$), CUT-4, CUT-6, and CUT-8 in 100 trials.

The RMSE and error distributions for the filters are shown in Figures 2.9 and 2.10, respectively. Overall, the differences in RMSE between the filters are small, and the RMSE is of the same order of magnitude for all of the filters. The filters' error quartiles are very similar as well. However, in the case with Pearson Type IV distributions, we observe a similar effect as with the projectile: HOUSE has fewer and less extreme outliers than the UKF or any of the CUT filters. Again, this is evidence that HOUSE is better equipped to process outliers because of its propagation of kurtosis.

Figure 2.11 shows the average runtimes for each of the filters. As in the projectile example, HOUSE is slightly slower than the UKF but signficantly faster than CUT-4, CUT-6, and especially CUT-8.

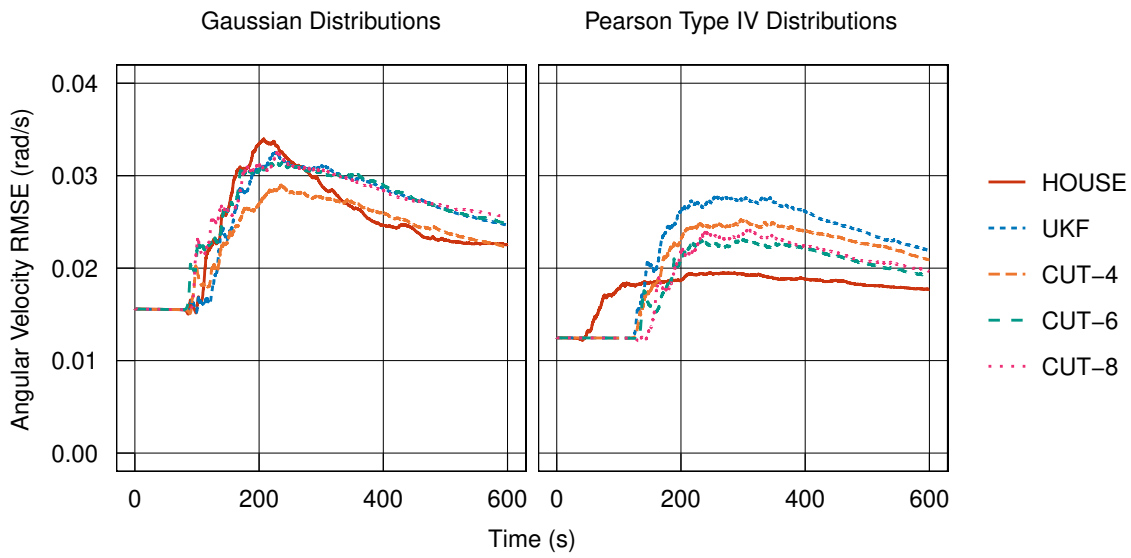

Figure 2.9: Root-mean-square estimation error for rigid body example.

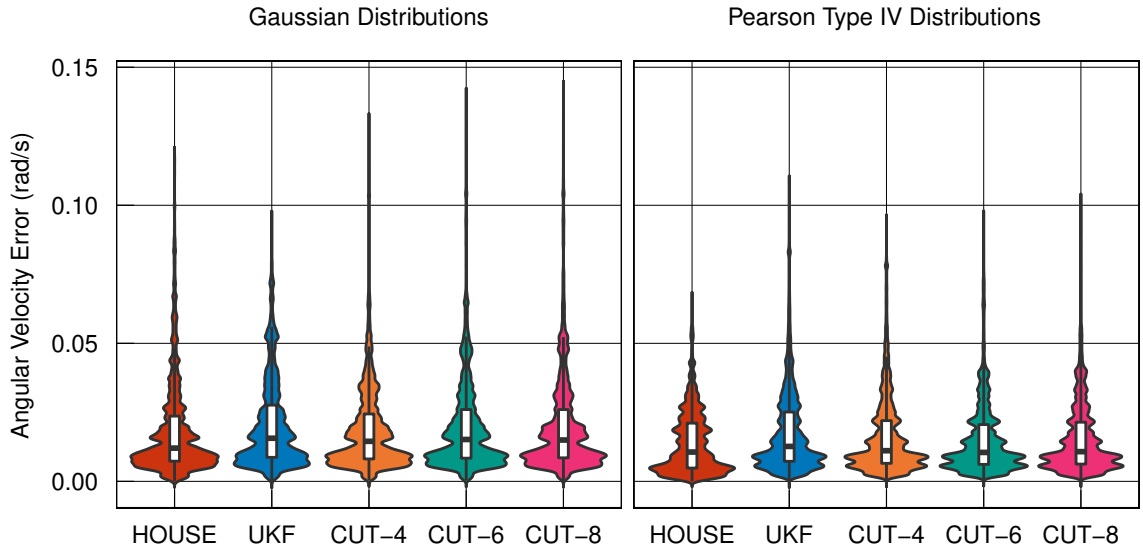

Figure 2.10: Distribution of estimation error for rigid body example.

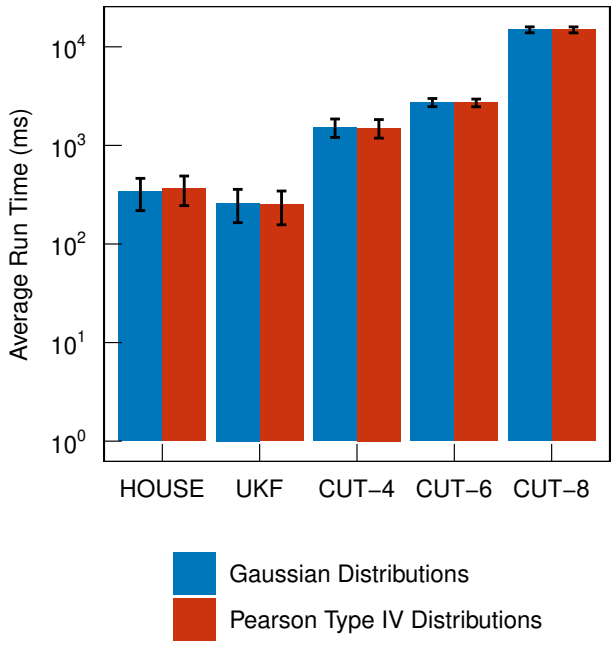

Figure 2.11: Filter run times for rigid body example.

2.4 Conclusion

We have proposed a new extension of the unscented Kalman filter with third and fourth order moment matching and explicit formulas for the sigma points and weights. Its computational cost is only slightly greater than that of the conventional UKF. We have also described conditions under which the sigma point formulas are applicable and the filter is operable. In simulations of dynamical systems, we have found that the new filter is more robust for distributions with high kurtosis than the conventional UKF or the CUT filters. Specifically, we have found that HOUSE generates many fewer outliers in the estimation error in cases where the initial conditions and noise have high kurtosis. Also, we have found that the run times for HOUSE are significantly shorter than for the CUT filters and only slightly longer than for the conventional UKF.

In our future work, we plan to further refine the HOUSE method to increase

its accuracy and range of applicability. In particular, we will investigate modifications to HOUSE that would minimize the distortion of moments and make the filter's performance less sensitive to the dimension of the system. Since these improvements would most likely require increasing the number of sigma points, and hence the computational complexity, we will also study in detail the trade-off between the speed and accuracy of higher-moment unscented filters.

Finally, we plan to apply HOUSE to the problem of satellite tracking using unconventional measurements—such as those obtained from ground imaging data—coupled with more conventional measurement techniques. Because the unconventional measurements might have highly non-Gaussian distributions, we expect that HOUSE would be well suited for this problem. We plan to test the performance of these new methods in spacecraft navigation and attitude determination using a realistic simulation of a satellite's dynamics and sensors.

CHAPTER 3

AUTONOMOUS CROSS-CALIBRATION FOR IMAGING SATELLITES

Cross-calibration between satellites is crucial to the performance of Earth-observing constellations. To this end, very rigorous cross-calibration schemes have been developed and implemented, e.g., for the Dove constellation [44]. However, these techniques rely heavily on communications with the ground station and humans-in-the-loop; as a result, they may not be feasible for larger constellations. To mitigate this problem, we propose a method called Autonomous Cross-Calibration for Imaging Satellites (ACCIS). This method uses measurements obtained from the primary mission images, combined with conventional position and attitude measurements, to estimate the states of satellites and their cameras, with key data transmitted between satellites to achieve accurate cross-calibration. This method was first presented in [69].

In this chapter, we describe in detail the concept of operation for ACCIS and analyze its performance using a realistic simulation. In ACCIS, each satellite computes a real-time estimate of its state, which includes its postion, attitude, and camera parameters. This is done using the square root sigma point filter, discussed in Section 1.4.3, which is robust and computationally efficient for a wide range of systems with nonlinear dynamics and measurements. Here, the filter processes not only conventional position and attitude measurements, but also measurements obtained from images. Specifically, it uses features, or key points, extracted from images using the Scale-Invariant Feature Transform (SIFT) introduced by Lowe [46]. For each key point found in the image, the SIFT algorithm computes a position, orientation, and scale, as well as a descriptor that is invariant under translation, rotation, and scaling. This allows us to develop a

model that maps changes in the imaging satellite's position, attitude, and camera parameters to changes in a key point's position, orientation, and scale.

During a mission, ACCIS would operate as follows. Whenever a satellite takes an image, it extracts the key points using SIFT. It then transmits the key points, along with the state estimate and covariance from its filter, to other satellites. Then, when two or more satellites have imaged approximately the same area, each of them can use the difference in the key points and the estimated states to update its own state estimate.

Compared to existing techniques, ACCIS has several features that could make cross-calibration faster and cheaper for large constellations of imaging satellites. First, it is fully autonomous, requiring no humans-in-the-loop. Also, it does not require uplinks nor downlinks, but only crosslinks between satellites, reducing the communication load for the ground station. Furthermore, the crosslinks transmit only the state estimates, covariance matrices, and SIFT key points and descriptors, which is much less expensive than transmitting full images. In addition, this method uses only the primary mission data, requiring no dedicated calibration measurements.

We test the performance of ACCIS by simulating a minimal constellation of two imaging satellites in low Earth orbit. Each of the satellites is equipped with a GPS receiver, a gyroscope, a star tracker, and a nominally nadir-pointing camera for imaging. The simulation features a detailed model of the satellites' rigid-body attitude dynamics and orbital motion, including high-fidelity models for perturbations such as non-spherical Earth gravity and aerodynamic drag. We also model the orientation, focus, and distortion parameters of the satellites' cameras. To emulate the raw data obtained by the satellites, we generate

synthetic images from Landsat data; the original Landsat images are trimmed, projected, and distorted based on the satellites' position, field of view, and camera distortion parameters.

3.1 State Models

In ACCIS, the satellite state is considered to include both its dynamical state and its camera parameters. In this section, we describe all of the state components, how they are propagated in time, and how they relate to the imaging operations of the satellite. The positions and reference frames used in our model are shown in Figure 3.1.

3.1.1 Satellite Dynamics

We model each satellite as a rigid body in a perturbed Keplerian orbit. Let G denote the satellite's center of mass and O' the center of the Earth. Furthermore, let $\mathcal{G}$ denote the Earth-centered inertial (ECI) frame and $\mathcal{B}$ a body-fixed frame. To simplify calculations, we choose $\mathcal{B}$ to be the principal axis frame, though in theory any other body-fixed frame could be used. The satellite's dynamical state consists of the position $\mathbf{r}_{G/O'}$, velocity $\mathbf{v}^{\mathcal{G}}_{G/O'}$, attitude quaternion $\mathbf{q}_{\mathcal{B}/\mathcal{G}}$, and angular velocity $\boldsymbol{\omega}_{\mathcal{B}/\mathcal{G}}$. The translational dynamics of the satellite are governed by

$$\dot{\mathbf{v}}^{\mathcal{G}}_{G/O'} = \mathbf{g} + \mathbf{u}_G + \mathbf{f}_G, \tag{3.1}$$

where $\mathbf{g}$ is the gravitational acceleration, $\mathbf{u}_G$ is the specific force due to controls (e.g., thrusters), and $\mathbf{f}_G$ is the sum of all other perturbing specific forces, including

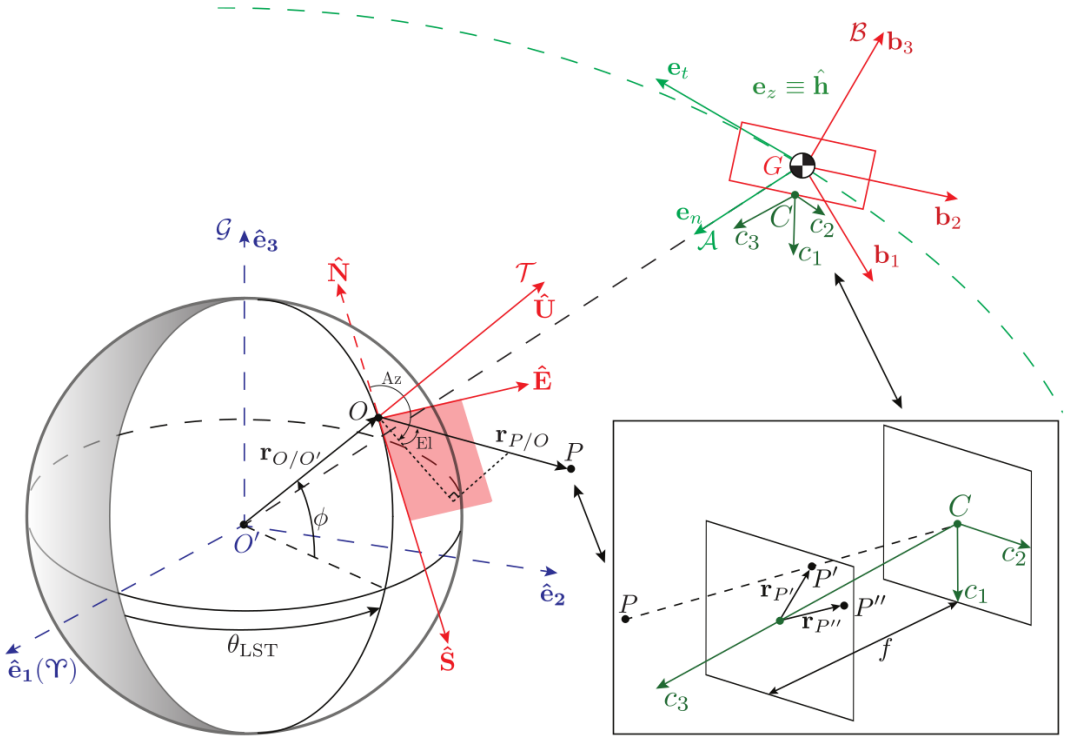

Figure 3.1: Positions and reference frames used in ACCIS [62]. We model the motion of the satellite's center of mass G in the Earth-centered inertial (ECI) frame $\mathcal{G}(O', \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$, and for the satellite's rotational dynamics, we use the body-fixed frame $\mathcal{B}(G, \mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3)$. The orbital frame $\mathcal{A}(G, \mathbf{e}_t, \mathbf{e}_n, \mathbf{e}_z)$ is defined by the nadir direction $\mathbf{e}_n$, the normal to the orbital plane $\mathbf{e}_z$ (which is the orbital angular momentum direction $\hat{\mathbf{h}}$), and the tangent vector $\mathbf{e}_t$. As for the camera, we use the frame $C(C, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3)$, in which C is the camera's position, and $\mathbf{c}_3$ is its boresight axis. The boresight points to a position O on the Earth's surface, with geodetic latitude ϕ and local sidereal time θ_{LST}. The position of a point P near O can be described by the azimuth, elevation, and range in the topocentric frame $\mathcal{T}(O, \hat{\mathbf{S}}, \hat{\mathbf{E}}, \hat{\mathbf{U}})$. The camera, with focal length f, maps P to the position P' in the image plane, which is observed at position P'' in the image due to distortion.

atmospheric drag. The rotational dynamics of the satellite are governed by

$$\dot{\boldsymbol{\omega}}^{\mathcal{B}}_{\mathcal{B}/\mathcal{G}} = \mathbb{I}_G^{-1}(\mathbf{M}_G + \boldsymbol{\tau}_G - \boldsymbol{\omega}_{\mathcal{B}/\mathcal{G}} \times (\mathbb{I}_G \cdot \boldsymbol{\omega}_{\mathcal{B}/\mathcal{G}})), \tag{3.2}$$

where $\mathbb{I}_G$ is the inertia tensor of the satellite about G, $\mathbf{M}_G$ is the total control moment due to ADCS actuators, and $\boldsymbol{\tau}_G$ is the total disturbance torque. In the filter, we model $\mathbf{f}_G$ and $\boldsymbol{\tau}_G$ as random process noise, due to the difficulty of accurately predicting the atmospheric density, the satellite's aerodynamics properties, and other factors.

We define an additional reference frame $\mathcal{A}$ as the nadir-pointing orbital frame, which depends on $\mathbf{r}_{G/O'}$ and $\mathbf{v}^{\mathcal{G}}_{G/O'}$. For filtering purposes, we use the Rodrigues parameter $\boldsymbol{\rho}_{\mathcal{B}/\mathcal{A}}$ rather than the quaternion $\mathbf{q}_{\mathcal{B}/\mathcal{G}}$. Thus, we avoid the "redundancy" of the fourth quaternion component. Also, in normal satellite operations, $\boldsymbol{\rho}_{\mathcal{B}/\mathcal{A}}$ is small and never approaches the singularity for Rodrigues parameters. To maintain an attitude in which the camera frame C is close to the nadir-pointing orbital frame $\mathcal{A}$, we implement a proportional-derivative (PD) law for the control torque (see [49]).

For the Earth's gravitational field, we use a truncated version of the EGM2008 [53], with the gravitational forces evaluated using the procedures described by Gottlieb [24]. For transformations between the ECI frame and the Earth-centered Earth-fixed frame (ECEF), we use the Naval Observatory Vector Astronomy Subroutines (NOVAS) [36]. When propagating the "ground truth" state in the simulation, we compute drag forces using the Jacchia-Bowman 2008 (JB2008) atmospheric model [9].

3.1.2 Camera Model

Let P be an arbitrary point in the camera's field of view. In an ideal pinhole camera at point C, this point is projected onto a point P' in the image plane, at the focal distance u from the camera origin and opposite of C. In a real camera, on the other hand, P is projected onto a different point P'' due to distortions. For simplicity, we assume a purely radial distortion model, i.e., that the positions $\mathbf{r}_{P'}$ and $\mathbf{r}_{P''}$ satisfy

$$\frac{\mathbf{r}_{P'}}{\|\mathbf{r}_{P'}\|} = \frac{\mathbf{r}_{P''}}{\|\mathbf{r}_{P''}\|}. \tag{3.3}$$

Such models are widely used to correct distortions in commercial lenses [23]. Specifically, we use a third-order radial distortion model,

$$\|\mathbf{r}_{P''}\| = \|\mathbf{r}_{P'}\| \left(1 - c_1 - c_2 - c_3 + c_1 \frac{\|\mathbf{r}_{P'}\|}{R} + c_2 \frac{\|\mathbf{r}_{P'}\|^2}{R^2} + c_3 \frac{\|\mathbf{r}_{P'}\|^3}{R^3}\right), \tag{3.4}$$

where c_1, c_2, and c_3 are the distortion parameters for a particular lens, and R is the radius of a circle circumscribed about the image; that is,

$$R = \frac{1}{2}\sqrt{W^2 + L^2}, \tag{3.5}$$

where W and L are the width and length of the image, respectively. This scaling ensures that the distortion parameters are of the same order of magnitude, which improves numerical stability in the filter. Furthermore, for filtering purposes, we combine the distortion parameters into a vector $\mathbf{c} = \begin{bmatrix} c_1 & c_2 & c_3 \end{bmatrix}^{\mathrm{T}}$.

In addition to the distortion, we account for imperfect focusing of the camera. This depends on the focal length f of the lens, the pupil diameter A, the distance u from the camera to the object, and the density ρ of pixels per unit length. The effects of defocusing can be modeled well by a Gaussian blur with standard deviation [48]

$$\sigma_G = \frac{\sigma}{\sqrt{2}}, \tag{3.6}$$

where

$$\sigma = \rho \frac{fs}{2N} \left(\frac{1}{f} - \frac{1}{u} - \frac{1}{s} \right) \tag{3.7}$$

and

$$N = \frac{f}{A}. \tag{3.8}$$

Finally, we account for the fact that the orientation of the camera frame C with respect to the body frame $\mathcal{B}$ is not known exactly. This can be due to structural tolerances, as well as dynamical or thermal deformation. We parametrize this attitude by a quaternion $\mathbf{q}_{C/\mathcal{B}}$, and in the filter, we use the Rodrigues parameter $\boldsymbol{\rho}_{C/\mathcal{B}}$.

All of the camera parameters that are estimated by the filter—namely, $\mathbf{c}$, f, and $\mathbf{q}_{C/\mathcal{B}}$—are assumed to be constant in time. Therefore, in the filter they are modeled as having unity dynamics with zero process noise. In the future, we will add a noise term to account for drift in the parameters, particularly for $\mathbf{q}_{C/\mathcal{B}}$.

In the simulation, we generate synthetic satellite images based on Landsat data and our camera model. Specifically, we obtain a sector of a Landsat image mosaic using the NASA World Wind library [54]. Then, we apply a perspective transform, based on the satellite's position, attitude, and camera parameters. Finally, we apply the radial distortion and defocusing blur. For the image transform, distortion, and blurring, we use the OpenCV library [10].

3.1.3 Combined State

By combining the dynamical state of the satellite with the camera parameters, we obtain the overall system state that is estimated by the filter in ACCIS.

$$\mathbf{x} = \begin{bmatrix} \mathbf{r}_{G/O'}{}^{\mathrm{T}} & \mathbf{v}_{G/O'}^{\mathcal{G}}{}^{\mathrm{T}} & \boldsymbol{\omega}_{\mathcal{B}/\mathcal{G}}{}^{\mathrm{T}} & \boldsymbol{\rho}_{\mathcal{B}/\mathcal{A}}{}^{\mathrm{T}} & \boldsymbol{\rho}_{C/\mathcal{B}}{}^{\mathrm{T}} & f & \mathbf{c}^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}. \tag{3.9}$$

The system state has a total of 19 components. While the rates $\mathbf{v}_{G/O'}^{\mathcal{G}}$ and $\boldsymbol{\omega}_{\mathcal{B}/\mathcal{G}}$ are not necessary for image calibration (assuming sufficiently short exposure time), they are needed for state prediction in the filter.

3.2 Measurement Models

The ACCIS framework features two categories of measurements: conventional measurements of the satellite's dynamical states and measurements derived from images.

3.2.1 Conventional Measurements

In our model, the satellite is equipped with a Global Positioning System (GPS) receiver, which provides measurements of its position and velocity; a gyroscope, to measure its angular velocity; and a star tracker, to obtain precise attitude measurements. The GPS and gyroscope measurements are assumed to have a Gaussian noise distribution. The star tracker measurements are assumed to have a cross-boresight error and a roll error [49], which we also model using Gaussian distributions.

3.2.2 Image-Based Measurements

While straightforward models exist for the conventional measurements, there are no standard methods for relating image data to the satellite's dynamical state and camera parameters. Here, we present a preliminary model for measurements based on features extracted from images.

The Scale-Invariant Feature Transform (SIFT), proposed by Lowe [46], provides a robust and efficient method for extracting and matching features from images. Each feature, or key point, extracted from an image by SIFT includes a position $\mathbf{r}_K$, orientation angle θ, and scale S, as well as a gradient-based descriptor that is invariant under translation, rotation, and scaling. Also, the effects of differences in illumination and perspective are relatively small. Thus, the descriptors can be matched between images taken from various distances, angles, etc. However, the SIFT descriptors are not fully invariant under affine transformations. In the simulation, we use the implementation of SIFT from the OpenCV library [10].

To reduce the chances of "false matches" between key points, we keep only the key points that are in the overlapping region of the two images on Earth's surface, based on the estimates of the satellites' positions and attitudes. Figure 3.2 shows an example of two satellite images with matched key points.

To make the SIFT key points more convenient to use in our model, we use an equivalent representation as two positions, $\mathbf{r}_{K_1'}$ and $\mathbf{r}_{K_2'}$, given by

$$\mathbf{r}_{K_{1,2}'} = \mathbf{r}_K \pm \frac{1}{2} S \mathbf{u}, \tag{3.10}$$

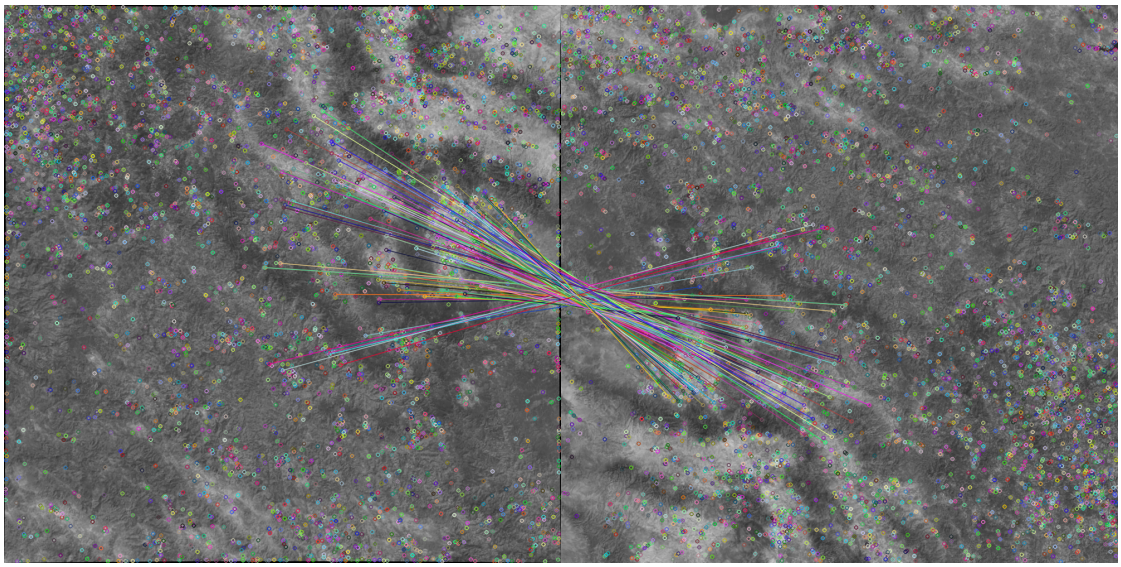

Figure 3.2: Example of two satellite images with matched SIFT key points.

where $\mathbf{u} = \begin{bmatrix} \cos\theta & \sin\theta \end{bmatrix}^{\mathrm{T}}$. Thus, a key point is given by

$$\mathbf{K} = \begin{bmatrix} \mathbf{r}_{K_1'}^{\mathrm{T}} & \mathbf{r}_{K_2'}^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}, \tag{3.11}$$

and we can directly convert between this and the standard key point representation. Given the time t and the satellite state $\mathbf{x}$, we can geometrically map $\mathbf{r}_{K_i'}$ to a point on the Earth's surface and vice versa. For this, we take into account the position and attitude of the satellite, as well as the attitude, distortion, etc. of the camera.

Suppose that a descriptor matching algorithm has matched two key points $\mathbf{K}_1$ and $\mathbf{K}_2$, from images taken at times t_1 and t_2 by satellites with states $\mathbf{x}_1$ and $\mathbf{x}_2$, respectively. We can map $\mathbf{K}_1$ to a pair of points on Earth's surface given t_1 and $\mathbf{x}_1$. Then, given t_2 and $\mathbf{x}_2$, we can map two ground positions to a "predicted" key point $\hat{\mathbf{K}}_2$. Thus, given estimate distributions of $\mathbf{x}_1$ and $\mathbf{x}_2$, we can obtain a distribution of $\hat{\mathbf{K}}$, adding an error term to both $\mathbf{K}_1$ and $\hat{\mathbf{K}}_2$. Then, we can apply $\mathbf{K}_2$ as the "true" measurement in the filter update step, to obtain a better estimate of $\mathbf{x}_2$.

To obtain an error estimate for this measurement model, we sampled 1,000 pairs of satellite states and generated corresponding images. We applied SIFT key point extraction and matching to each pair of images, and we computed the differences between the true key points and the ones predicted by the measurement model. From these differences, we obtained a covariance matrix, which we used in the filter.

3.3 Results

We tested the ACCIS methodology using a simulation of two imaging Earth-imaging satellites. Figure 3.3 shows a flowchart of the entire simulation framework. The spacecraft are in coplanar circular orbits at a nominal altitude of 500 km, with an inclination of 30°. Satellite 2 follows behind Satellite 1, with a 4° difference in phase. Both satellites take one image per minute. We simulated 4000 minutes of flight time, which is approximately 42 orbital periods. The complete simulation code is available on GitHub [70].

Using conventional measurements, both satellites were able to obtain accurate measurements of their dynamical state, as shown in Figures 3.4 and 3.5. Also, they are able to accurately maintain a nadir-pointing attitude, as shown in Figure 3.6.

Figure 3.7 shows the estimation errors for the camera attitude and focal length. These errors remain nearly constant throughout the simulation. For the focal length, this was to be expected, since the SIFT-based measurement model described in Section 3.2.2 does not account for it. As for the camera attitude error, it starts at the same order of magnitude as the body attitude estimation

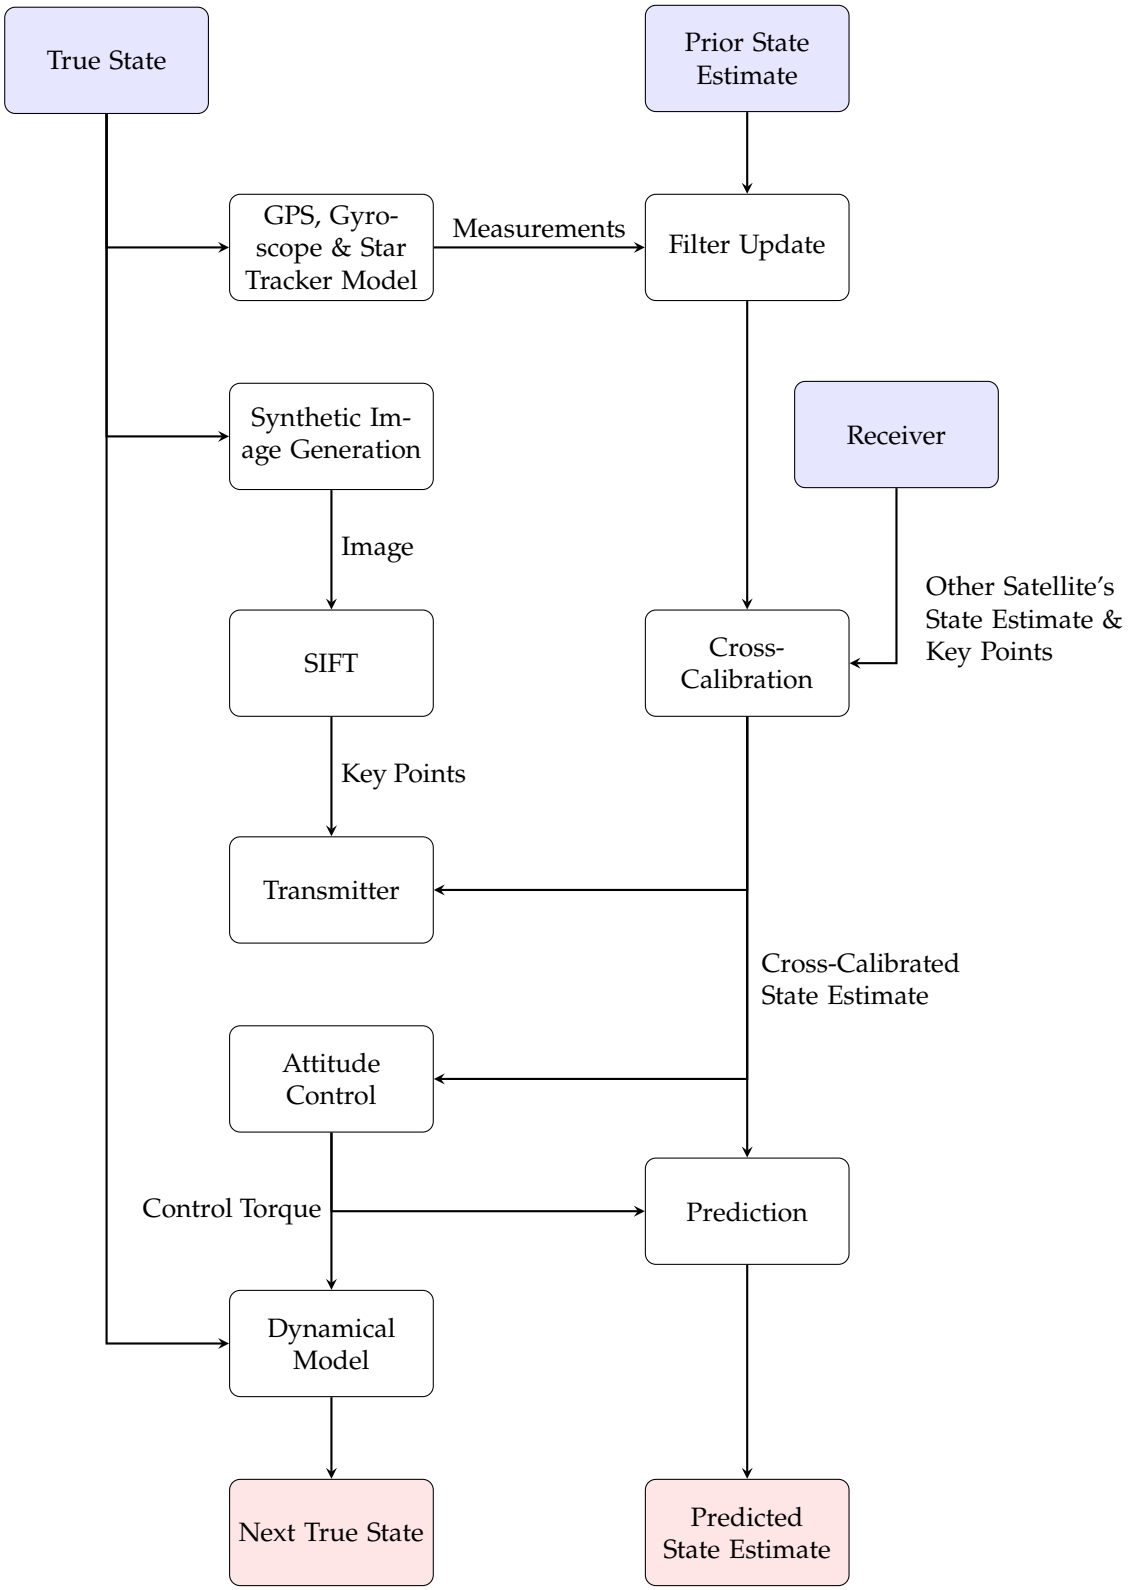

Figure 3.3: Flowchart of the ACCIS simulation framework.

error, shown in Figure 3.5, and the pointing error, shown in Figure 3.6. Thus, the estimator cannot determine whether attitude discrepancies detected using the image-based measurements are due to errors in the body attitude or the camera attitude. Perhaps if the attitude estimates were made even more accurate, then the camera attitude could be estimated more accurately as well.

Figure 3.8 shows the estimation errors for the three camera distortion parameters. For Satellite 2, the error for c_1 and c_2 drops rapidly early in the simulation and then remains nearly constant. This shows that the cross-calibration process using SIFT-based measurements can significantly improve estimates of the distortion parameters—at least for the lower-order ones. For c_3, on the other hand, the error increases slightly. As for Satellite 1, the estimation errors remain nearly constant throughout the simulation, except for very small changes near the beginning. This is because Satellite 1 had fewer opportunities for cross-calibration, since it was always "ahead" of Satellite 2.

3.4 Conclusion

We have demonstrated the feasibility of fully autonomous, real-time cross-calibration for Earth-imaging satellites using features extracted from images. To this end, we have developed a detailed simulation framework for the dynamics, measurements, imaging, and state estimation for a constellation of satellites. In our future work, we plan to further refine the cross-calibration methodology, particularly in the emerging area of image-based measurements.

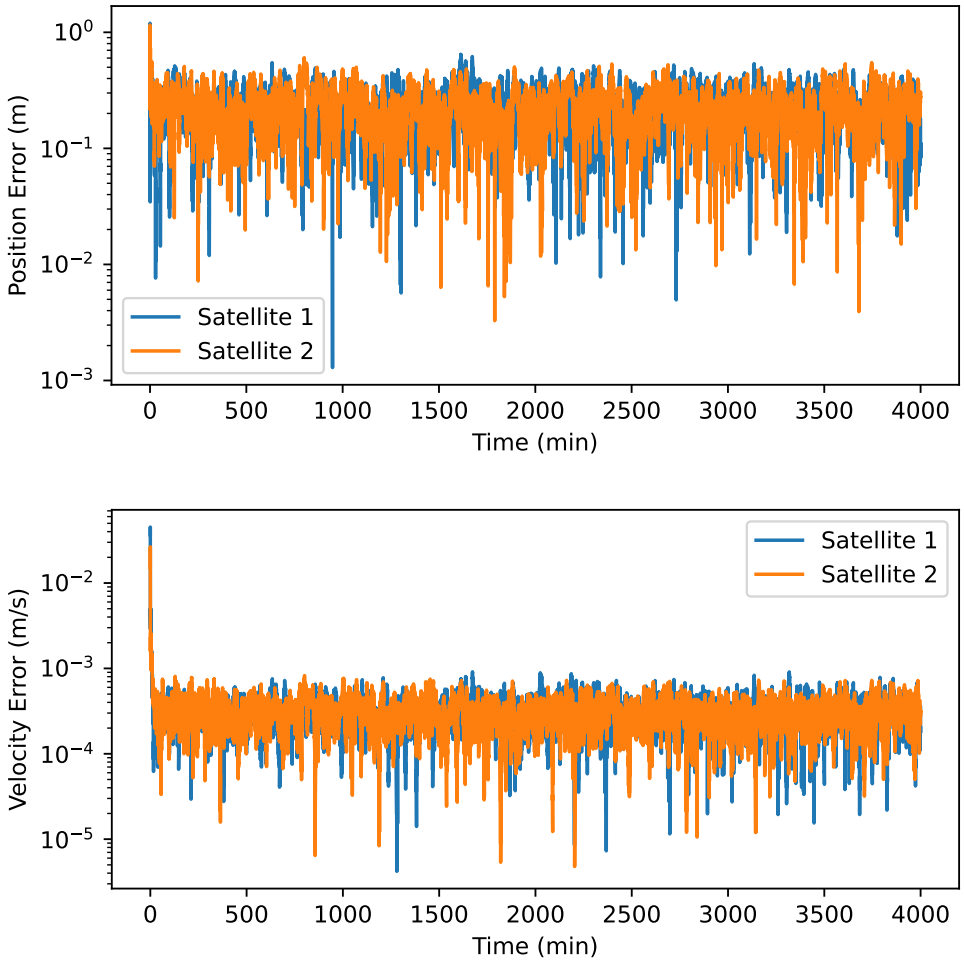

Figure 3.4: Estimation errors for translational motion.

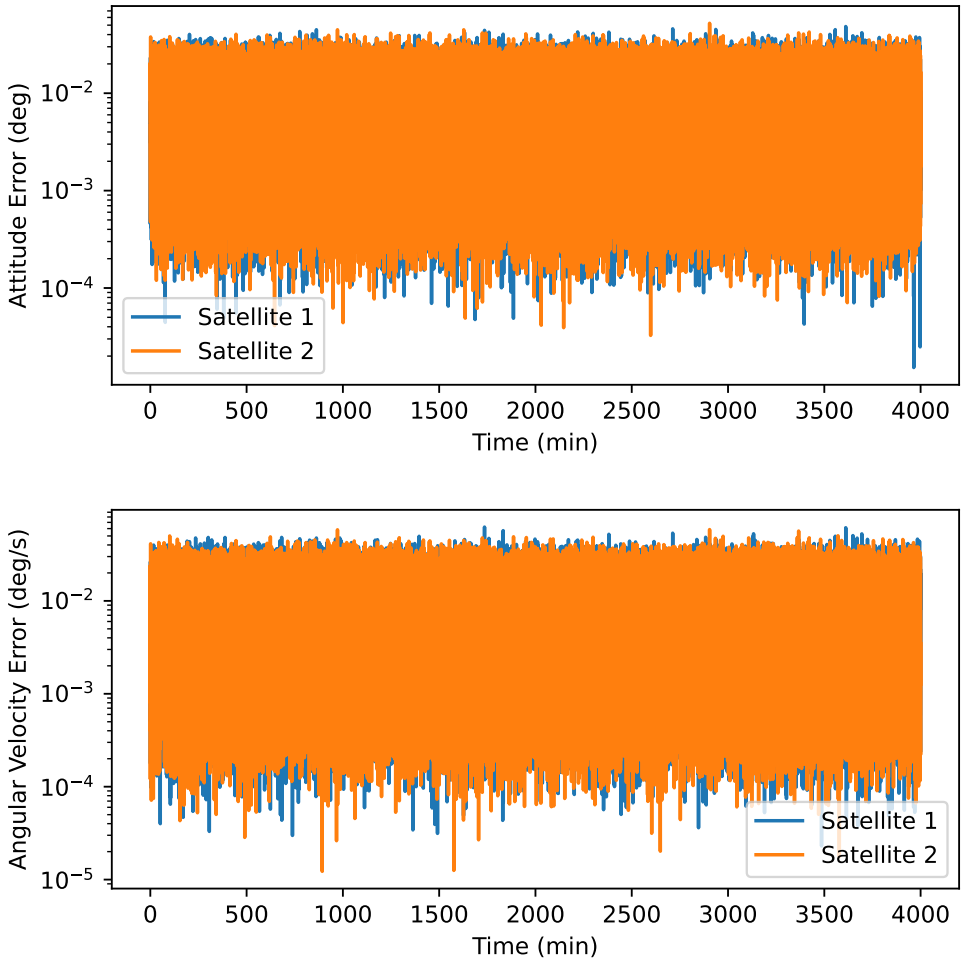

Figure 3.5: Estimation errors for rotational motion.

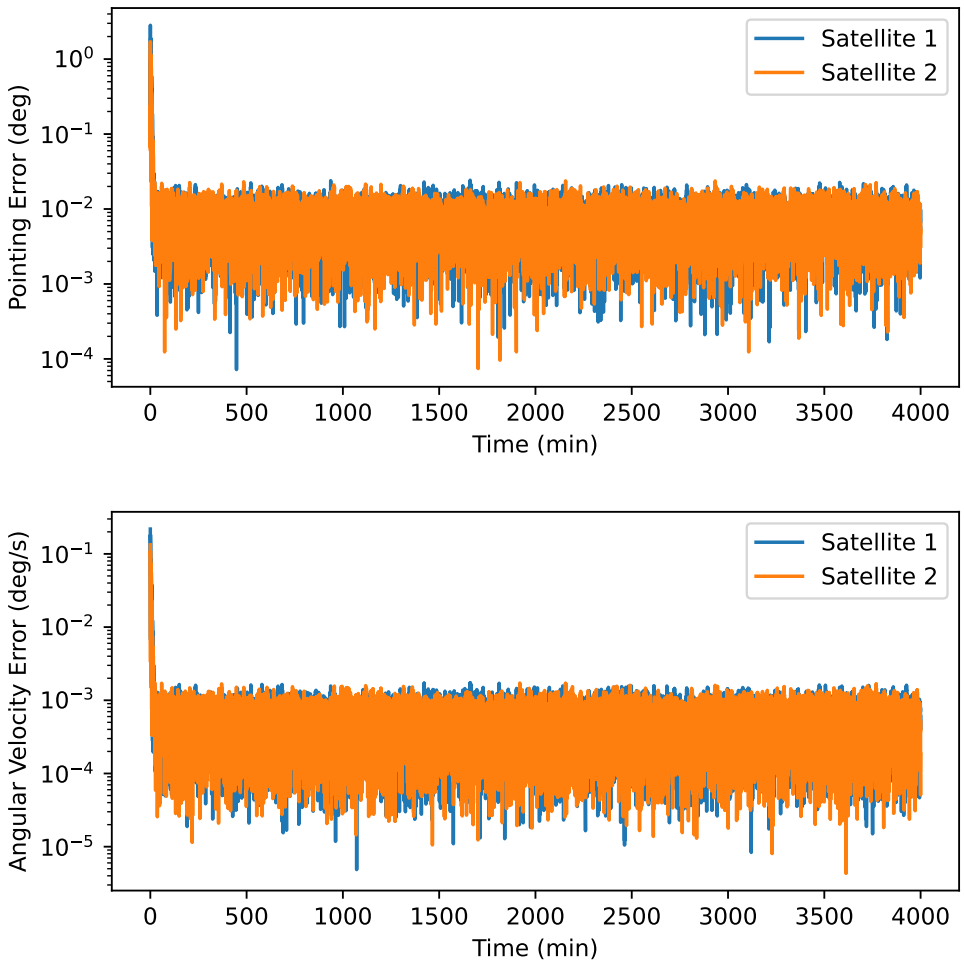

Figure 3.6: Attitude control errors.

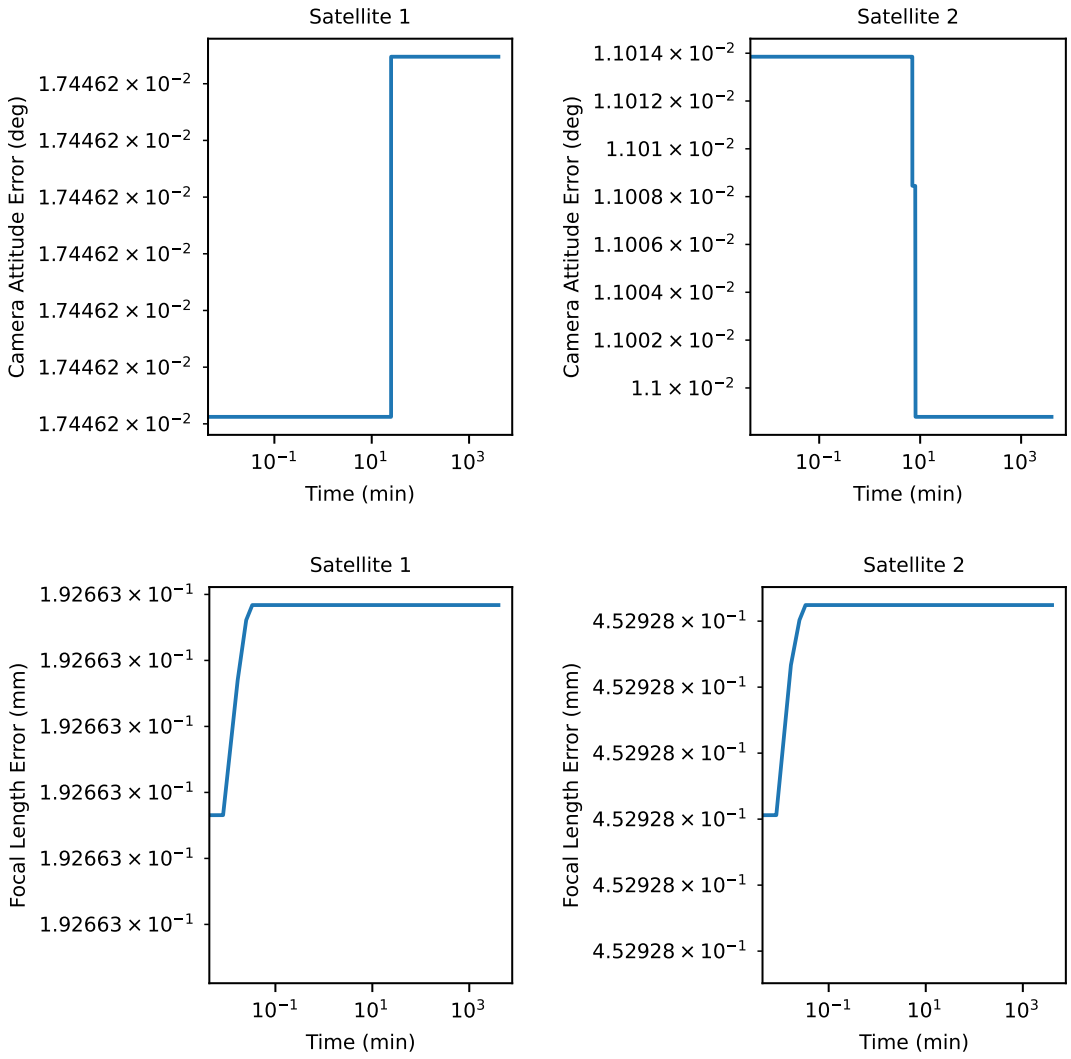

Figure 3.7: Camera attitude and focal length errors. The effects of the cross-calibration on the estimates of these parameters are negligible.

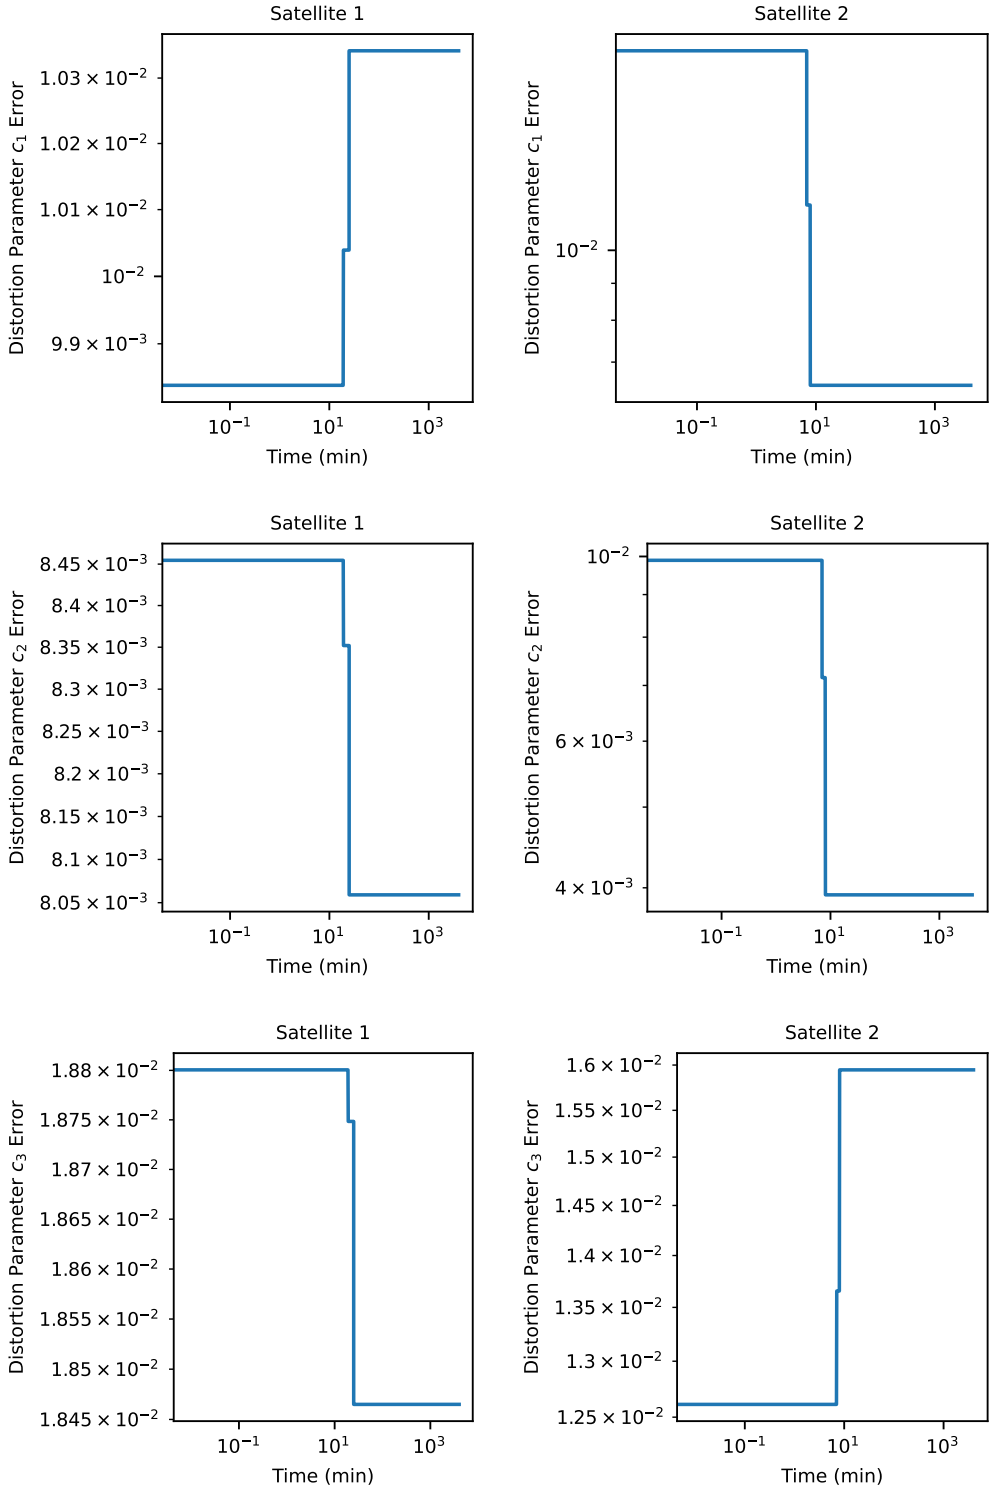

Figure 3.8: Camera distortion parameter estimation errors. Early in the simulation, the cross-calibration reduces the error for c_1 and c_2 for Satellite 2 by an order of magnitude; however, the c_3 error increases slightly. Since Satellite 1 has fewer cross-calibration opportunities, its errors remain nearly constant.

CHAPTER 4

UNSCENTED FILTERING FOR DIRECTLY-OBSERVED EXOPLANET ORBITS

An exoplanet is a planet outside of our solar system. When an exoplanet is directly observed, we can measure a projection of its position onto the image plane, as shown in Figure 4.1. Recent developments in imaging instruments, such as the Gemini Planet Imager (GPI) [47], have enabled increasingly precise astrometric measurements of directly-imaged exoplanets, with measurement errors within a few milliarcseconds [40, 51, 60, 77]. Still, the problem of fitting orbits to astrometric data remains challenging, due to the sparsity of the measurements and the highly nonlinear dynamics. Some of the difficulties of the exoplanet orbit fitting problem are illustrated in Figure 4.2. In this chapter, we present a new method for exoplanet orbit fitting from direct observations, based on the unscented Kalman filter and a new set of orbital elements.

4.1 Review of Exoplanet Orbit Fitting Methods

In this section, we briefly review the orbit fitting methods that have been used for exoplanets. While Monte Carlo methods are the most widely applied, the popularity of Bayesian rejection sampling methods is increasing. Our method, on the other hand, is based on nonlinear filtering, which was previously applied to the exoplanet orbit fitting problem by Savransky and Kasdin [64].

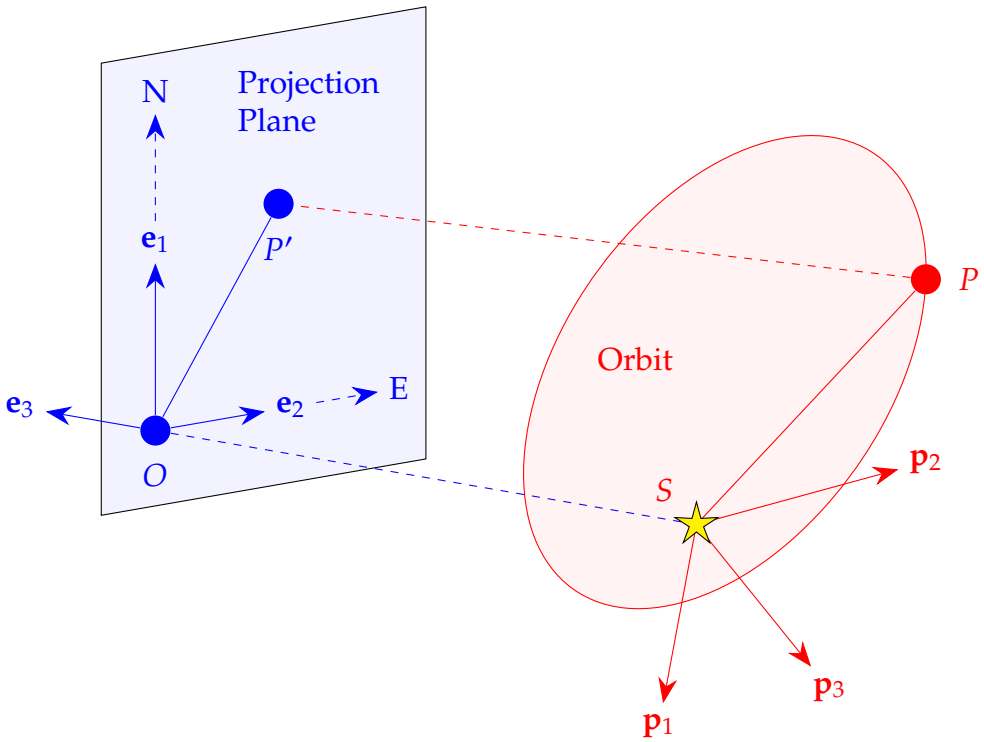

Figure 4.1: In the direct imaging of exoplanets, we observe the projection of the exoplanet's position P with respect to the star S onto the image plane.

4.1.1 Monte Carlo Methods

The majority of exoplanet fitting techniques in the literature are based on Monte Carlo methods, which are very versatile, though they tend to be computationally expensive. These include Bayesian Markov chain Monte Carlo (MCMC) methods, first applied to this problem by Ford [22, 21] and subsequently used by Rameau et al. [60], and Millar-Blanchaer et al. [51] among others. Others, such as Chauvin et al. [13], have used least-squares Monte Carlo (LSMC) methods, which were introduced by Press, Rybicki, and Hewitt. [58]

The following are the basic steps of MCMC for exoplanet orbit fitting based on the Metropolis-Hastings algorithm, as described in [52]. It works with Keplerian orbits, using a classical orbit parametrization (see Section 4.2.1).

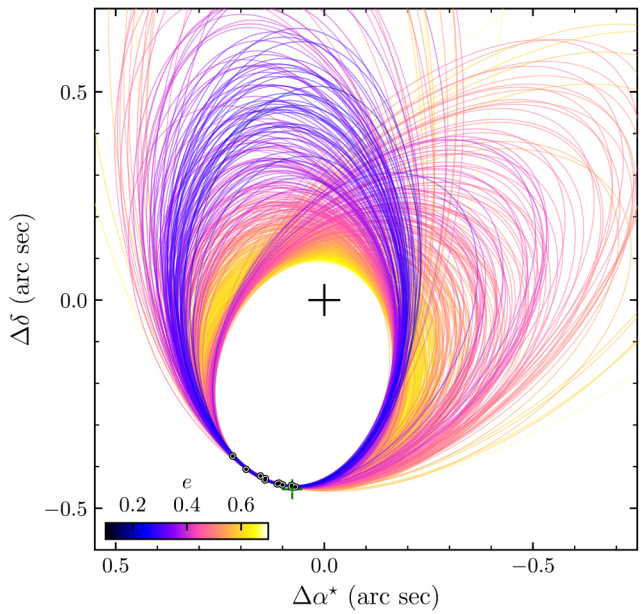

Figure 4.2: Orbit fits from direct imaging measurements for 51 Eridani b by De Rosa et al. using a Markov chain Monte Carlo method [16]. Here, as with most exoplanets, it is possible to fit a wide range of orbits to the same measurements. This is because the few measurements that are available have significant uncertainty and cover only a small fraction of the orbital period.

1. *The start of the chain*: Initial values are chosen for each free parameter.

2. *Generating the chain*:

   (a) For each parameter, a displacement is randomly drawn from the Gaussian centered on the current value of the parameter with fixed standard deviations.

   (b) Based on the ratio between the measurement likelihood of the current and trial parameters and a random number, the trial parameters are adopted or rejected.

4.1.2 Bayesian Rejection Sampling Methods

Another class of methods, such as those used by Konopacky et al. [39], Price-Whelan et al. [59], are based on rejection sampling. Here, we briefly describe Orbits for the Impatient, a Bayesian rejection sampling method for orbit fitting introduced by Blunt et al. in [7]. Like the MCMC method previously discussed, it works with Keplerian orbits and classical elements.

The basic OFTI algorithm consists of the following steps.

1. *Monte Carlo orbit generation from priors:* Trial orbits are drawn randomly from priors, i.e., initial sets of seven random orbital parameters are drawn from the prior probability distributions.
2. *Scale-and-rotate:* The trial orbits are adjusted to match the data, by scaling the semimajor axis and rotating. For each adjusted orbit, the algorithm solves Kepler's equation and computes a likelihood.
3. *Rejection sampling:* Each orbit is either accepted or rejected by comparing its likelihood to a uniform random number.

This process is repeated until the number of accepted orbits reaches a desired value. How quickly this happens depends on the acceptance rate, which is inversely proportional to the number of measurements. That is, if more orbit measurements are available, a much smaller set of orbits will fit the data and more sample orbits will be rejected. According to [7], the OFTI algorithm is most efficient for astrometry covering a short fraction of an orbit, typically less than 15% of the full orbital period. An open-source software library for running the OFTI algorithm, as well as a Markov chain Monte Carlo algorithm, is presented in [8].

4.1.3 Filtering Methods

All of the aforementioned orbit fitting methods are batch estimators; that is, they process all measurements simultaneously. Another approach is recursive estimation, or filtering, in which measurements are processed one by one, and the state estimate is gradually refined [5]. Savransky and Kasdin [64] took the latter approach, applying the extended Kalman filter to the problem of exoplanet orbit fitting. Due to the highly nonlinear dynamics and large measurement errors in this problem, they suggested that using an unscented filter could improve the orbit estimates. In this work, we follow their lead and develop a methodology for fitting exoplanet orbits based on direct observations using a higher-order unscented Kalman filter. This is a continuation of our work presented in [73].

4.2 Orbit Parametrizations and Measurements

This section starts with a review of the relations between the classical orbital elements, the Thiele-innes parameters, and the astrometric measurements, followed by a discussion of the Cohen-Hubbard nonsingular elements. We then present a new set of orbital elements, combining features of the Thiele-Innes parameters and the Cohen-Hubbard elements, which we developed specifically for exoplanet orbit fitting using filters. In addition, we present techniques for solving Kepler's equation in terms of the nonsingular elements. Finally, we describe how measurements of the host star's mass are related to the exoplanet's orbital elements.

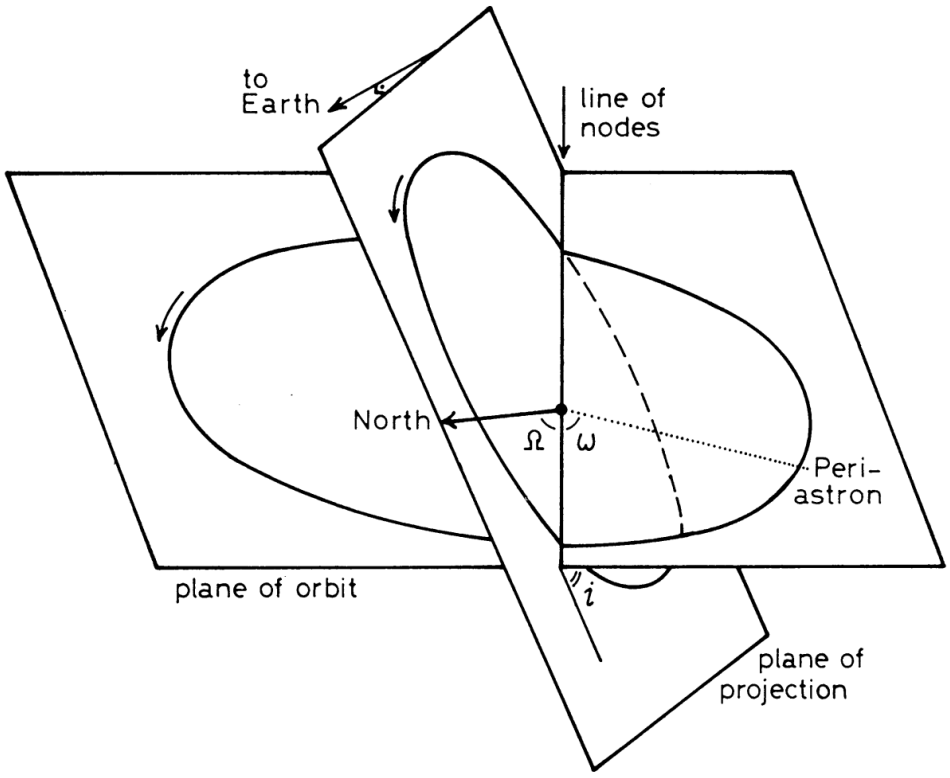

Figure 4.3: Orbital elements for binary stars, illustrated by Heintz [29]. The same definitions of the orbital elements are used for exoplanets. (This diagram denotes the inclination by i; in the rest of this work, we denote it by I, to avoid confusion with the imaginary unit.)

4.2.1 Classical Orbital Elements and the Measurement Model

Traditionally, Keplerian orbits are parametrized by the semi-major axis a, eccentricity e, inclination I, longitude of the ascending node Ω, argument of periapsis ω, mean anomaly at epoch M_0, and period P. The definitions of I, Ω, and ω used for exoplanet orbits are shown in Figure 4.3.

We measure the projected position of the exoplanet with respect to its host star in terms of the declination difference x and right ascension difference y, or alternatively, the angular separation ρ and the position angle θ. These two

parametrizations of the measurements are related by

$$x = \rho \cos\theta, \tag{4.1}$$

$$y = \rho \sin\theta. \tag{4.2}$$

To relate these measurements to Keplerian orbital motion, we use an approach that was originally developed for binary stars but is equally applicable to exoplanets. In addition to the Keplerian orbital elements, the astrometric measurements depend on the parallax ϖ of the host star. The stellar parallax is defined as the maximum angular distance between lines-of-sight from the Earth and the Sun to the star; because the parallax is very small, it can be accurately approximated as being inversely proportional to the star's distance from the Sun [6]. Here, it is useful to introduce the Thiele-Innes parameters, which are defined in terms of the classical elements as [29]

$$A = a(\cos\omega\cos\Omega - \sin\omega\sin\Omega\cos I), \tag{4.3}$$

$$B = a(\cos\omega\sin\Omega + \sin\omega\cos\Omega\cos I), \tag{4.4}$$

$$F = a(-\sin\omega\cos\Omega - \cos\omega\sin\Omega\cos I), \tag{4.5}$$

$$G = a(-\sin\omega\sin\Omega + \cos\omega\cos\Omega\cos I). \tag{4.6}$$

In terms of these quantities, the measurements are given by

$$x = \varpi(AX_{\mathcal{P}} + FY_{\mathcal{P}}), \tag{4.7}$$

$$y = \varpi(BX_{\mathcal{P}} + GY_{\mathcal{P}}), \tag{4.8}$$

where $(X_{\mathcal{P}}, Y_{\mathcal{P}})$ is the position of the planet relative to the star in the perifocal frame $\mathcal{P}$ scaled by $1/a$. Thus, the astrometric measurements are linear with respect to the Thiele-Innes parameters, making them a convenient fitting basis for orbit estimation.

As for the scaled perifocal positions $X_{\mathcal{P}}$ and $Y_{\mathcal{P}}$, they can be expressed most simply in terms of the eccentric anomaly E as

$$X_{\mathcal{P}} = \cos E - e, \tag{4.9}$$

$$Y_{\mathcal{P}} = \sqrt{1 - e^2} \sin E. \tag{4.10}$$

To relate the orbital position to time, we use the mean anomaly, which is defined as

$$M = n(t - t_p), \tag{4.11}$$

where t_p is the time of periapse passage and $n = 2\pi/P$ is the mean motion. This is equivalent to

$$M = M_0 + n(t - t_0), \tag{4.12}$$

where t_0 is the reference epoch. Finally, the eccentric anomaly and the mean anomaly are related by Kepler's equation

$$M = E - e \sin E, \tag{4.13}$$

which is a consequence of Kepler's second law.

Kepler's third law relates the mean motion or period to the semi-major axis; these elements obey the relation [55]

$$\mu = a^3 n^2 = \frac{4\pi^2 a^3}{P^2}, \tag{4.14}$$

where μ is the gravitational parameter, defined as

$$\mu = G M_{\text{tot}}, \tag{4.15}$$

where M_{tot} is the combined mass of the star and the exoplanet, and G is the universal gravitational constant.

4.2.2 The Cohen-Hubbard Nonsingular Elements

It is well known that the classical Keplerian elements produce a number of singularities, such as when $e = 0$, $I = 0$, and $I = 180°$. Thus, if an estimator uses the classical elements, it may behave badly at or near these singularities. Furthermore, while the Thiele-Innes parameters offer a significant improvement over the angles Ω, I, and ω for orbit fitting, they still do not remove the singularities associated with the eccentricity. Because of this, the fully nonsingular set of orbital elements introduced by Cohen and Hubbard [14] is of particular interest to us. These elements are defined as

$$q_0 = p^{\frac{1}{4}} \cos\left(\frac{I}{2}\right) \cos\left(\frac{\Omega + \omega + M_0}{2}\right), \tag{4.16}$$

$$q_1 = p^{\frac{1}{4}} \sin\left(\frac{I}{2}\right) \cos\left(\frac{\Omega - \omega - M_0}{2}\right), \tag{4.17}$$

$$q_2 = p^{\frac{1}{4}} \sin\left(\frac{I}{2}\right) \sin\left(\frac{\Omega - \omega - M_0}{2}\right), \tag{4.18}$$

$$q_3 = p^{\frac{1}{4}} \cos\left(\frac{I}{2}\right) \sin\left(\frac{\Omega + \omega + M_0}{2}\right), \tag{4.19}$$

$$e_X = e \cos M_0, \tag{4.20}$$

$$e_Y = -e \sin M_0, \tag{4.21}$$

where p is the semilatus rectum, which is related to a and e by

$$p = a(1 - e^2). \tag{4.22}$$

These elements are based on a reference frame, which we call $\mathcal{Q}$, that is rotatated from the perifocal frame $\mathcal{P}$ by the angle M_0 about the perifocal z-axis (the direction of the orbital angular momentum vector), as illustrated in Figure 4.4. The q elements, when normalized, form a quaternion representation of $\mathcal{Q}$ with respect to the celestial frame in which Ω, I, and ω are defined.

To compute the orbital position, Cohen and Hubbard use the angle φ defined as

$$\varphi = E - M_0, \tag{4.23}$$

which was proposed by Herget [14]. Substituting this into Kepler's equation (4.13) along with the relation (4.12) and the definitions (4.20–4.21) gives

$$\varphi = n(t - t_0) + e_X \sin\varphi - e_Y \cos\varphi. \tag{4.24}$$

Methods of solving for φ as a function of time are discussed in Section 4.2.4. The orbital position in Q scaled by $1/a$ is then given by

$$X_Q = \cos\varphi - e_X + e_Y s, \tag{4.25}$$

$$Y_Q = \sin\varphi - e_Y - e_X s, \tag{4.26}$$

where

$$s = \frac{e_X \sin\varphi - e_Y \cos\varphi}{1 + f}, \tag{4.27}$$

and

$$f = \sqrt{1 - e_X^2 - e_Y^2}. \tag{4.28}$$

4.2.3 Nonsingular Elements for Astrometric Orbit Fitting

For the purposes of exoplanet orbit fitting, in an attempt to combine the best features of the Cohen-Hubbard elements and the Thiele-Innes parameters, we

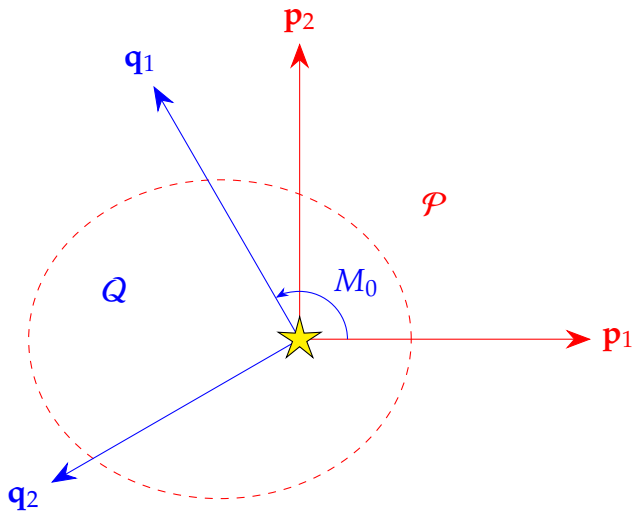

Figure 4.4: The perifocal frame $\mathcal{P}$ and the auxiliary frame $\mathcal{Q}$.

define a new set of nonsingular orbital elements:

$$\Xi_{11} = \varpi a(\cos(\omega + M_0)\cos\Omega - \sin(\omega + M_0)\sin\Omega\cos I), \tag{4.29}$$

$$\Xi_{21} = \varpi a(\cos(\omega + M_0)\sin\Omega + \sin(\omega + M_0)\cos\Omega\cos I), \tag{4.30}$$

$$\Xi_{12} = \varpi a(-\sin(\omega + M_0)\cos\Omega - \cos(\omega + M_0)\sin\Omega\cos I), \tag{4.31}$$

$$\Xi_{22} = \varpi a(-\sin(\omega + M_0)\sin\Omega + \cos(\omega + M_0)\cos\Omega\cos I), \tag{4.32}$$

$$\eta_1 = \frac{e\cos M_0}{\sqrt{1-e^2}}, \tag{4.33}$$

$$\eta_2 = -\frac{e\sin M_0}{\sqrt{1-e^2}}. \tag{4.34}$$

Since the gravitational parameter for an exoplanetary system is assumed not to be known *a priori*, Kepler's third law cannot be applied exactly, so the orbital period is considered to be independent of the semi-major axis. Therefore, we introduce an additional parameter

$$\lambda = \ln\left(\frac{P}{P_0}\right), \tag{4.35}$$

where P is the period, and P_0 is an arbitrary time scale. Throughout this work, we set P_0 to one year.

Properties of the Eccentricity-Related Parameters

We can recover e_X and e_Y from the η_i by

$$e_X = \frac{\eta_1}{\sqrt{1+\eta_1^2+\eta_2^2}}, \tag{4.36}$$

$$e_Y = \frac{\eta_2}{\sqrt{1+\eta_1^2+\eta_2^2}}. \tag{4.37}$$

Thus, for any $\boldsymbol{\eta} \in \mathbb{R}^2$, we have

$$e = \sqrt{e_X^2 + e_Y^2} = \sqrt{\frac{\eta_1^2+\eta_2^2}{1+\eta_1^2+\eta_2^2}} < 1, \tag{4.38}$$

which describes an elliptic orbit. Furthermore, we have $e = 0$ for $\boldsymbol{\eta} = 0$, and $e \to 1$ as $\|\boldsymbol{\eta}\| \to \infty$.

The Measurement Model

The Ξ_{ij} are closely related to the Thiele-Innes parameters, but are defined in frame $\mathcal{Q}$ rather than $\mathcal{P}$ and are scaled by the parallax. In other words, we take (4.3–4.6) and replace ω with $\omega + M_0$ and a with ϖa. Because of this scaling, the Ξ_{ij} are measured in angular units. Thus, the measurement model in Eqs. (4.7–4.8) can be modified to give

$$\mathbf{z} = \boldsymbol{\Xi}\boldsymbol{\zeta} + \mathbf{w}, \tag{4.39}$$

where

$$\boldsymbol{\zeta} = \begin{bmatrix} X_{\mathcal{Q}} \\ Y_{\mathcal{Q}} \end{bmatrix}, \tag{4.40}$$

in which $X_{\mathcal{Q}}$ and $Y_{\mathcal{Q}}$ are given by (4.25–4.26), and $\mathbf{w}$ is an additive noise term.

Recovering the Classical Elements

From (4.35), we can obtain the period,

$$P = P_0 \exp \lambda, \tag{4.41}$$

and from (4.33–4.34), we can solve for the eccentricity,

$$e = \sqrt{\frac{\eta_1^2 + \eta_2^2}{1 + \eta_1^2 + \eta_2^2}}, \tag{4.42}$$

and the mean anomaly at epoch,

$$M_0 = \text{atan2}(-\eta_2, \eta_1). \tag{4.43}$$

For the remaining elements, we develop formulae based on the ones for the Thiele-Innes parameters. To obtain Ω and ω, we can use the relations [29]

$$A + G = a(1 + \cos I)\cos(\Omega + \omega), \tag{4.44}$$

$$A - G = a(1 - \cos I)\cos(\Omega - \omega), \tag{4.45}$$

$$B - F = a(1 + \cos I)\sin(\Omega + \omega), \tag{4.46}$$

$$B + F = a(1 - \cos I)\sin(\Omega - \omega), \tag{4.47}$$

which give

$$\Omega + \omega = \text{atan2}(B - F, A + G), \tag{4.48}$$

$$\Omega - \omega = \text{atan2}(B + F, A - G). \tag{4.49}$$

Since the Ξ_{ij} replace ω with $\omega + M_0$, and the scaling by ϖ is independent of the angles, we have

$$\Omega + \omega + M_0 = \text{atan2}(\Xi_{21} - \Xi_{12}, \Xi_{11} + \Xi_{22}), \tag{4.50}$$

$$\Omega - \omega - M_0 = \text{atan2}(\Xi_{21} + \Xi_{12}, \Xi_{11} - \Xi_{22}). \tag{4.51}$$

Note that these solution leave an ambiguity of ±180° for Ω. This is not a fault of the parameterization, but rather of the measurements themselves: using only astrometry, we can determine the line of nodes (i.e., where the orbital plane intersects with a plane parallel to the image), but we cannot distinguish the ascending and descending node (i.e., at which node the object is moving towards or away from the observer). However, this ambiguity may be resolved with additional radial velocity measurements. When solving for Ω from the Thiele-Innes parameters alone, by convention, we take $0 \leq \Omega < 180°$, and we adopt the same convention for the Ξ_{ij}.

As for a and I, we can use the relations [29]

$$a^2(1 + \cos^2 I) = A^2 + B^2 + F^2 + G^2, \tag{4.52}$$

$$a^2 \cos I = AG - BF. \tag{4.53}$$

This is equivalent to

$$a^2(1 + \cos^2 I) = \|\mathbf{T}\|^2, \tag{4.54}$$

$$a^2 \cos I = \det \mathbf{T}, \tag{4.55}$$

where $\mathbf{T}$ is the matrix

$$\mathbf{T} = \begin{bmatrix} A & F \\ B & G \end{bmatrix}, \tag{4.56}$$

and $\|\cdot\|$ denotes the Frobenius norm. Then, defining

$$\gamma = \frac{\|\mathbf{T}\|^2}{2 \det \mathbf{T}}, \tag{4.57}$$

we obtain the quadratic equation

$$\cos^2 I - 2\gamma \cos I + 1 = 0. \tag{4.58}$$

The solution to this equation, with the sign chosen so that it lie in the interval $[-1, 1]$, is

$$\cos I = \gamma - \operatorname{sgn}(\gamma)\sqrt{\gamma^2 - 1}. \tag{4.59}$$

Finally, we can solve (4.54) to obtain a. An alternative expression for a is [6]

$$a^2 = u + \sqrt{(u+v)(u-v)}, \tag{4.60}$$

where

$$u = \frac{1}{2}\|\mathbf{T}\|^2, \tag{4.61}$$

$$v = \det \mathbf{T}. \tag{4.62}$$

We can use the same approaches to obtain a and I from $\boldsymbol{\Xi}$, with the substitutions

$$\|\mathbf{T}\|^2 = \frac{1}{\varpi^2}\|\boldsymbol{\Xi}\|^2, \tag{4.63}$$

$$\det \mathbf{T} = \frac{1}{\varpi^2} \det \boldsymbol{\Xi}. \tag{4.64}$$

These relations hold because $\boldsymbol{\Xi}$ has the same form as $\mathbf{T}$ except for the rotation by M_0 and the scaling by ϖ, but only the scaling affects the norm and the determinant. In addition, from (4.60), we have

$$\varpi^2 a^2 = u + \sqrt{(u+v)(u-v)}, \tag{4.65}$$

where

$$u = \frac{1}{2}\|\boldsymbol{\Xi}\|^2, \tag{4.66}$$

$$v = \det \boldsymbol{\Xi}. \tag{4.67}$$

4.2.4 Solutions of the Nonsingular Kepler Equation

Despite its simplicity, the classical Kepler equation (4.13) has no closed-form solution for E as a function of M; however, it can be reliably solved using Newton-

Raphson iteration, or alternatively using Fourier-Bessel series. The same is true for the equation's nonsingular counterpart (4.24), which can be expressed as

$$\varphi = \Delta M + e_X \sin\varphi - e_Y \cos\varphi, \tag{4.68}$$

where

$$\Delta M = n(t - t_0). \tag{4.69}$$

Here, we present analogous methods for solving for φ as a function of ΔM.

Newton-Raphson Iteration

Newton-Raphson iteration is a simple but often very efficient method for solving equations of the form $f(x) = 0$ for x, where f is a continuously differentiable function. We take an initial guess x_0 for x, and we iteratively refine the estimate by

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}. \tag{4.70}$$

This process is repeated until $|x_{n+1} - x_n| < \varepsilon$, where ε is a specified tolerance. Thus, we can solve for E given M using the iteration [76]

$$E_{n+1} = E_n + \frac{M - E_n + e\sin E_n}{1 - e\cos E_n}, \tag{4.71}$$

and for the initial guess, we can use

$$E_0 = M + e\sin M, \tag{4.72}$$

which is an approximation to the first order in e [55]. As for φ, we solve the equation

$$f(\varphi) = \Delta M + e_X \sin\varphi - e_Y \cos\varphi - \varphi = 0, \tag{4.73}$$

which gives the Newton-Raphson iteration

$$\varphi_{n+1} = \varphi_n - \frac{\Delta M + e_X \sin\varphi_n - e_Y \cos\varphi_n - \varphi_n}{e_X \cos\varphi_n + e_Y \sin\varphi_n - 1}, \tag{4.74}$$

and for the first estimate, we can express (4.72) in terms of the nonsingular variables to obtain

$$\varphi_0 = \Delta M + e_X \sin \Delta M - e_Y \cos \Delta M. \tag{4.75}$$

Fourier-Bessel Series

While iterative techniques are the most useful in practice, it is worth noting that we can obtain exact solutions to Kepler's equation using infinite series. (In practice, of course, these series must be truncated to a finite number of terms.) Since E is a periodic function of M, and therefore of time, we can express functions of E as Fourier series in M. The coefficients in these Fourier series give rise to the Bessel functions of the first kind, denoted by J_k, which Bessel represented by the integrals

$$J_k(z) = \frac{1}{\pi} \int_0^\pi \cos(z \sin \theta - k\theta) \mathrm{d}\theta \tag{4.76}$$

for integer k and complex z [31].

The Bessel functions have the following useful property [31]: if a, b, c, β, and γ are real numbers such that

$$c \exp(i\beta) = a - b \exp(-i\gamma), \tag{4.77}$$

where i is the imaginary unit, then

$$J_m(c) \exp(im\beta) = \sum_{l=-\infty}^{\infty} J_{m+l}(a) J_l(b) \exp(il\gamma). \tag{4.78}$$

In the special case $\gamma = \pi/2$, we have a somewhat simpler result: if a, b, c, and β are real numbers such that

$$a = c \cos \beta, \tag{4.79}$$

$$b = c \sin \beta, \tag{4.80}$$

then

$$J_k(c)\cos(k\beta) = \sum_{m=-\infty}^{\infty} (-1)^m J_{k+2m}(a) J_{2m}(b), \tag{4.81}$$

$$J_k(c)\sin(k\beta) = \sum_{m=-\infty}^{\infty} (-1)^m J_{k+2m+1}(a) J_{2m+1}(b). \tag{4.82}$$

For E as a function of M, Plummer [55] gives the Fourier-Bessel series

$$E = M + 2\sum_{k=1}^{\infty} \frac{1}{k} J_k(ke)\sin(kM). \tag{4.83}$$

Substituting $\varphi = E - M_0$ and $M = M_0 + \Delta M$ and applying trigonometric identities, we can obtain a similar series for φ, namely

$$\varphi = \Delta M + \sum_{k=1}^{\infty} (a_k \cos(k\Delta M) + b_k \sin(k\Delta M)), \tag{4.84}$$

where the coefficients a_k and b_k are given by

$$a_k = \frac{2}{k} J_k(ke)\sin(kM_0), \tag{4.85}$$

$$b_k = \frac{2}{k} J_k(ke)\cos(kM_0). \tag{4.86}$$

To express the coefficients in terms of e_X and e_Y rather than e and M_0, we use the definitions (4.20–4.21) and the Bessel function identities (4.81–4.82) to obtain

$$a_k = -\frac{2}{k} \sum_{m=-\infty}^{\infty} (-1)^m J_{k+2m+1}(e_X) J_{2m+1}(e_Y), \tag{4.87}$$

$$b_k = \frac{2}{k} \sum_{m=-\infty}^{\infty} (-1)^m J_{k+2m}(e_X) J_{2m}(e_Y). \tag{4.88}$$

Using a similar approach, we can obtain a series expansion for $\exp(ik\varphi)$, where k is an integer. In the classical elements, for $\exp(ikE)$ as a function M, Plummer [55] gives the series

$$\exp(ikE) = \sum_{m=-\infty}^{\infty} A_{km} \exp(imM), \tag{4.89}$$

with the coefficients

$$A_{km} = \begin{cases} \frac{k}{m} J_{m-k}(me), & m \neq 0 \\ 1, & m = k = 0 \\ -\frac{1}{2}e, & m = 0,\ k = \pm 1 \\ 0, & m = 0,\ |k| > 1 \end{cases}. \tag{4.90}$$

Now, we wish to find the coefficients C_{km} of an analogous series for $\exp(ik\varphi)$ of the form

$$\exp(ik\varphi) = \sum_{m=-\infty}^{\infty} C_{km} \exp(im\Delta M). \tag{4.91}$$

Again using the relations $\varphi = E - M_0$ and $M = M_0 + \Delta M$, we have

$$C_{km} = A_{km} \exp(i(m-k)M_0), \tag{4.92}$$

and substituting (4.90) gives

$$C_{km} = \begin{cases} \frac{k}{m} J_{m-k}(me) \exp(i(m-k)M_0), & m \neq 0 \\ 1, & m = k = 0 \\ -\frac{1}{2}e \exp(-ikM_0), & m = 0,\ k = \pm 1 \\ 0, & m = 0,\ |k| > 1 \end{cases}. \tag{4.93}$$

For the case $m = 0$ and $k = \pm 1$, we can apply Euler's formula and the definitions of e_X and e_Y (4.20–4.21) to obtain

$$C_{km} = -\frac{1}{2}(e_X + ike_Y). \tag{4.94}$$

As for $m \neq 0$, we can apply the general property (4.78) of Bessel functions: in this case, we have

$$me \exp(iM_0) = me_X - ime_Y, \tag{4.95}$$

and it follows that

$$J_{m-k}(me)\exp(i(m-k)M_0) = \sum_{l=-\infty}^{\infty}(-i)^l J_{m-k+l}(me_X)J_l(me_Y). \tag{4.96}$$

Thus, the coefficients for the series (4.91) can be expressed in terms of the nonsingular parameters e_X and e_Y as

$$C_{mk} = \begin{cases} \frac{k}{m}\sum_{l=-\infty}^{\infty}(-i)^l J_{m-k+l}(me_X)J_l(me_Y), & m \neq 0 \\ 1, & m = k = 0 \\ -\frac{1}{2}(e_X + ike_Y) & m = 0,\ k = \pm 1 \\ 0, & m = 0,\ |k| > 1 \end{cases}. \tag{4.97}$$

Convergence

We numerically tested the convergence of both iterative and series solutions for φ as a function of ΔM. For the Newton-Raphson iteration (4.74), we counted the number of steps needed to achieve a tolerance $\varepsilon = 10^{-9}$. For the Fourier-Bessel series (4.84), we found the number of terms K required to achieve the same tolerance, with the series (4.84) truncated to $k \leq K$ and (4.87–4.88) to $-K \leq m \leq K$. Since the nonsingular form of Kepler's equation depends on e, M_0, and ΔM, we tested each method with fixed M_0 and varying ΔM, and vice versa. In both cases, we varied the eccentricity, with $0 \leq e \leq 0.99$, covering almost the entire range for elliptic orbits.

Figures 4.5 and 4.6 show the convergence results for the Newton-Raphson iteration. In all cases, the desired tolerance was achieved with fewer than nine iterations, and for $e < 0.2$, with three or fewer. As for the Fourier-Bessel series, for which the convergence results are shown in Figures 4.7 and 4.8, between

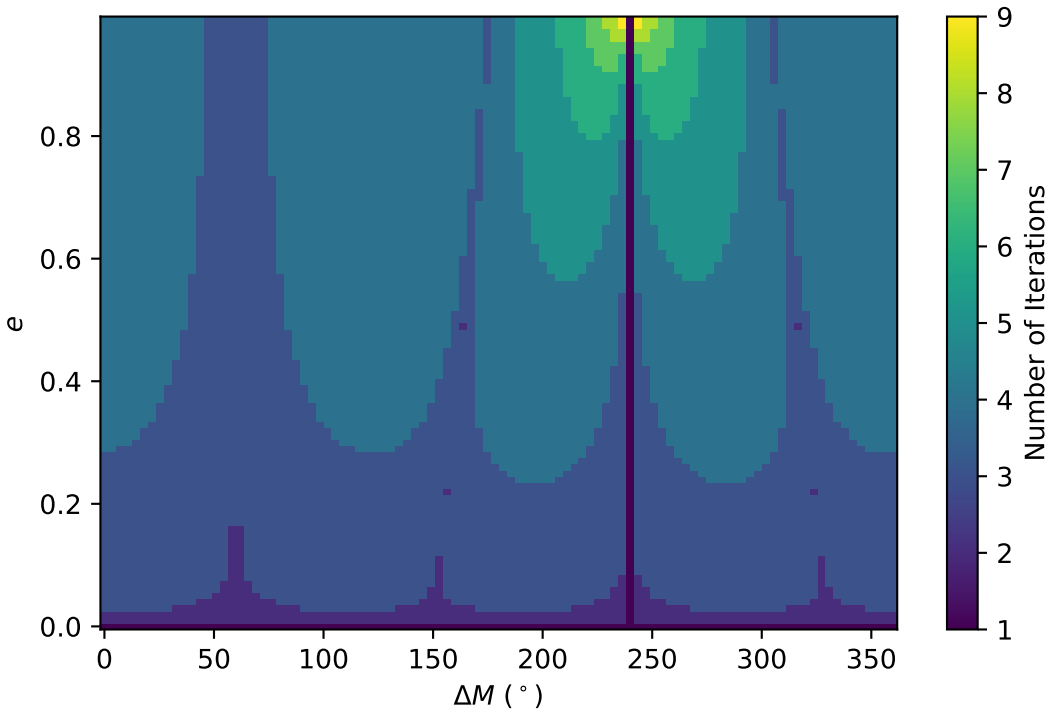

Figure 4.5: Convergence of Newton-Raphson iteration for φ with $M_0 = 120°$.

three and five series terms were required in all cases. These results indicate that both methods converge quickly.

4.2.5 Mass Measurements

As noted in Section 4.2.3, we cannot apply Kepler's third law exactly to the problem of exoplanet orbit fitting, since the gravitational parameter of the system is not known *a priori*. However, we do have some prior knowledge of the gravitational parameter, based on estimates of the star's mass. Here, we treat the mass of the star as M_{tot}, since the mass of the planet is assumed to be much smaller and well within the margin of error for the star's mass.

A rough estimate of a star's mass M can be obtained from its luminosity L,

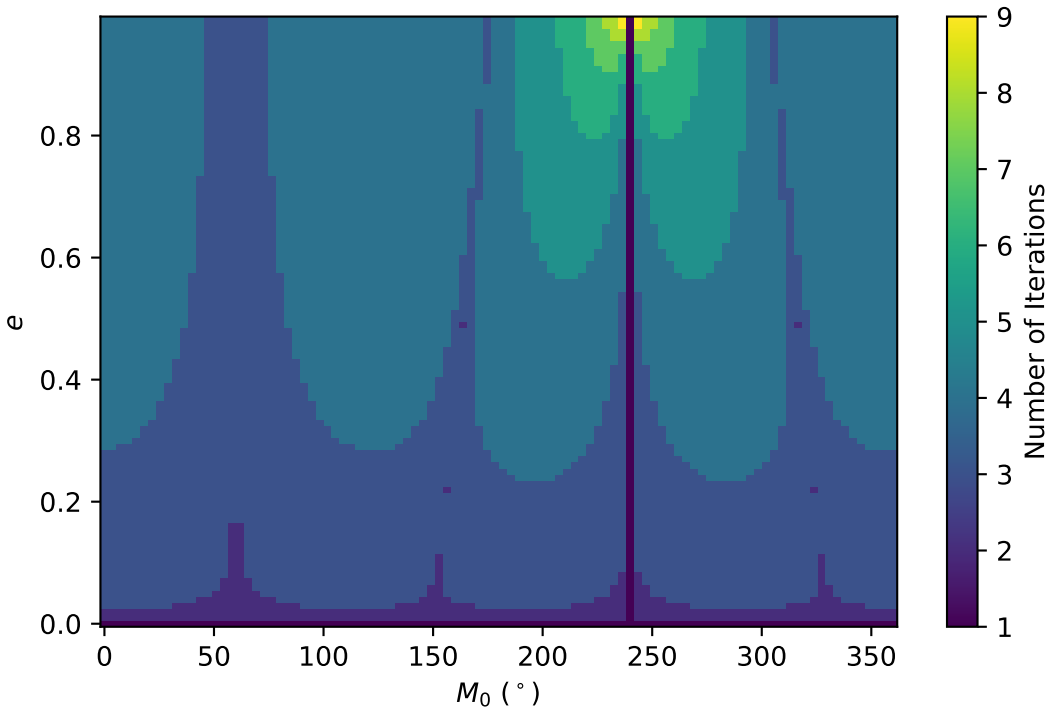

Figure 4.6: Convergence of Newton-Raphson iteration for φ with $\Delta M = 120°$.

using a mass-luminosity relation of the form

$$L = L_\odot \left(\frac{M}{M_\odot} \right)^\alpha , \tag{4.98}$$

where $L_\odot$ and $M_\odot$ are the luminosity and mass of the Sun, and α is an empirical constant. For main-sequence stars, $3 \lesssim \alpha \lesssim 4$ [28].

For the purposes of orbit fitting, we account for estimates of the star's mass using a measurement that is a function of the parallax and the total mass, defined as

$$z_{\mathrm{PM}} = \ln \left(\frac{M_{\mathrm{tot}} \varpi^3}{M_\odot \varpi_0^3} \right), \tag{4.99}$$

where ϖ_0 is a constant having units of parallax (e.g., $\varpi_0 = 1$ as). To express this measurement in terms of the orbital elements, we substitute (4.14) and (4.15) to

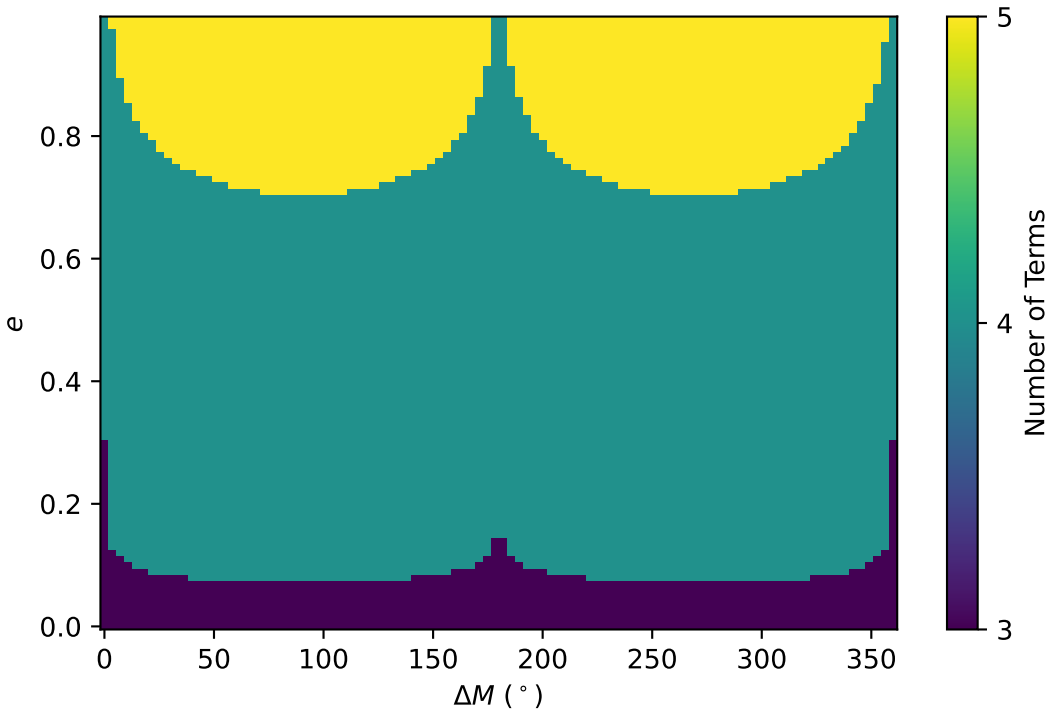

Figure 4.7: Convergence of Fourier-Bessel series for φ with $M_0 = 120°$.

obtain

$$z_{\mathrm{PM}} = \ln\left(\frac{4\pi^2\varpi^3 a^3}{GM_\odot P^2 \varpi_0^3}\right), \tag{4.100}$$

and substituting (4.35) gives

$$z_{\mathrm{PM}} = \ln\left(\frac{4\pi^2}{GM_\odot P_0^2 \varpi_0^3}\right) + \frac{3}{2}\ln(\varpi^2 a^2) - 2\lambda, \tag{4.101}$$

where the second term depends only on $\boldsymbol{\Xi}$ and can be evaluated using (4.65). Finally, given the mean values $(\bar{M}, \bar{\varpi})$ and standard deviations $(\sigma_M, \sigma_\varpi)$ of M_{tot} and ϖ, we have a first-order estimate of the variance of z_{PM}, given by

$$\sigma_{\mathrm{PM}}^2 = \left(\frac{\sigma_M}{\bar{M}}\right)^2 + \left(\frac{3\sigma_\varpi}{\bar{\varpi}}\right)^2. \tag{4.102}$$

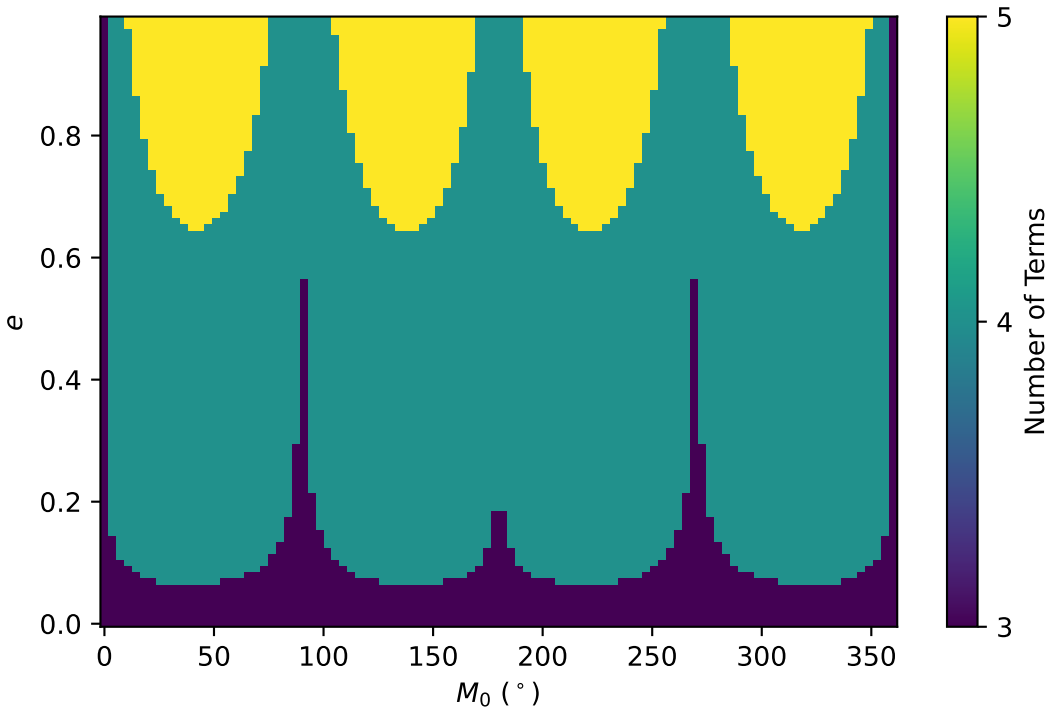

Figure 4.8: Convergence of Fourier-Bessel series for φ with $\Delta M = 120°$.

4.3 The Nonsingular Estimator for Exoplanet Orbits

In this section, we present the Nonsingular Estimator for Exoplanet Orbits (NEXO), a new method for exoplanet orbit fitting that combines the unscented Kalman filter with a batch estimator. It applies Gaussian mixture models, described in Section 1.4.4, with the the square-root sigma point filter (SRSPF), described in Section 1.4.3. The use of Gaussian mixtures was suggested to us by Trevor Wolf. Furthermore, we compute the sigma points using a fifth-order cubature rule for Gaussian distributions, and we run the filter multiple times with the same measurements to refine the estimates.

For the fitting basis, we use the nonsingular orbital elements described in

Section 4.2.3, and we define the state vector as

$$\mathbf{x} = \begin{bmatrix} \lambda & \eta_1 & \eta_2 & \Xi_{11} & \Xi_{21} & \Xi_{12} & \Xi_{22} \end{bmatrix}^{\mathrm{T}}. \tag{4.103}$$

For an unperturbed Keplerian orbit, this state vector is constant in time. Here, we assume that all perturbations on the exoplanet's orbit are negligible. Therefore, we only use the correction step of the filter, and there is no need for the prediction step.

4.3.1 Motivation for Gaussian mixtures

Among the components of $\mathbf{x}$, the strongest nonlinearities in the astrometric measurements are due to variations in the log-period λ. As shown in Section 4.2.3, the astrometric measurements are linear in $\boldsymbol{\Xi}$, and their dependence on $\boldsymbol{\eta}$ can be approximated well by low-order polynomials. On the other hand, λ causes oscillations, and the frequency of these oscillations increases exponentially as $\lambda \to -\infty$. Thus, if the variance of λ is large, the unscented transform does not accurately represent the distribution of the measurements. Because of this, our early tests found the plain SRSPF to be highly sensitive to the prior distribution parameters.

While the nonlinearities in λ can be problematic for a simple unscented transform, they can be handled much more robustly using Gaussian mixtures. Specifically, we can represent a diffuse distribution of λ using a mixture of Gaussian distributions with different values for the mean of λ and small values for the variance of λ. Then, because of the small variances, each mixture component effectively covers only a small range of λ values, and the dependence of the measurements on λ can be accurately approximated by low-order polynomials.

Thus, the unscented transform for each component of the mixture can represent the measurement distribution very accurately. In addition, since the unscented transform is accurate for $\boldsymbol{\eta}$ and $\boldsymbol{\Xi}$ even when their covariances are large, their distributions do not have to be "broken up" using Gaussian mixtures.

4.3.2 Sigma Points

To obtain more accurate approximations of $\bar{\mathbf{z}}$, $\mathbf{P}_{zz}$, and $\mathbf{P}_{xz}$ for each component of the mixture, we use a set of sigma points that matches statistical moments up to the fifth order for Gaussian distributions. It is based on Stroud's fifth-order cubature rule of dimension $n \geq 3$ for the weighting function $\exp(-\|\mathbf{x}\|^2)$ with $N_S = 2^n + 2n$ points [74]. Given Stroud's points $\tilde{\mathbf{x}}^{(j)}$ and weights $\tilde{w}_j$, we compute the sigma points

$$\mathbf{x}^{(j)} = \bar{\mathbf{x}} + \sqrt{2\mathbf{P}_{xx}}\tilde{\mathbf{x}}^{(j)}, \tag{4.104}$$

and the corresponding weights

$$w_j = \frac{\tilde{w}_j}{\sum_{j=1}^{N_S} \tilde{w}_j}. \tag{4.105}$$

Then, the $\mathbf{x}^{(j)}$ and w_j form a fifth-order cubature rule, whose weighting function is a Gaussian probability density with mean $\bar{\mathbf{x}}$ and covariance $\mathbf{P}_{xx}$.

4.3.3 The Prior Distribution

For the prior distribution of λ, we assume that it is uniform on the interval $[\ln(P_{\min}/P_0), \ln(P_{\max}/P_0)]$, where $P_{\min}$ and $P_{\max}$ are minimum and maximum values for the period, chosen to cover the entire range of realistic values of

P. Then, we form the prior Gaussian mixture based on a Gauss-Legendre quadrature rule with N nodes on this interval: for the prior mean values of λ for the components of the mixture, we take the abscissae of the quadrature rule, and we take the corresponding quadrature weights to be the prior weights of the mixture components. This is effectively treating the prior mean of λ as a nuisance parameter and integrating over it using a quadrature approximation. As for the prior covariances of the mixture components, we take

$$\sqrt{\mathbf{P}_{xx}} = \mathrm{diag}(\sigma_\lambda, \sigma_\eta, \sigma_\eta, \sigma_\Xi, \sigma_\Xi, \sigma_\Xi, \sigma_\Xi), \tag{4.106}$$

where σ_λ, σ_η, and σ_Ξ are the same for all components of the mixture.

With this form of the prior distribution, we have to choose the parameters $P_{\min}$, $P_{\max}$, σ_η, and σ_Ξ. As discussed in Chapter 1, the purpose of the prior distribution is to represent our knowledge of the system before any measurements are obtained. This is particularly difficult for exoplanets; for an individual exoplanet, the prior distribution is necessarily diffuse. However, we do have some prior knowledge of exoplanet populations. Several such population models are implemented in EXOSIMS, a simulation framework for exoplanet imaging missions developed by Savransky et al. in the Space Imaging and Optical Systems (SIOS) Lab [17, 63]. This open-source library is available on GitHub. We use the "Kepler-like" distributions of orbital elements and masses for exoplanets from EXOSIMS, which are based on the exoplanets detected using transit photometry methods by the Kepler space telescope [38]. Based on the statistics of 10,000 orbits sampled from the "Kepler-like" population model, we chose $P_{\min} = 0.1$ yr, $P_{\max} = 10,000$ yr, $\sigma_\lambda = 0.5$ and $\sigma_\eta = 0.15$, with $N = 1000$ Gaussian mixture components. In addition, for the semi-major axis, we chose $\sigma_a = 100$ au. Since $|\Xi_{ij}| \leq \varpi a$, we set $\sigma_\Xi = \bar{\varpi}\sigma_a$, where $\bar{\varpi}$ is defined in Section 4.2.5.

4.3.4 Repeated Filtering

We found that running the filter repeatedly with the same measurements, stepping through the measurements forward in time at each iteration, increased the accuracy of its estimates. Specifically, we found that running 10 such filter passes gave lower root-mean-square error values than simply running the filter once or using a smoothing procedure (i.e., running the filter forward in time, then backward, then forward again). Because of this, we ran 10 filter passes in all of the examples in this work. However, determining the optimal number of filter passes will require further testing and analysis.

4.4 Validation

To test the performance of our estimation method over a wide range of possible orbits where the "ground truth" is known, we used large sets of synthetic orbital parameters and astrometric measurements.

4.4.1 Generation of Sample Orbits and Measurements

Using the "Kepler-like" population model from EXOSIMS, we generated 100 sets of random orbital elements and planet masses, with the mass of the host star fixed at one solar mass. Then, for each orbit, we computed five pairs of astrometric measurements (x, y), evenly spaced in time over two years, using the *orbitize!* library [8]. Then, we added simulated measurement errors sampled from a central Gaussian distribution with a standard deviation of 5 mas.

4.4.2 Testing and Comparison of Estimation Methods

We ran NEXO, MCMC, and OFTI with the 100 sets of synthetic measurements. With MCMC and OFTI, we attempted to sample 10,000 sample orbits for each set of measurements. If an estimator took longer than 15 minutes to fit one orbit, the run was stopped. This occurred in 13% of cases for OFTI, but in no case for NEXO and MCMC. We compared the three estimators based on the credible intervals that they generated for the classical orbital elements. For NEXO, we computed the mean estimates and credible intervals using the unscented transform and the Chebyshev and Markov inequalities, as described in Section 1.4.5. As for MCMC and OFTI, we applied the same method, but slightly modified to use their outputs: instead of sigma points, we used the 10,000 generated orbits with equal weights. We computed non-angular credible intervals for a, e, and P, and angular credible intervals for I, Ω, ω, and M_0.

Figures 4.9–4.15 show the means and 95% credible intervals for each of the seven classical elements using the three estimation methods. The accuracy of the estimates varies greatly between the orbital parameters for all three estimation techniques. This is likely due to the varying observability of the elements from the measurements.

The estimates of the semi-major axis are shown in Figure 4.9. This parameter is highly observable, since it is closely related to the measured angular separation. Because of this, the expected values of a obtained using all three methods follow the true values closely. However, the methods appear to have different biases: NEXO and MCMC tend to underestimate a, while OFTI often overestimates it. As for the credible intervals, they are very wide for MCMC and even wider for OFTI, in some cases covering the entire range of a values of the gener-

ated orbits. On the other hand, the credible intervals are too narrow for NEXO, since many of them do not even contain the true value. In this regard, none of the three methods generates credible intervals that are informative. This trend—that the credible intervals are too wide for MCMC and OFTI and too narrow for NEXO—is present for all seven orbital elements.

Figure 4.10 shows the estimates of the eccentricity, which is more difficult to observe than the semi-major axis, especially when the orbit is nearly circular. For this parameter, NEXO gives significantly more accurate estimates than MCMC or OFTI. A possible explanation is that the process of filtering, which recursively refines the estimates, can better capture the effects of e on the measurements than methods based on random sampling.

The estimates of the inclination are presented in Figure 4.11. The most noticeable feature in the MCMC and OFTI estimates is that many of them have an expected value near 90°. However, this does not necessarily mean that these estimators favor edge-on orbits: in particular, a uniform distribution for I has a mean of 90°. In many of the cases where the inclination estimate is near 90°, the credible interval covers the entire range of I from 0 to 180°, which suggests that the distribution could be close to uniform. A similar but less prominent phenomenon appears for the longitude of the ascending node, which has the same range, shown in Figure 4.12. NEXO, on the other hand, does not appear to have this bias, and it gives much more accurate estimates of both I and Ω overall.

Perhaps the least observable orbital elements are the argument of periapsis and the mean anomaly at epoch. Consequently, none of the three estimators gives accurate results for these two parameters, as shown in Figures 4.13 and 4.14. For both ω and M_0, MCMC exhibits a bias towards 180°, with credible intervals

covering most of the range from 0 to 360°; as with I and Ω, this could be due to nearly uniform distributions. However, this does not occur for NEXO or OFTI.

The estimates of the period, shown in Figure 4.15, exhibit similar trends to those of the semi-major axis, shown in Figure 4.15. This is not surprising, since a and P are related by Kepler's third law (4.14), and the estimators are given a mean and standard deviation for M_{tot}. Despite the uncertainty in M_{tot}, estimating the orbital period would be much more difficult without this information. For example, some tests of an early variant of NEXO without the mass measurements described in Section 4.2.5 produced orbits in which multiple periods elapsed between measurements. Of course, such orbits are extremely unlikely.

Table 4.1 shows the root-mean-square error (RMSE) values of the mean estimates for each orbit element and estimation method. For most of the parameters, NEXO has a lower RMSE than MCMC or OFTI; the only exceptions are a, for which MCMC has a slightly lower error, and M_0 for which OFTI has a lower error. In particular, NEXO has a significantly lower RMSE than MCMC or OFTI for the eccentricity and the period.

Table 4.2 shows percentages of cases where the true values lie within the 95% credible intervals. Unlike for MCMC or OFTI, these values are significantly lower than 95% for NEXO. This suggests that NEXO can become "overconfident" in its estimates.

Table 4.3 shows the average run time for successful orbit fits for each of the three estimators, excluding the cases where OFTI was terminated because it failed to generate 10,000 orbits in 15 minutes. NEXO is over times 40 faster than OFTI, which in turn is faster than MCMC.

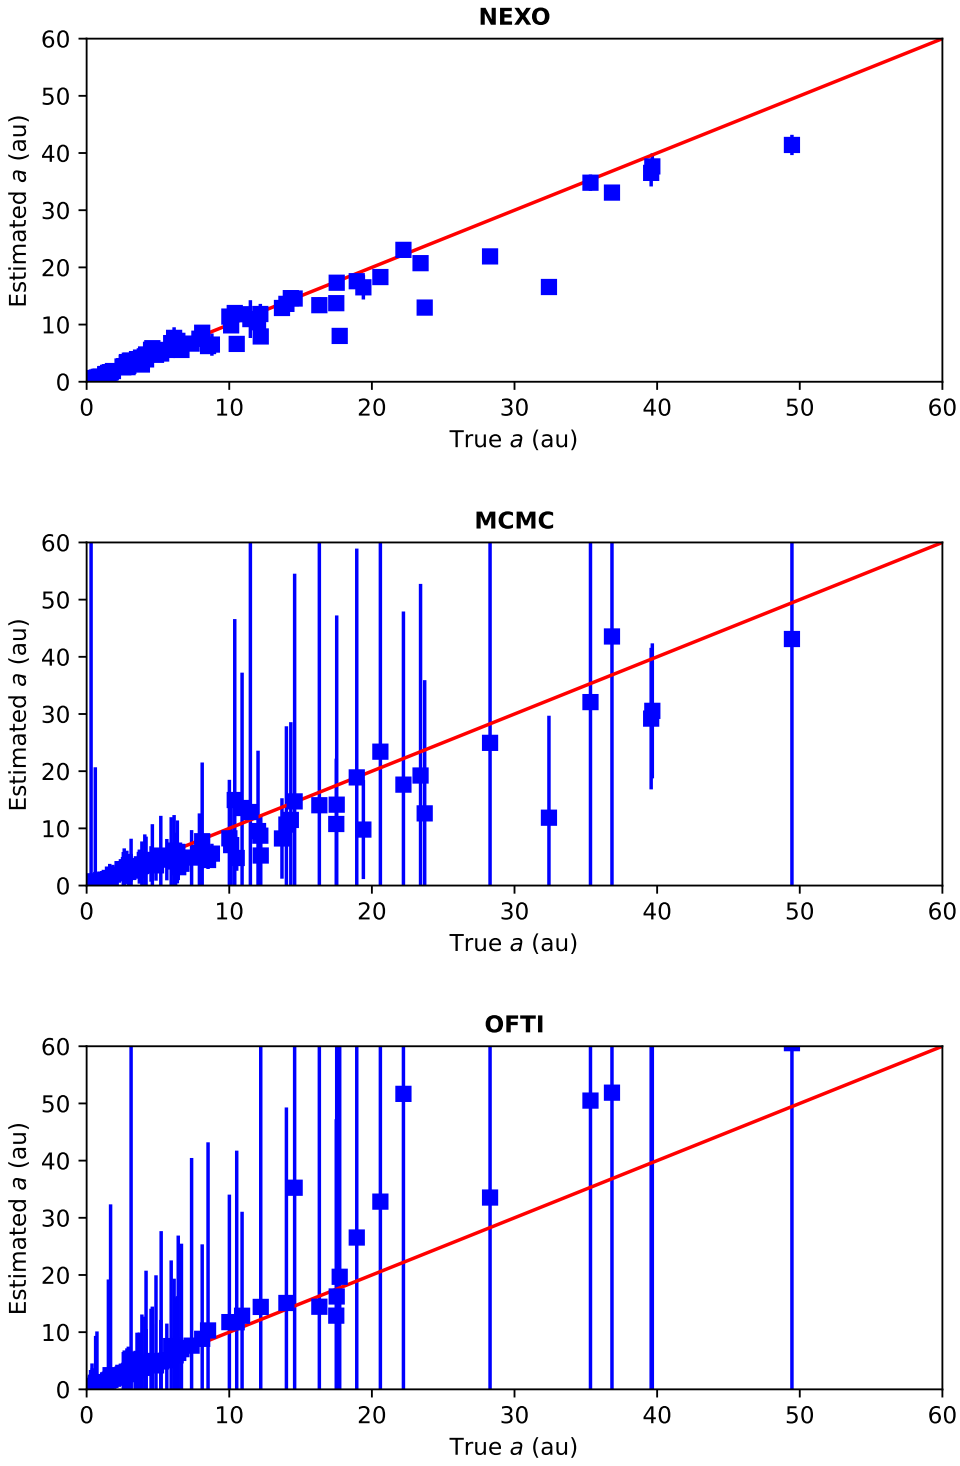

Figure 4.9: Credible intervals for the semi-major axis. All three methods give accurate central estimates. However, the credible intervals for NEXO are much narrower than for MCMC or OFTI, and they sometimes fail to include the true value.

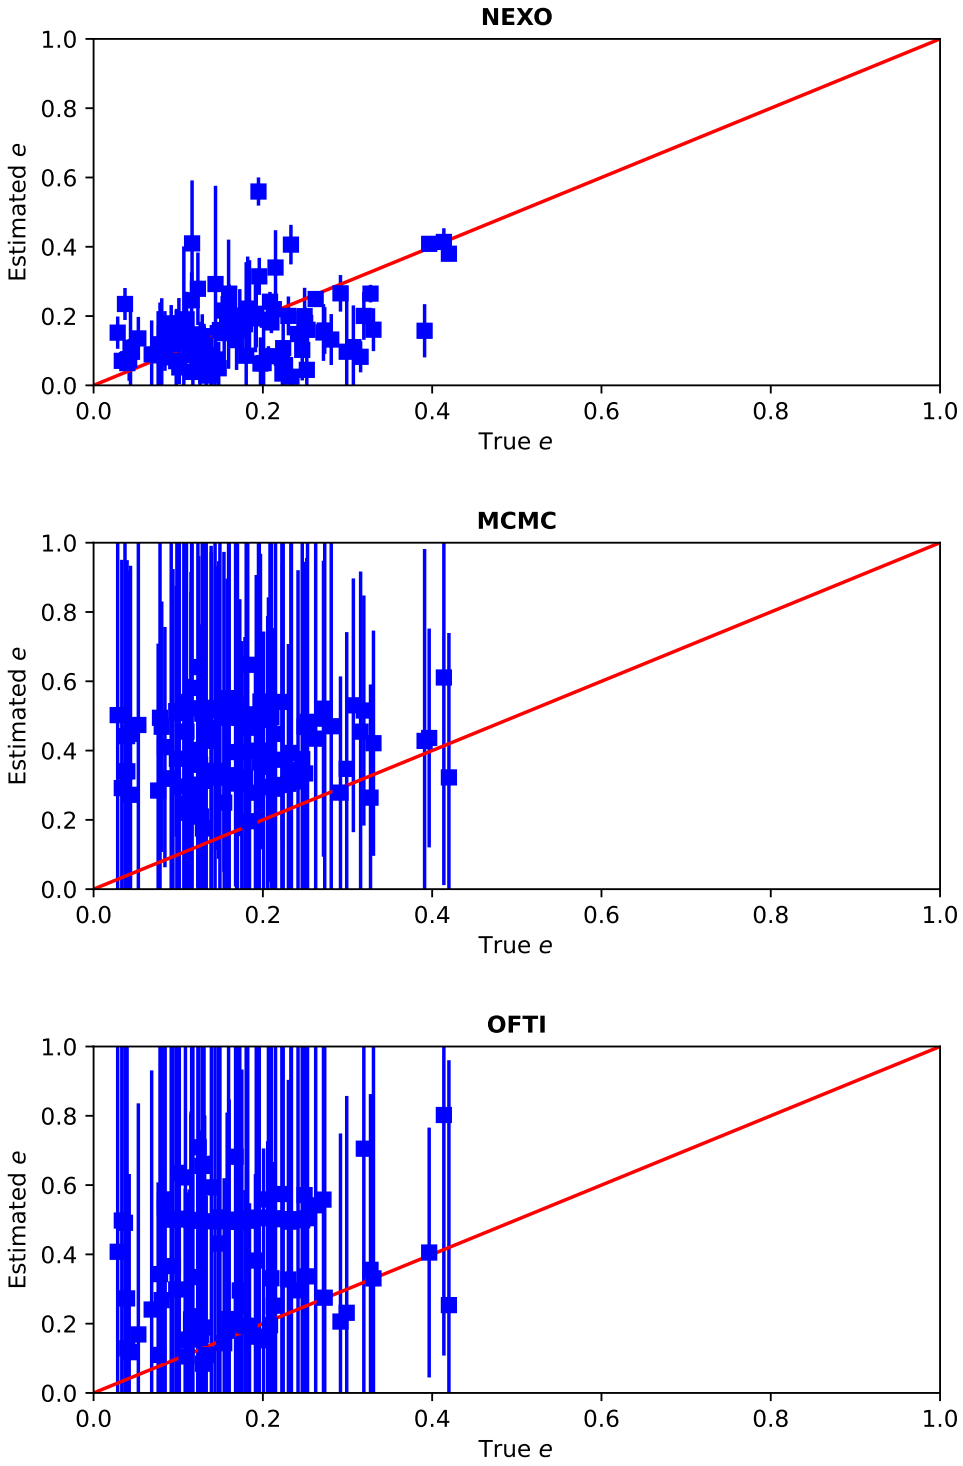

Figure 4.10: Credible intervals for the eccentricity. NEXO gives the best estimates overall; MCMC tends to overestimate it, while the eccentricity errors for OFTI appear to be spread out more or less uniformly.

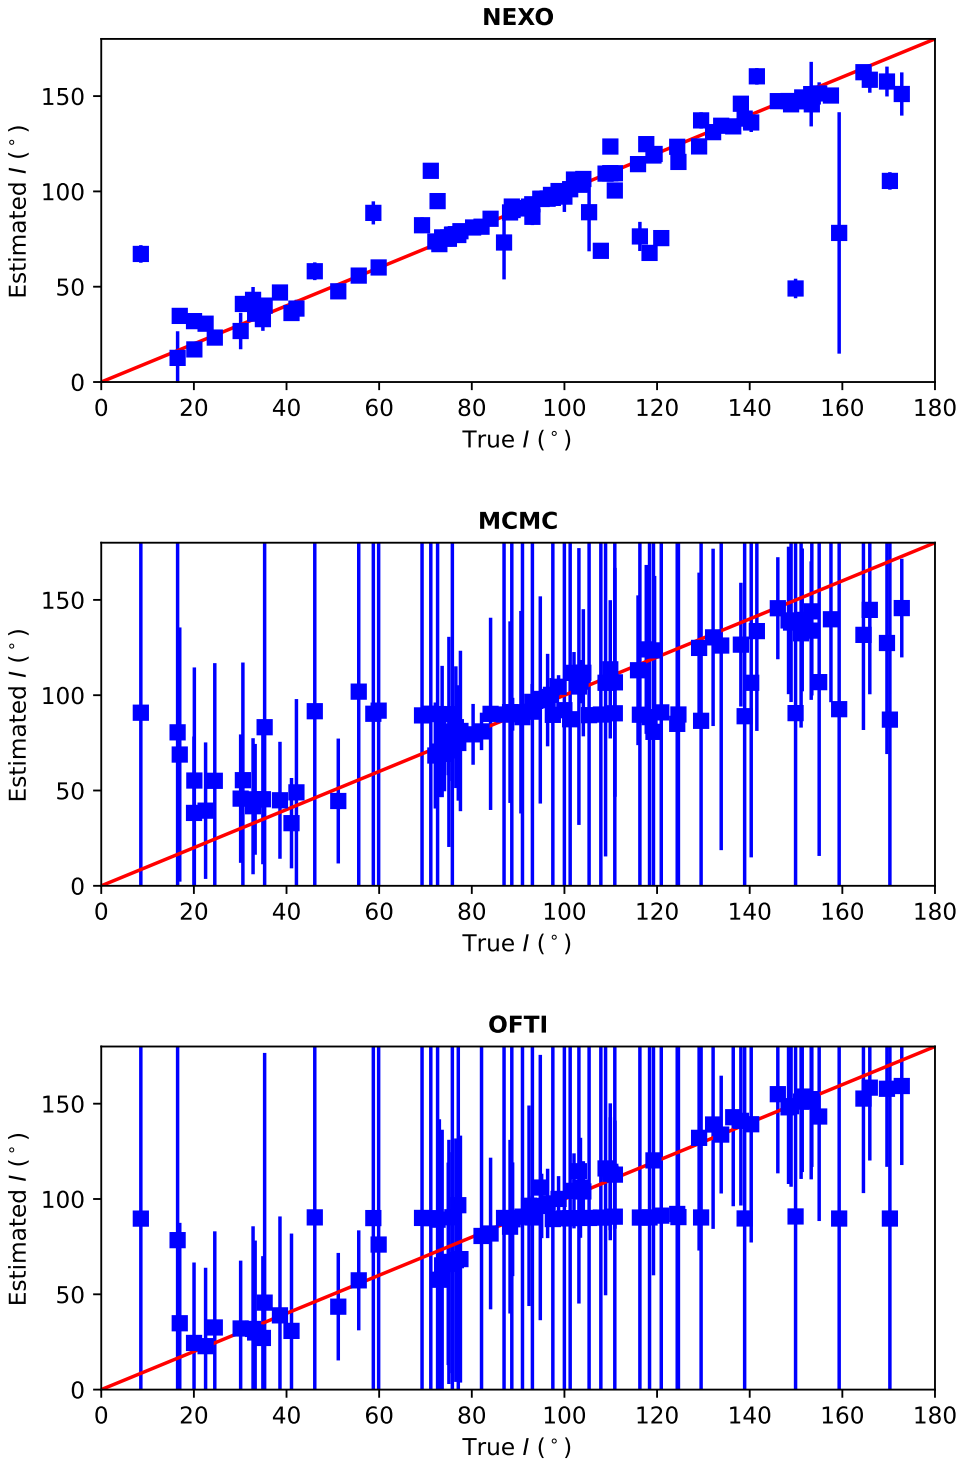

Figure 4.11: Credible intervals for the inclination. NEXO gives the most accurate estimates, while MCMC and OFTI exhibit a strong bias towards edge-on orbits, or $I = 90°$.

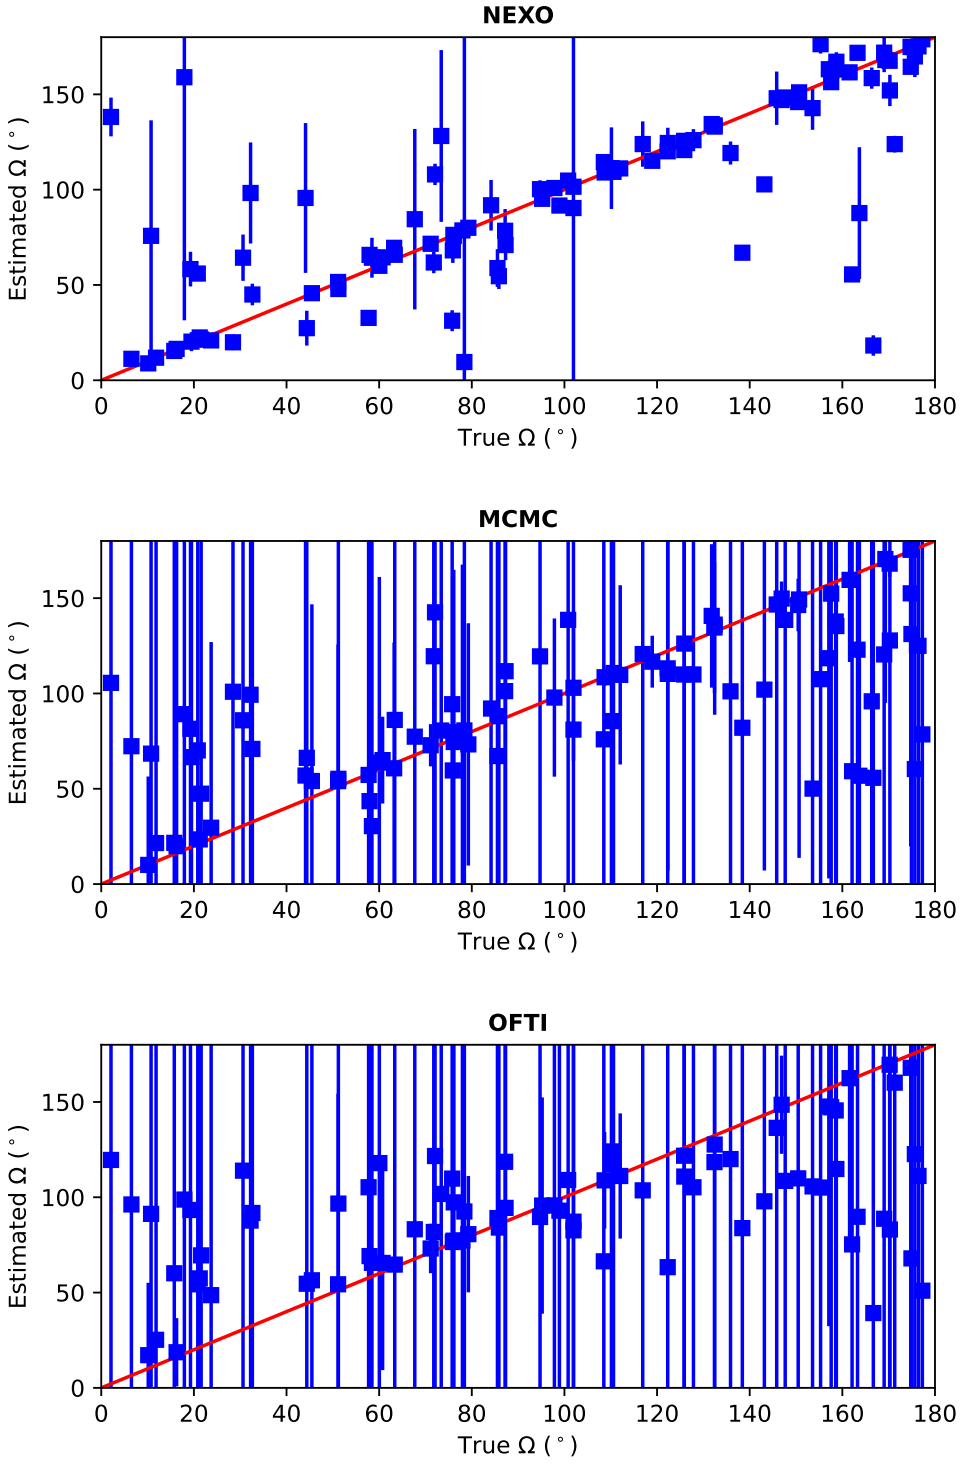

Figure 4.12: Credible intervals for the longitude of the ascending node. NEXO gives the most accurate estimates, followed by MCMC.

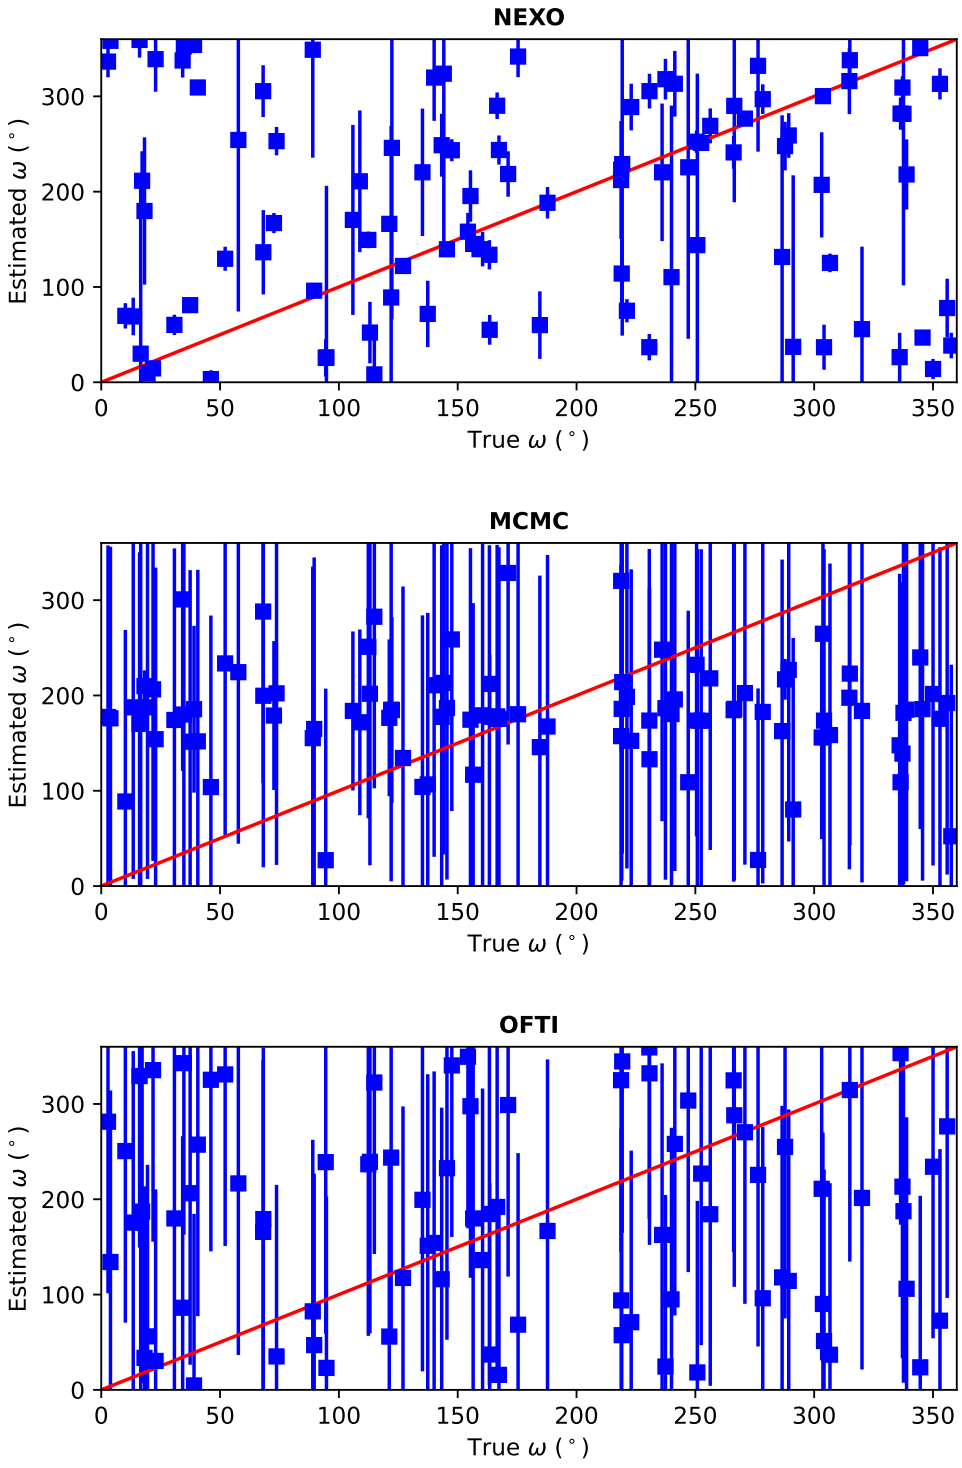

Figure 4.13: Credible intervals for the argument of periapsis. None of the three methods gives reliable estimates.

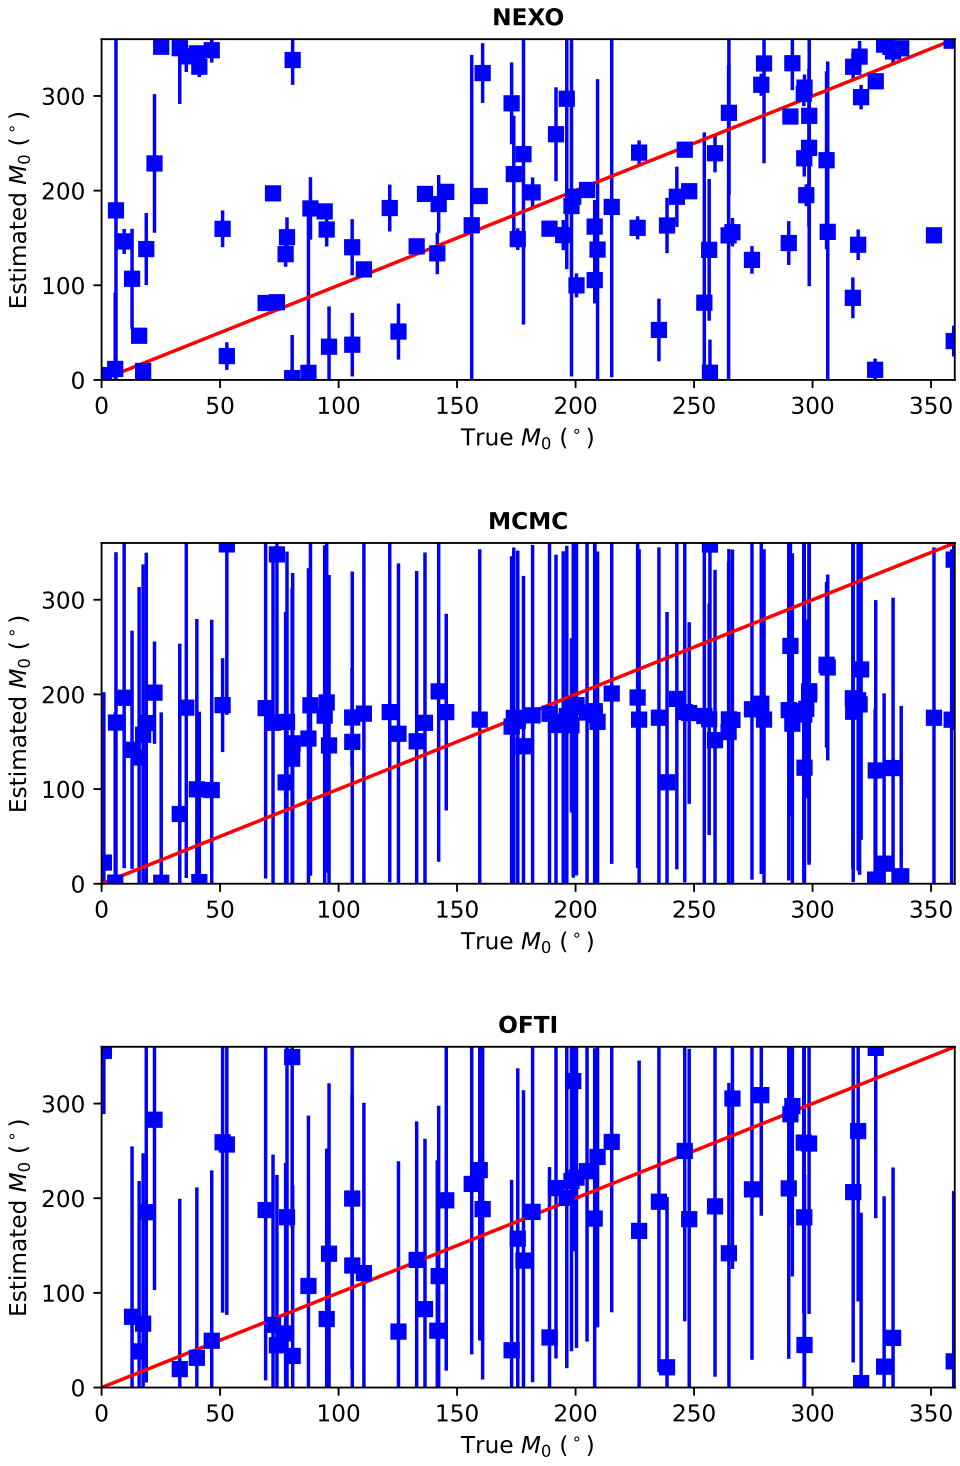

Figure 4.14: Credible intervals for the mean anomaly at epoch. None of the three methods gives reliable estimates, but OFTI gives the closest estimates.

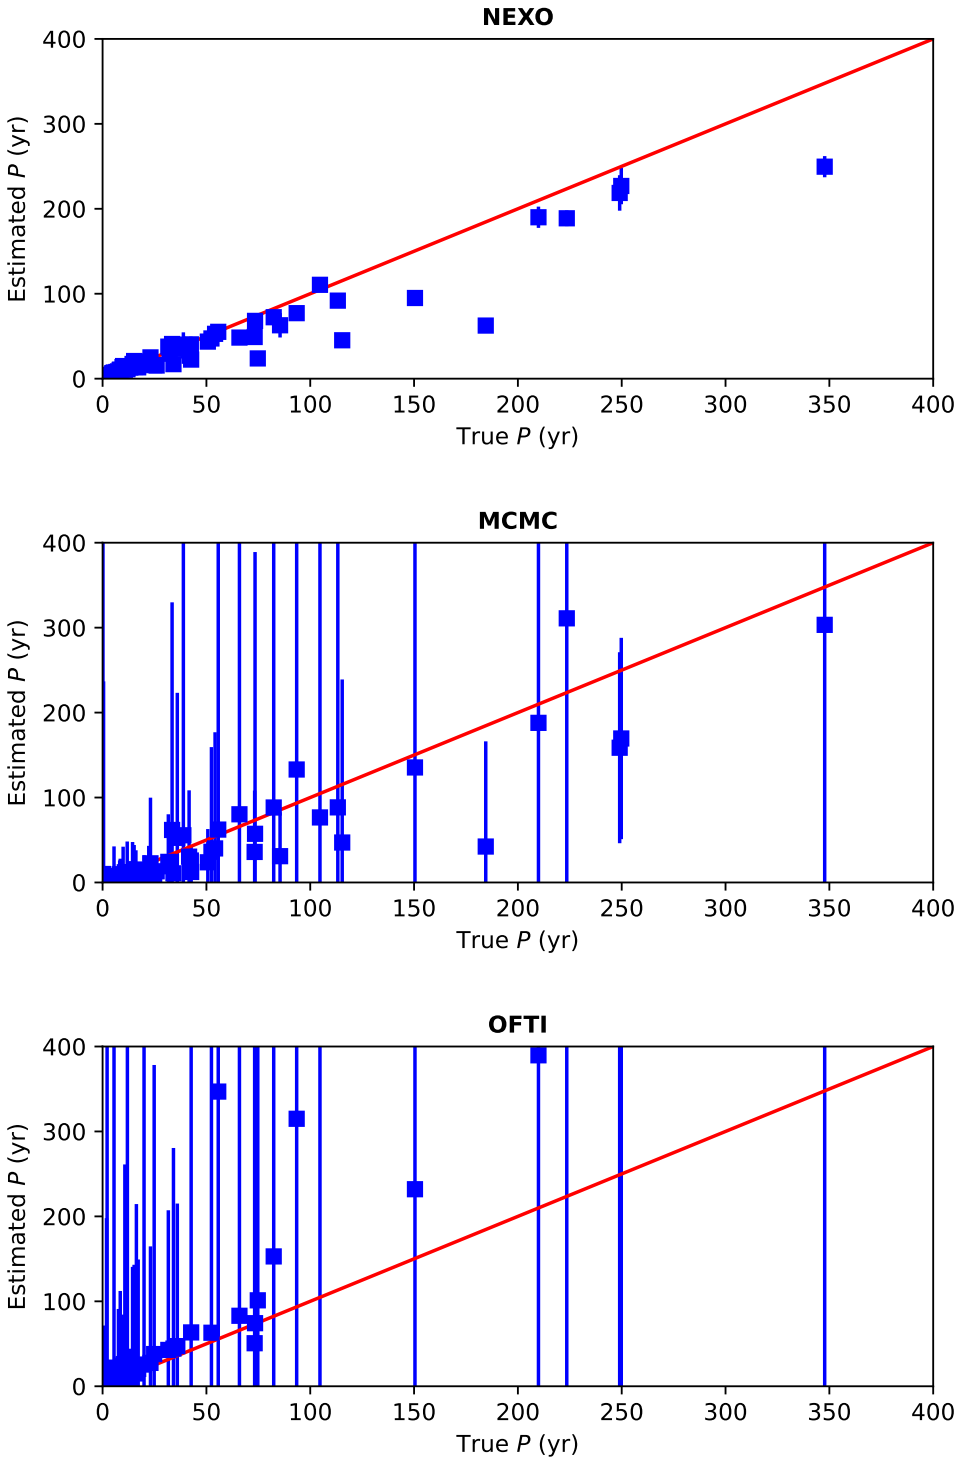

Figure 4.15: Credible intervals for the period. Unlike MCMC or OFTI, NEXO never significantly overestimates the period. As with the semi-major axis, however, the credible intervals for NEXO appear to be too narrow.

Table 4.1: Root-mean-square errors for the orbit estimators.

		NEXO	MCMC	OFTI
a	(%)	32.0	29.2	38.2
e	(%)	93.1	336.2	328.9
I	(°)	19.7	25.8	22.1
Ω	(°)	34.4	40.4	45.1
ω	(°)	82.1	108.0	105.0
M_0	(°)	78.6	89.0	69.8
P	(%)	59.7	571.9	143.9

Table 4.2: Percentages of orbital elements within credible intervals.

	NEXO	MCMC	OFTI
a	42.0	90.7	100.0
e	44.0	96.9	100.0
I	42.0	99.0	100.0
Ω	40.0	100.0	100.0
ω	29.0	93.8	100.0
M_0	26.0	94.8	100.0
P	33.0	90.7	100.0

4.5 Examples of Application

In this section, we present orbit fits using NEXO for two exoplanets and compare the results to previous orbit fits. The first planet is β Pictoris b, with measurements by Lagrange et al. [43] and Nielsen et al. [52]. The second is GJ 504 b, with measurements by Kuzuhara et al. [42]. The orbits of these two planets were previously fitted, respectively, using MCMC by Nielsen et al. [52] and using OFTI by Blunt et al. [8].

Table 4.3: Average run times for the orbit estimators.

Method	Time (s)
NEXO	1.5
MCMC	353.0
OFTI	66.1

Figures 4.16 and 4.17 show the true measurements and the predicted measurements based on the NEXO orbit fit for the two exoplanets. In both cases, the predicted measurements follow the true measurements very closely. Figure 4.18 compares the orbital element estimates for β Pictoris b using NEXO and MCMC, and Figure 4.19 does the same for GJ 504 b using NEXO and OFTI. Overall, the central estimates using NEXO are quite close to the previous fits, particularly in the first case. In both cases, the credible intervals obtained from NEXO are significantly narrower than the confidence intervals from the MCMC or OFTI fits; however, as discussed in Section 4.4.2, it is possible that the NEXO intervals are too narrow.

4.6 Conclusion

We have presented a new method for estimating exoplanet orbits, based on nonlinear filtering and a set of nonsingular orbital elements developed specifically for this purpose. Using both simulated and real astrometric measurements, we have shown that NEXO obtains estimates with accuracy comparable to or higher than that of the MCMC and OFTI methods, but with much faster execution. The next step in the development of NEXO would be to ensure that its output covariances and credible intervals reflect the true uncertainty of the state estimate. In the future, we plan to work on further improving the accuracy of NEXO by fine-tuning the existing filter or using other filtering techniques.

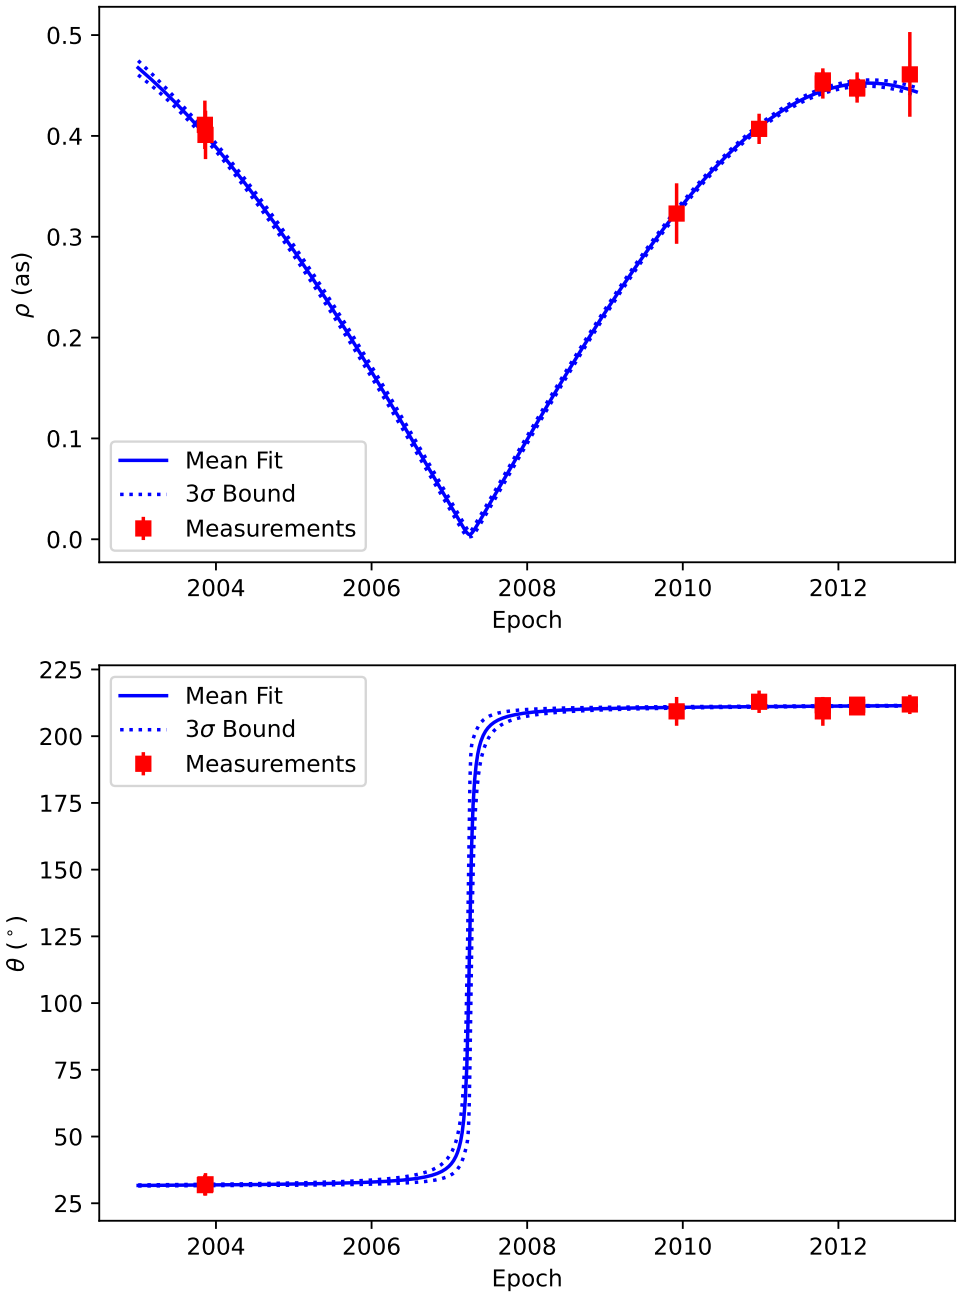

Figure 4.16: Astrometric measurements and orbit fits for β Pictoris b. The predicted measurements follow the true measurements very closely, despite the apparent "jump" in both ρ and θ.

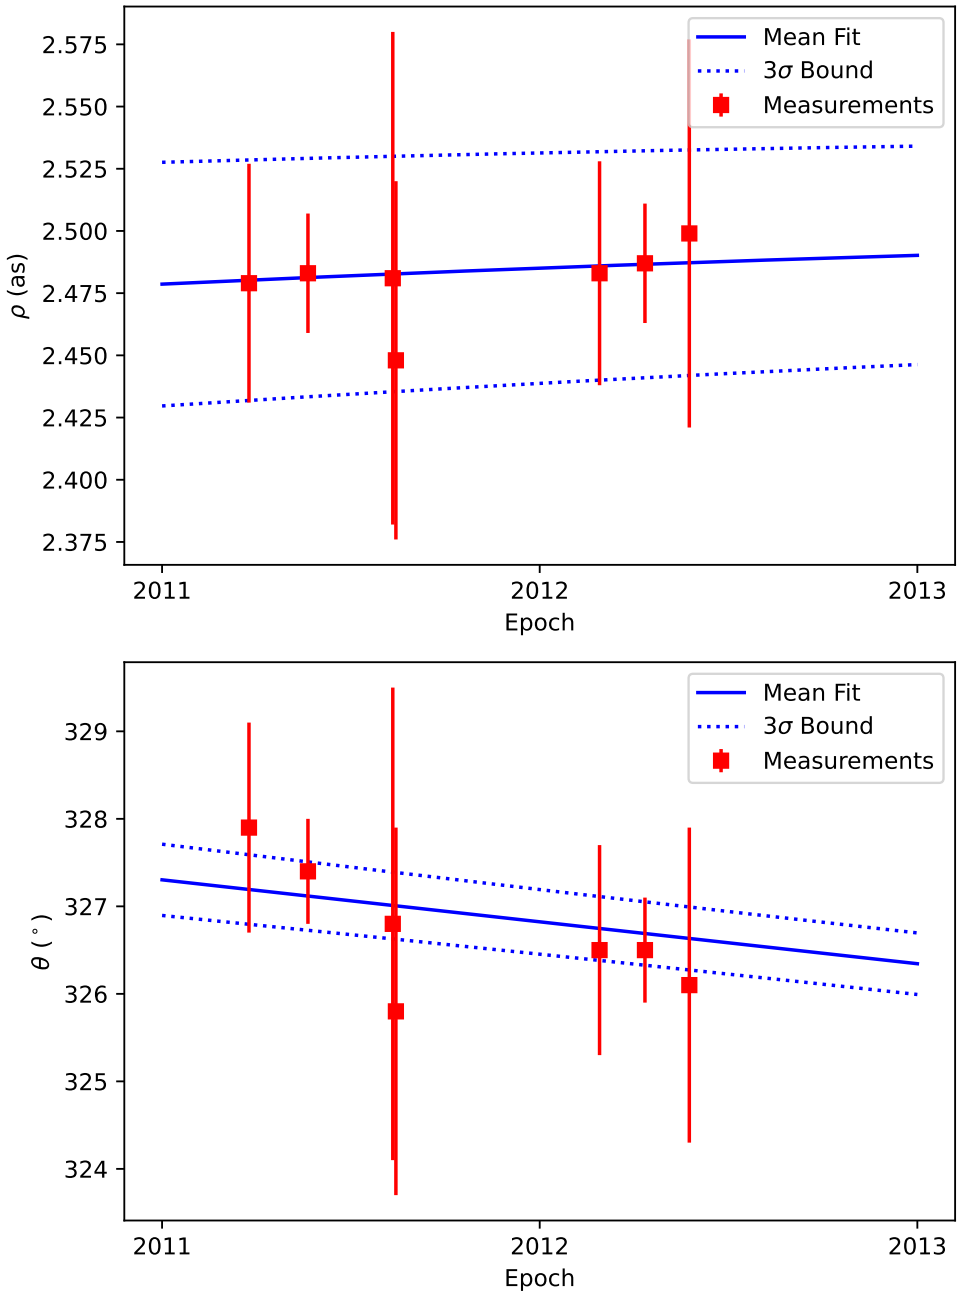

Figure 4.17: Astrometric measurements and orbit fits for GJ 504 b. We see good agreement between the orbit fit and the true measurements. However, the measurements appear to cover a smaller fraction of the orbit than for β Pictoris b, making this a more challenging fit.

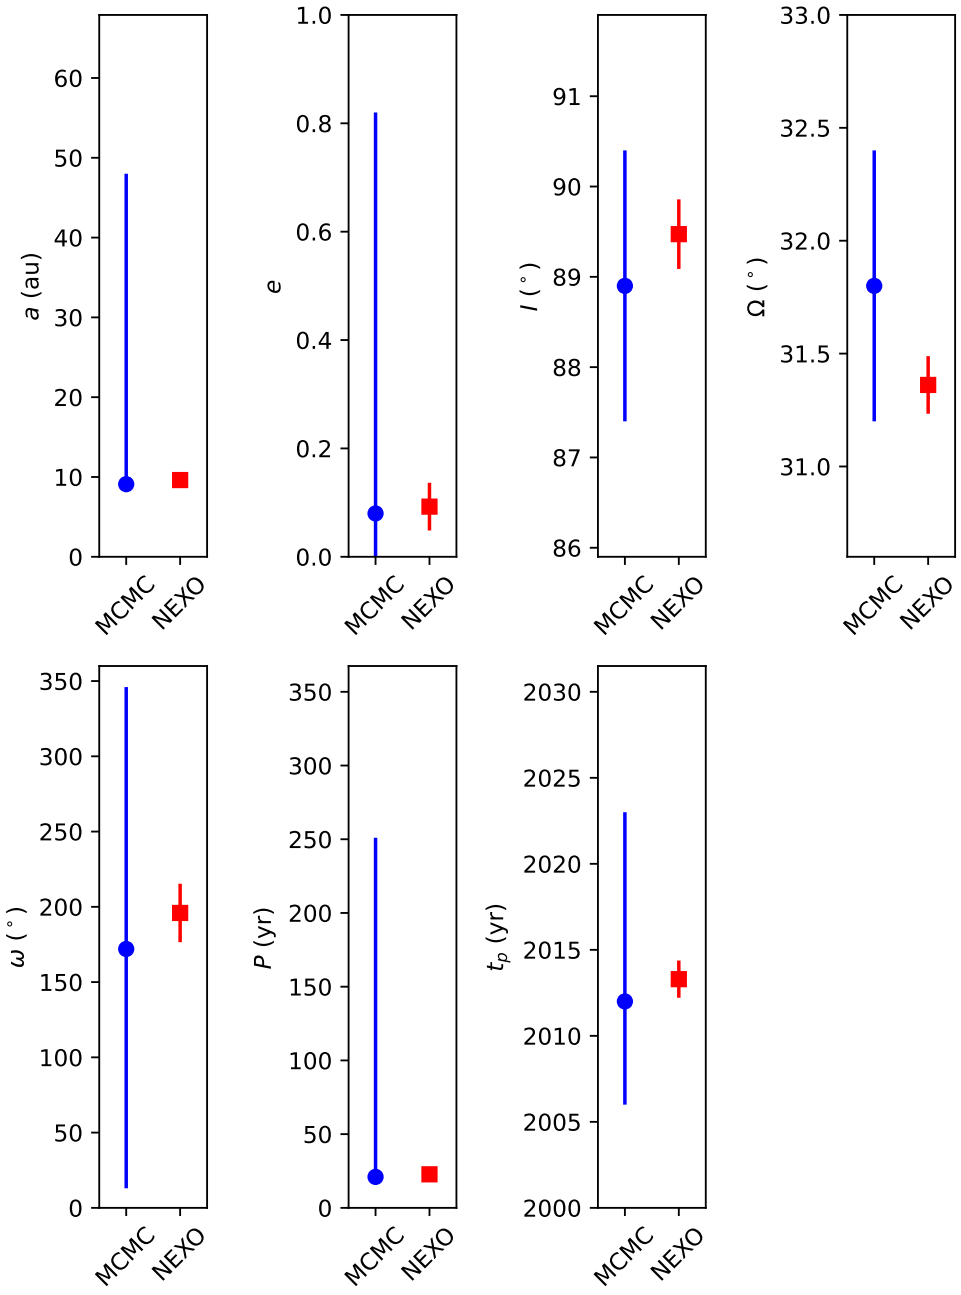

Figure 4.18: Orbital parameter estimates for β Pictoris b. For NEXO, we show the means and 95% credible intervals, and for MCMC we show the medians and 95% confidence intervals reported by Nielsen et al. using the fixed-mass method [52]. We shift Ω and ω in the MCMC fit by 180° for consistency. The central estimates using the two methods are remarkably close for all seven elements. Furthermore, the credible intervals for NEXO are contained within the confidence intervals for MCMC for all elements.

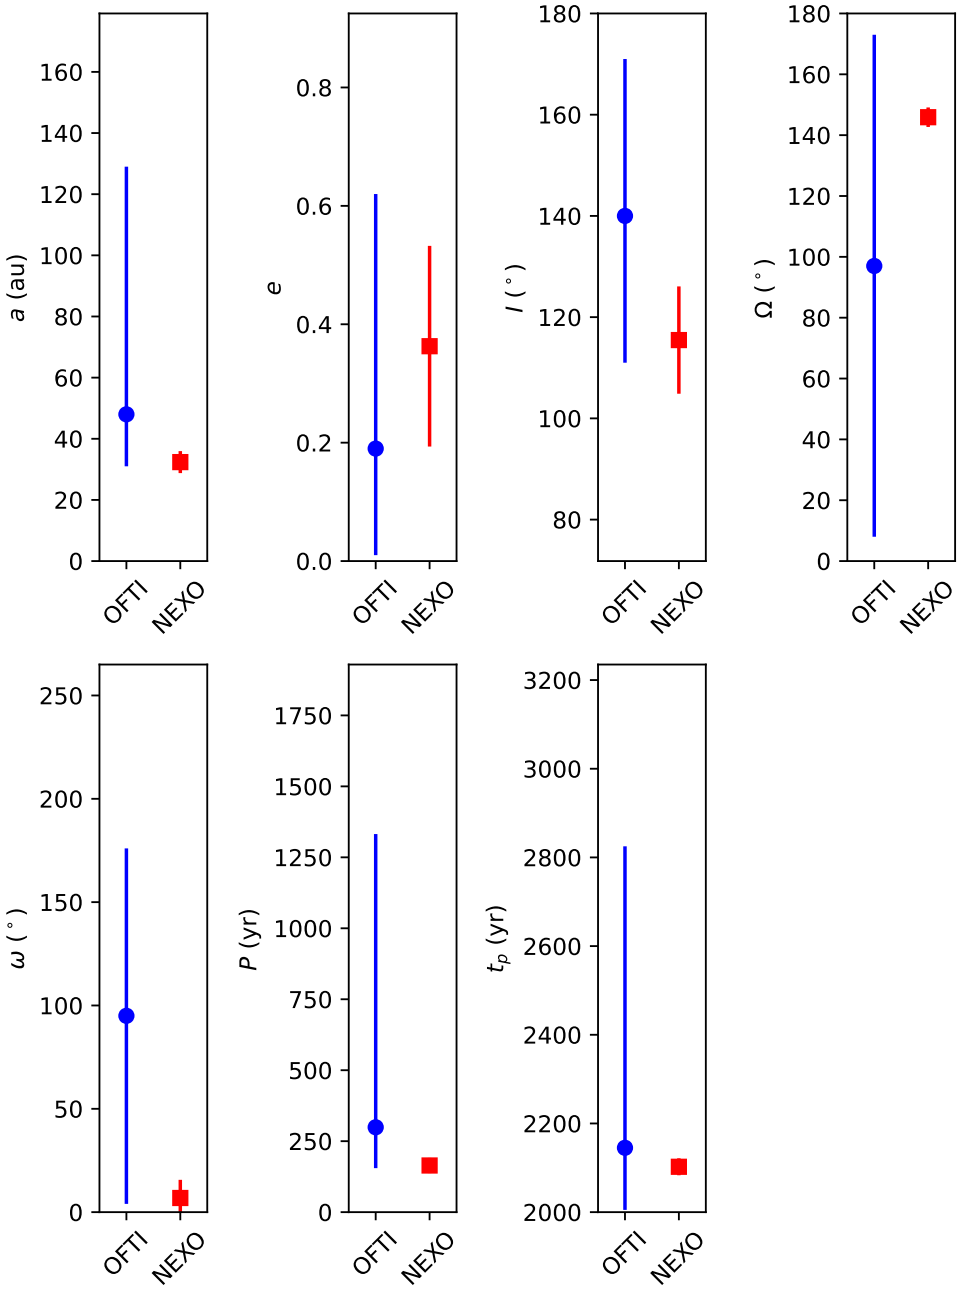

Figure 4.19: Orbital parameter estimates for GJ 504 b. For NEXO, we show the means and 95% credible intervals, and for OFTI we show the medians and 95% confidence intervals reported by Blunt et al. [7]. NEXO and OFTI give the closest central estimates for a, P, and t_p; the larger differences in the other elements may be due to the long period of the orbit.

CHAPTER 5

CONCLUSION

This thesis presents advancements in nonlinear estimation, particularly for problems in astrodynamics. It includes several projects with the common goal of obtaining more accurate state estimates in a broader range of nonlinear dynamical systems. The contributions include the development of new nonlinear filtering methods, as well as new applications of nonlinear filtering. We hope that the results of our work might contribute to the wider applicability of nonlinear filters in aerospace and astronomy.

First, we developed the Higher-Order Unscented Estimator (HOUSE), which extends the unscented Kalman filter (UKF) to account for skewness and kurtosis in addition to the mean and covariance. In simulations of nonlinear dynamical systems, HOUSE performed better than other state-of-the-art estimators when the noise produced many outliers. Thus, we have shown that a relatively small modification to the UKF can significantly improve the accuracy of state estimates in certain cases. With this in mind, it is worth investigating what other extensions to the UKF and HOUSE could be developed. For example, it is possible—at least in theory—to develop a filter that uses sigma points to propagate any finite set of moments, which could include marginal and mixed moments of the third, fourth, and higher orders. We expect such a filter to be even more accurate and applicable to a wider range of systems. However, increasing the number of moments that the filter propagates inevitably increases its computational complexity, and the trade-off between accuracy and computational efficiency has to be considered, particularly for real-time applications. Furthermore, it is necessary to investigate which moments are significant for particular systems

and noise distributions.

The second project is called Autonomous Cross-Calibration for Imaging Satellites (ACCIS). It applies nonlinear filtering methods with a combination of conventional and image-based measurements to cross-calibrate a constellation of Earth-imaging satellites. In this framework, satellites transmit features extracted from images along with their state estimates to other members of the constellations, and they use these data to achieve cross-calibration in real time. Using a detailed simulation, we showed that this method is feasible. However, further development of the ACCIS methodology is required to achieve a useful cross-calibration accuracy. In particular, since there are currently no standard models relating the state of a satellite to the image features, we have thus far only used a basic image-based measurement model in the ACCIS filters. We expect that, with appropriate refinements of this model, the accuracy of ACCIS can be significantly increased.

In the third project, we applied nonlinear filtering techniques to the problem of fitting exoplanet orbits to measurements from direct observations, developing the Nonsingular Estimator for Exoplanet Orbits (NEXO). Some of the main difficulties in this estimation problem are that the prior distribution of the orbital parameters is very diffuse and that the measurements are highly nonlinear. Consequently, the representation of the state distribution using a single mean and covariance has limitations, which we overcame by using Gaussian mixtures. Furthermore, with the development of a new set of nonsingular orbital elements, we demonstrated how the parametrization of the system state can be tailored to the estimation problem at hand. In particular, we were able to obtain orbital parameters with no singularities or bounds and with a partly linear mea-

surement model, both of which are useful features for filtering. Our testing of NEXO showed that filter-based methods can provide faster and more accurate exoplanet orbit fits than Monte Carlo methods. Still, we expect that the filtering techniques in NEXO can be further improved to give more accurate orbit fits and more reliable error estimates.

Throughout this work, we have shown that there is plenty of room for developing new filtering techniques and applying them to new dynamical systems. We plan to continue to explore these questions in our future work.

BIBLIOGRAPHY

[1] N. Adurthi, P. Singla, and T. Singh. *Conjugate Unscented Transform points*. GitHub repository. 2017. URL: `https://github.com/nadurthi/CUTpoints`.

[2] N. Adurthi, P. Singla, and T. Singh. "Conjugate Unscented Transform Rules for Uniform Probability Density Functions." *2013 American Control Conference* (2013), pp. 2454–2459.

[3] N. Adurthi, P. Singla, and T. Singh. "Conjugate Unscented Transformation: Applications to Estimation and Control." *Journal of Dynamic Systems, Measurement, and Control* 140.3 (2018), pp. 1–22.

[4] I. Arasaratnam and S. Haykin. "Cubature Kalman Filters." *Institute of Electrical and Electronics Engineers Transactions on Automatic Control* 54.6 (2009), pp. 1254–1269.

[5] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, 2001.

[6] L. Binnendijk. *Properties of Double Stars: A Survey of Parallaxes and Orbits*. University of Pennsylvania Press, 1960.

[7] S. Blunt, E. L. Nielsen, R. J. D. Rosa, et al. "Orbits for the Impatient: A Bayesian Rejection-sampling Method for Quickly Fitting the Orbits of Long-period Exoplanets." *The Astronomical Journal* 153.5 (2017).

[8] S. Blunt et al. "*orbitize!*: A Comprehensive Orbit-fitting Software Package for the High-contrast Imaging Community." *The Astronomical Journal* 159.3 (2020).

[9] B. Bowman et al. "A New Empirical Thermospheric Density Model JB2008 Using New Solar and Geomagnetic Indices." *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*. 2008.

[10] G. Bradski. "The OpenCV Library." *Dr. Dobb's Journal of Software Tools* (2000).

[11] S. Brunke and M. E. Campbell. "Square Root Sigma Point Filtering for Real-Time, Nonlinear Estimation." *Journal of Guidance, Control, and Dynamics* 27.2 (2004), pp. 314–317.

[12] J. Burkardt. *ODE: Shampine and Gordon ODE Solver*. Software library. 2020. URL: `https://people.sc.fsu.edu/~jburkardt/cpp_src/ode/ode.html`.

[13] Chauvin, G., Lagrange, A.-M., Beust, H., et al. "Orbital characterization of the β Pictoris b giant planet." *Astronomy & Astrophysics* 542 (2012).

[14] C. J. Cohen and E. C. Hubbard. "A Nonsingular Set of Orbital Elements." *The Astronomical Journal* 67.1 (1962).

[15] J. L. Crassidis and F. L. Markley. "Unscented Filtering for Spacecraft Attitude Estimation." *Journal of Guidance, Control, and Dynamics* 26.4 (2003), pp. 536–542.

[16] R. J. De Rosa, E. L. Nielsen, J. J. Wang, et al. "An Updated Visual Orbit of the Directly Imaged Exoplanet 51 Eridani b and Prospects for a Dynamical Mass Measurement with *Gaia*." *The Astronomical Journal* 159.1 (2020).

[17] C. Delacroix et al. "Science yield modeling with the Exoplanet Open-Source Imaging Mission Simulator (EXOSIMS)." *Modeling, Systems Engineering, and Project Management for Astronomy VI*. Ed. by G. Z. Angeli and P. Dierickx. Vol. 9911. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. Aug. 2016, p. 991119.

[18] W. P. Elderton and N. L. Johnson. *Systems of Frequency Curves*. Cambridge University Press, 1969.

[19] H. Fang et al. "Nonlinear Bayesian estimation: from Kalman filtering to a broader horizon." *Journal of Automatica Sinica* 5.2 (2018), pp. 401–417.

[20] N. I. Fisher. *Statistical Analysis of Circular Data*. Cambridge University Press, 1993.

[21] E. B. Ford. "Improving the Efficiency of Markov Chain Monte Carlo for Analyzing the Orbits of Extrasolar Planets." *The Astrophysical Journal* 642.1 (2006), pp. 505–522.

[22] E. B. Ford. "Quantifying the Uncertainty in the Orbits of Extrasolar Planets with Markov Chain Monte Carlo." *The Search for Other Worlds*. Ed. by S. S. Holt and D. Deming. Vol. 713. American Institute of Physics Conference Series. 2004, pp. 27–30.

[23] R. G. v. Gioi et al. "A Precision Analysis of Camera Distortion Models." *IEEE Transactions on Image Processing* 26.6 (2017), p. 2694.

[24] R. G. Gottlieb. *Fast Gravity, Gravity Partials, Normalized Gravity, Gravity Gradient Torque and Magnetic Field: Derivation, Code and Data*. NASA Contractor Report 188243, 1993.

[25] M. S. Grewal and A. P. Andrews. "Applications of Kalman Filtering in Aerospace 1960 to the Present [Historical Perspectives]." *IEEE Control Systems Magazine* 30.3 (2010), pp. 69–78.

[26] O. Grothe. "A higher order correlation unscented Kalman filter." *Applied Mathematics and Computation* 219.17 (2013), pp. 9033–9042.

[27] G. Guennebaud and B. Jacob. *Eigen 3*. Software library. 2010. URL: `http://eigen.tuxfamily.org`.

[28] M. Harwit. *Astrophysical Concepts*. 4th. Springer, 2006.

[29] W. D. Heintz. *Double Stars*. 1st. D. Reidel Publishing Company, 1978.

[30] K. Ito and K. Xiong. "Gaussian filters for nonlinear filtering problems." *Institute of Electrical and Electronics Engineers Transactions on Automatic Control* 45.5 (2000), pp. 910–927.

[31] E. Jahnke and F. Emde. *Tables of Functions with Formulae and Curves [Funktionentafeln mit Formeln und Kurven]*. 4th. Dover, 1945.

[32] B. Jia, M. Xin, and Y. Cheng. "High-degree cubature Kalman filter." *Automatica* 49.2 (2013), pp. 510–518.

[33] S. Julier and J. Uhlmann. "Reduced sigma point filters for the propagation of means and covariances through nonlinear transformations." *Proceedings of the 2002 American Control Conference* 2 (2002), pp. 887–892.

[34] S. J. Julier and J. K. Uhlmann. "New extension of the Kalman filter to nonlinear systems." *Signal Processing, Sensor Fusion, and Target Recognition VI*. Ed. by I. Kadar. Vol. 3068. Society of Photo-Optical Instrumentation Engineers, 1997, pp. 182–193.

[35] R. E. Kalman. "A New Approach to Linear Filtering and Prediction Problems." *Journal of Basic Engineering* 82 (1960), pp. 35–45.

[36] G. Kaplan et al. *User's Guide to NOVAS Version F3.1*. United States Naval Observatory, 2011.

[37] B. D. Keister. "Multidimensional quadrature algorithms." *Computers in Physics* 10.2 (1996), pp. 119–122.

[38] D. Keithly. "Evaluating and Scheduling Exoplanet Direct Imaging Missions." PhD thesis. Cornell University, Dec. 2021.

[39] Q. M. Konopacky, C. Marois, B. A. Macintosh, et al. "Astrometric Monitoring of the HR 8799 Planets: Orbit Constraints from Self-consistent Measurements." *The Astronomical Journal* 152.2 (2016).

[40] Q. M. Konopacky, S. J. Thomas, B. A. Macintosh, et al. "Gemini planet imager observational calibrations V: Astrometry and distortion." *Ground-based and Airborne Instrumentation for Astronomy V*. Ed. by S. K. Ramsay, I. S. McLean, and H. Takami. Vol. 9147. International Society for Optics and Photonics. SPIE, 2014, pp. 2585–2600.

[41] C. Krafft, A. Volokitin, and G. Gauthier. "Turbulence and microprocesses in inhomogeneous solar wind plasmas." *Fluids* 4.2 (2019), pp. 43–45.

[42] M. Kuzuhara, M. Tamura, T. Kudo, et al. "Direct Imaging of a Cold Jovian Exoplanet in Orbit Around the Sun-Like Star GJ 504." *The Astrophysical Journal* 774.1 (2013).

[43] A.-M. Lagrange, D. Gratadour, G. Chauvin, et al. "A probable giant planet imaged in the β Pictoris disk: VLT/NaCo deep L'-band imaging." *Astronomy & Astrophysics* 493.2 (2009), pp. L21–L25.

[44] L. Leung, V. Beukelaers, S. Chesi, et al. "ADCS at Scale: Calibrating and Monitoring the Dove Constellation." *32nd Annual AIAA/USU Conference on Small Satellites* (2018).

[45] J.-F. Lévesque. "Second-Order Simplex Sigma Points for Nonlinear Estimation." *AIAA Guidance, Navigation, and Control Conference and Exhibit* (2006), pp. 1–12.

[46] D. G. Lowe. "Distinctive Image Features from Scale-Invariant Keypoints." *International Journal of Computer Vision* 60.2 (2004), pp. 91–110.

[47] B. Macintosh, J. R. Graham, P. Ingraham, et al. "First Light of the Gemini Planet Imager." *Proceedings of the National Academy of Sciences* 111.35 (2014), pp. 12661–12666.

[48] F. Mannan and M. S. Langer. "What is a Good Model for Depth from Defocus?" *13th Conference on Computer and Robot Vision (CRV)*. 2016, pp. 273–280.

[49] F. L. Markley and J. L. Crassidis. *Fundamentals of Spacecraft Attitude Determination and Control*. Springer, 2014.

[50] E. J. McGrath and D. C. Iriving. *Random Number Generation for Selected Probability Distributions*. Office of Naval Research, 1975.

[51] M. A. Millar-Blanchaer, J. R. Graham, L. Pueyo, et al. "β Pictoris' Inner Disk in Polarized Light and New Orbital Parameters for β Pictoris b." *The Astrophysical Journal* 811.1 (2015).

[52] E. L. Nielsen, M. C. Liu, Z. Wahhaj, et al. "The Gemini/NICI Planet-Finding Campaign: The Orbit of the Young Exoplanet β Pictoris b." *The Astrophysical Journal* 794.2 (2014).

[53] N. K. Pavlis et al. "An Earth Gravitational Model to Degree 2160: EGM2008." *Presented at the 2008 General Assembly of the European Geosciences Union* (2008).

[54] F. Pirotti et al. "An open source virtual globe rendering engine for 3D applications: NASA WorldWind." *Open Geospatial Data, Software and Standards* 2.4 (2017).

[55] H. Plummer. *An Introductory Treatise on Dynamical Astronomy*. Cambridge University Press, 1918.

[56] K. Ponomareva and P. Date. "Higher order sigma point filter: A new heuristic for nonlinear time series filtering." *Applied Mathematics and Computation* 221 (2013), pp. 662–671.

[57] K. Ponomareva, P. Date, and Z. Wang. "A new unscented Kalman filter with higher order moment-matching." *Proceedings of the 19th International Symposium on Mathematical Theory of Networks and Systems* (2010).

[58] W. H. Press, G. B. Rybicki, and J. N. Hewitt. "The Time Delay of Gravitational Lens 0957+561. II. Analysis of Radio Data and Combined Optical-Radio Analysis." *The Astrophysical Journal* 385 (1992).

[59] A. M. Price-Whelan et al. "The Joker: A Custom Monte Carlo Sampler for Binary-star and Exoplanet Radial Velocity Data." *The Astrophysical Journal* 837.1 (2017), p. 20.

[60] J. Rameau, E. L. Nielsen, R. J. D. Rosa, et al. "Constraints on the Architecture if the HD 95086 Planetary System with the Gemini Planet Imager." *The Astrophysical Journal* 822.2 (2016).

[61] J. V. Ramsdell. "Wind Shear Fluctuations Downwind of Large Surface Roughness Elements." *Journal of Applied Meteorology* 17.4 (1978), pp. 436–443.

[62] D. Savransky. "Sub-Pixel Inter-Satellite Imagery Cross-Calibration via Image Decomposition and Dynamic Filtering" (2019).

[63] D. Savransky and D. Garrett. "WFIRST-AFTA coronagraph science yield modeling with EXOSIMS." *Journal of Astronomical Telescopes, Instruments, and Systems* 2 (Jan. 2016), p. 011006.

[64] D. Savransky and N. J. Kasdin. "Dynamic Filtering for the Analysis of Astrometric and Radial Velocity Data Sets for the Detection of Exoplanets." *AIAA Guidance, Navigation, and Control Conference*. 2009.

[65] L. Shampine and M. Gordon. *Computer Solution of Ordinary Differential Equations: The Initial Value Problem*. Freeman, 1975.

[66] J. Shapiro. "Using Modern Mathematical and Computational Tools for Image Processing." PhD thesis. Cornell University, Dec. 2020.

[67] M. Šimandl and J. Duník. "Sigma Point Gaussian Sum Filter Design Using Square Root Unscented Filters." *IFAC Proceedings Volumes* 38.1 (2005). 16th IFAC World Congress, pp. 1000–1005.

[68] A. Steinwolf and S. A. Rizzi. "Non-Gaussian Analysis of Turbulent Boundary Layer Fluctuating Pressure on Aircraft Skin Panels." *AIAA Journal of Aircraft* 43.6 (2012), pp. 1662–1675.

[69] Z. Stojanovski and D. Savransky. "Autonomous Cross-Calibration for Imaging Satellites." *Astrodynamics Specialists Conference* (2021).

[70] Z. Stojanovski and D. Savransky. *ACCIS: Autonomous Cross-Calibration for Imaging Satellites*. GitHub repository. 2023. URL: `https://github.com/SIOSlab/ACCIS`.

[71] Z. Stojanovski and D. Savransky. "Higher-Order Unscented Estimator." *Journal of Guidance, Control, and Dynamics* 44.12 (2021).

[72] Z. Stojanovski and D. Savransky. *The Higher-Order Unscented Estimator*. GitHub repository. 2021. URL: `https://github.com/SIOSlab/HOUSE`.

[73] Z. Stojanovski and D. Savransky. "Unscented filtering for directly observed exoplanet orbits." *Space Telescopes and Instrumentation 2022: Optical, Infrared, and Millimeter Wave*. Ed. by L. E. Coyle, S. Matsuura, and M. D. Perrin.

Vol. 12180. International Society for Optics and Photonics. SPIE, 2022, 121805G.

[74] A. H. Stroud. *Approximate Calculation of Multiple Integrals*. Prentice-Hall, 1971.

[75] D. Tenne and T. Singh. "The higher order unscented filter." *Proceedings of the 2003 American Control Conference* 3 (2003), pp. 2441–2446.

[76] D. A. Vallado. *Fundamentals of Astrodynamics and Applications*. Microcosm Press, 2014.

[77] J. J. Wang, J. R. Graham, L. Pueyo, et al. "The Orbit and Transit Prospects for β Pictoris b Constrained with One Milliarcsecond Astrometry." *The Astronomical Journal* 152.4 (2016).

[78] P. H. Westfall. "Kurtosis as Peakedness, 1905–2014. R.I.P." *The American Statistician* 68.3 (2014), pp. 191–195.