Senior Design Report: Unscented Filtering for Satellite State Estimation

Nathaniel Kinzly (nrk42)

Spring 2020

1 Introduction

State estimation of a satellite is an important part of any mission. An incorrect state estimate can lead to errors in satellite measurements or controls. In order to compare new filtering schemes with existing schemes. I implemented 3 different Unscented Kalman Filters (UKFs) and compared their performance first on toy problems and then on satellite dynamics.

2 Filters Implemented

2.1 Standard UKF

The first filter I implemented was the a modified form of the standard UKF presented by Julier and Uhlmann [1]. The filter generates a relatively small set of *sigma points* with corresponding weights such that the weighted average of the sigma points is the same as the current mean and the weighted covariance calculation gives the current covariance. These sigma points are then propagated through the dynamics and measurement function to give a set of updated state and measurement values. After calculating an effective mean updated state and mean updated measurement, the final mean and covariance are calculated using the normal Kalman filter equations.

One small difference between my initial implementation and Julier and Uhlmann's formulation is I used a non-augmented formulation of the filter [2]. I decided to use this formulation first since it was more intuitive for me to understand. The main difference between these two formulation is the augmented formulation extends the dimensionality of the sigma points to include the measurement and process noise and allows for any relationship between state/measurement and noise, and the non-augmented version uses a smaller dimensionality sigma point and adds the process and measurement noise to the state and measurement respectively. Additionally, due to some initial difficulties getting the filter to function, I also made use of the parameters used in [3] to help tune my filter.

2.2 Augmented UKF

After I had the non-augmented formulation working, I was able to better understand the augmented formulation. Since the next higher moment version of the filter (see next section) used the augmented formulation, I decided to implement a normal augmented formulation of the filter. Implementing this allowed me to be sure I understood the augmented formulation and compare my augmented and non-augmented filters to ensure they performed the same.

2.3 Higher Moment UKF

The final filter I implemented was the higher moment UKF presented by Ponomareva [4]. This filter uses a higher dimensionality sigma points like that used by the augmented version of the UKF, but it uses a slightly modified set of sigma points and weights. Both the sigma points and weights are dependent on two parameters, α and β that are propagated in time like the covariance. These parameters contain data

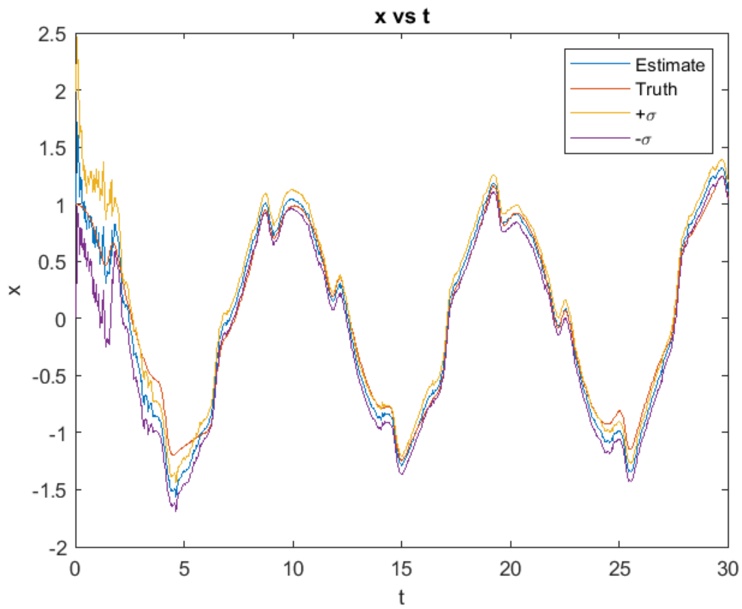

Figure 1: The cart position estimate as a function of time from the **non-augmented UKF**

about the third and fourth moments of the state distribution, which allows the filter to capture these higher moments. For the application of satellite filtering, this is especially useful since the dynamics are highly nonlinear.

3 Implementation Approach

3.1 Guiding Principles

Since the satellite dynamics functionality was being developed in parallel with the filter, and since the dynamics and measurement functions had the potential to change drastically, the code I wrote had to be modular. In particular, no part of the dynamics, measurement, or state dimensionality could be hard coded into the filtering. Using this as a guiding principal, I made no assumptions about the size of input arrays, and I made sure it would work as long as the dimensions of the inputs agree. Additionally, pointers to the dynamics and measurement function were also inputs to the filters so that they could be swapped easily. The resulting code is easily adaptable in the event something changes.

3.2 Toy Problem

In order to validate my work, I initially looked to test my results on a toy problem. In a previous semester, I implemented an extended Kalman Filter (EKF) to estimate the state of a pendulum on a cart. A schematic of the system is displayed below. INSERT DIAGRAM. The state of the system was x, the position of the cart, θ, the angle of the pendulum with the vertical axis with an angle of 0 corresponding to the pendulum point to the positive $\hat{j}$ axis, the velocity of the cart , $\dot{x}$, and the angular velocity of the pendulum, $th\dot{e}ta$. There were two measurements for this system: the carts acceleration, $\ddot{x}$, and $\sin\theta$.

Evaluating my filters on a toy problem and comparing my result with a prior result which I knew to be correct gave me more confidence in the correctness of my UKF. The results of executing both the non-augmented UKF, augmented UKF, and higher moment UKF on the toy problem are given in figures 1 through 20.

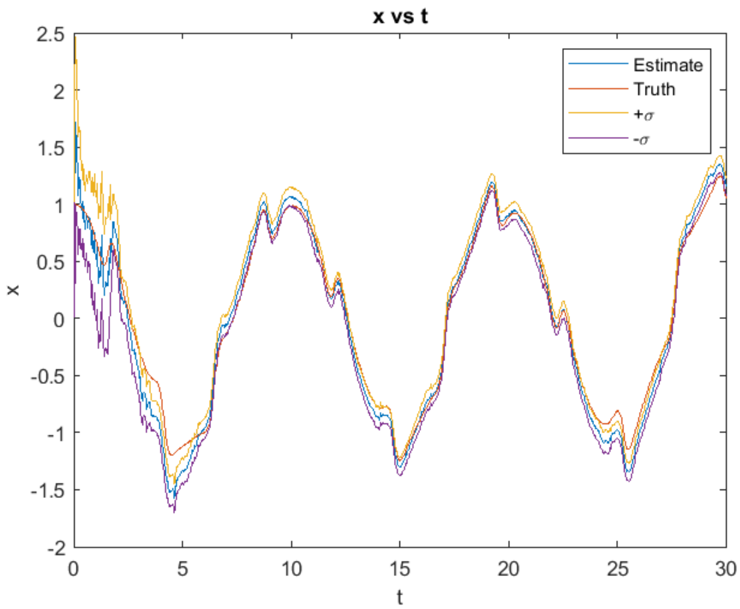

Figure 2: The cart position estimate as a function of time from the **augmented UKF**

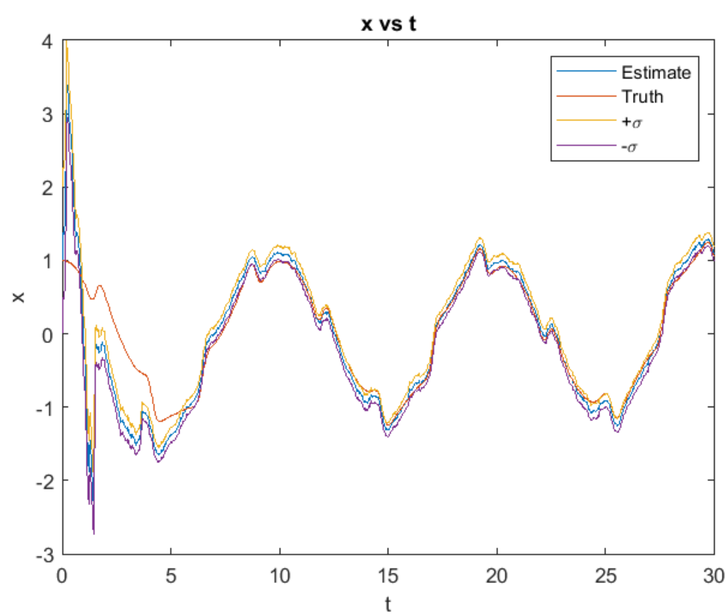

Figure 3: The cart position estimate as a function of time from the **higher moment UKF**

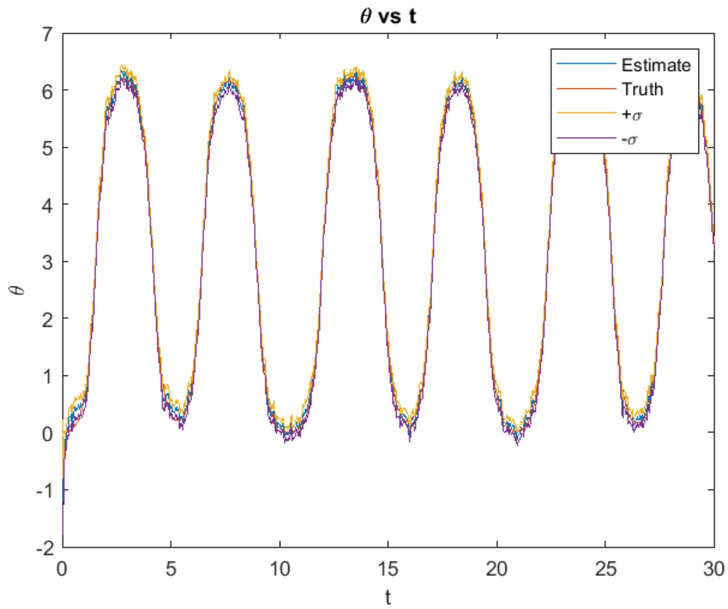

Figure 4: The pendulum angle estimate as a function of time from the **non-augmented UKF**

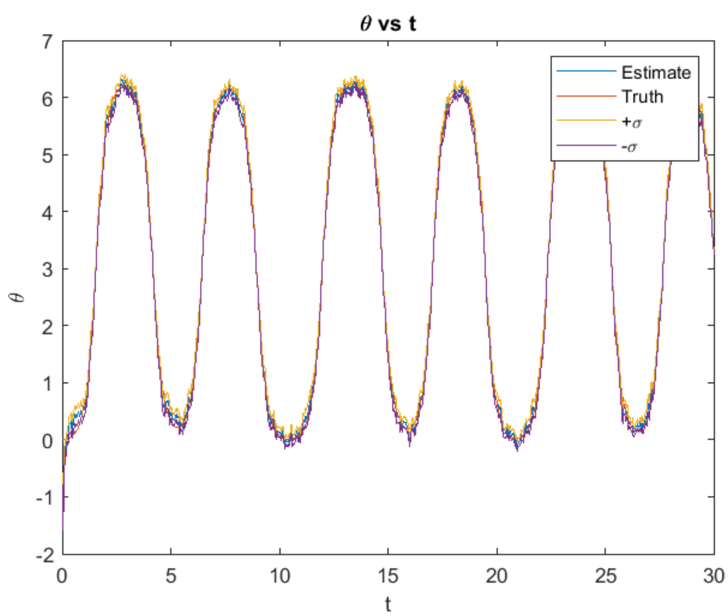

Figure 5: The pendulum angle estimate as a function of time from the **augmented UKF**

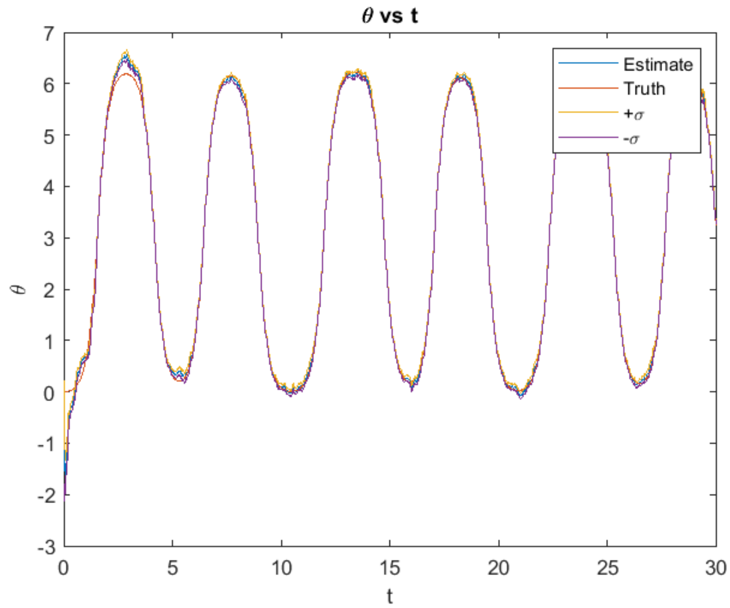

Figure 6: The pendulum angle estimate as a function of time from the **higher moment UKF**

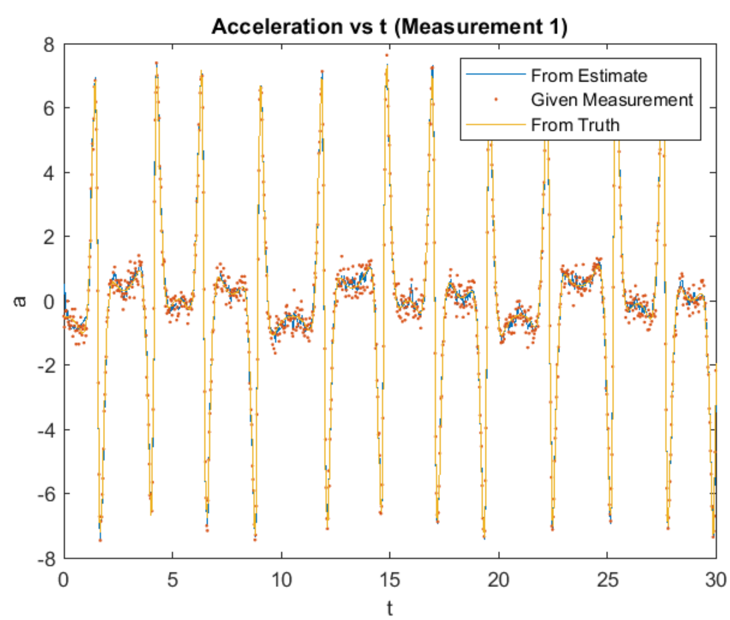

Figure 7: The cart acceleration based on the estimated state, true state, and measurement from the **non-augmented UKF**

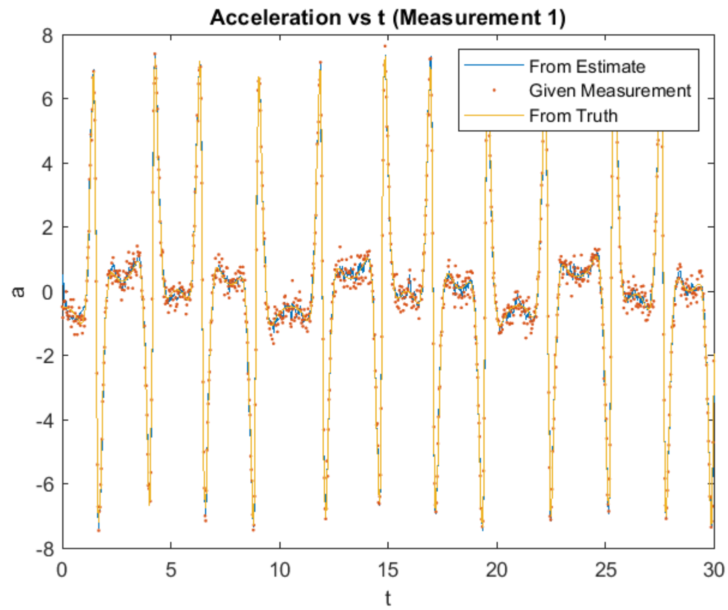

Figure 8: The cart acceleration based on the estimated state, true state, and measurement from the **augmented UKF**

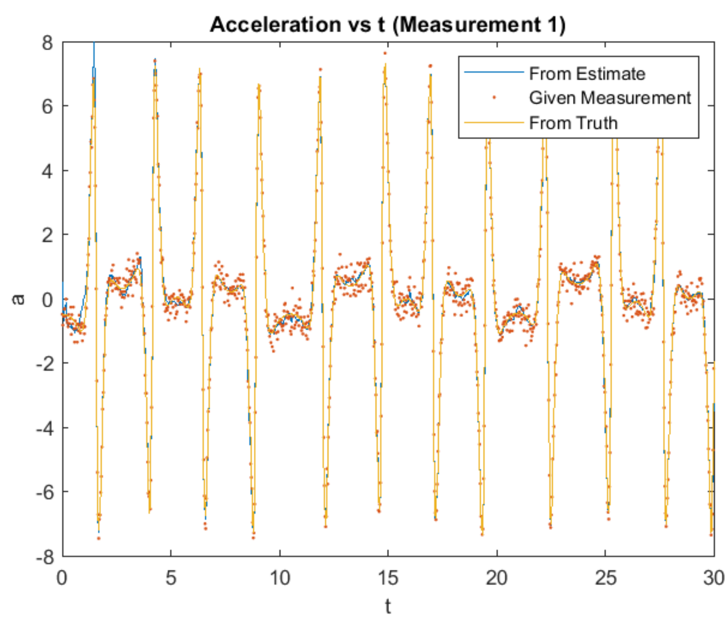

Figure 9: The cart acceleration based on the estimated state, true state, and measurement from the **higher moment UKF**

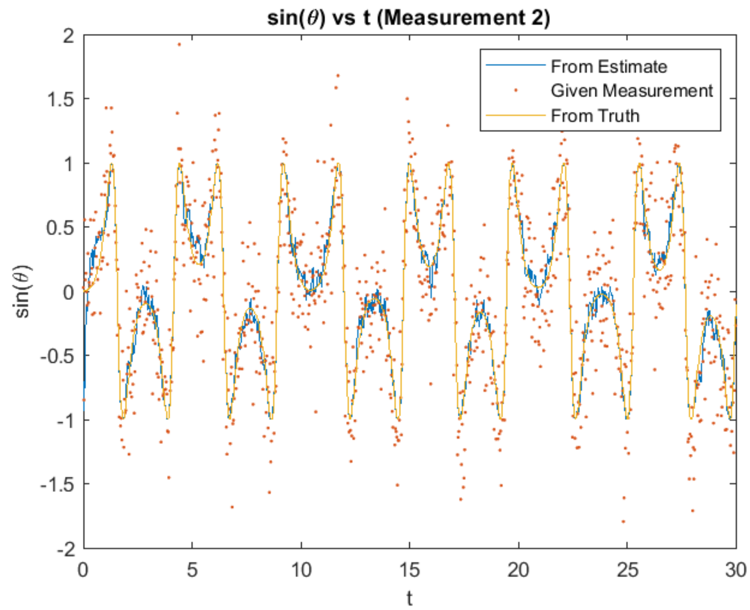

Figure 10: The sin of the pendulum angle based on the estimated state, true state, and measurement from the **non-augmented UKF**

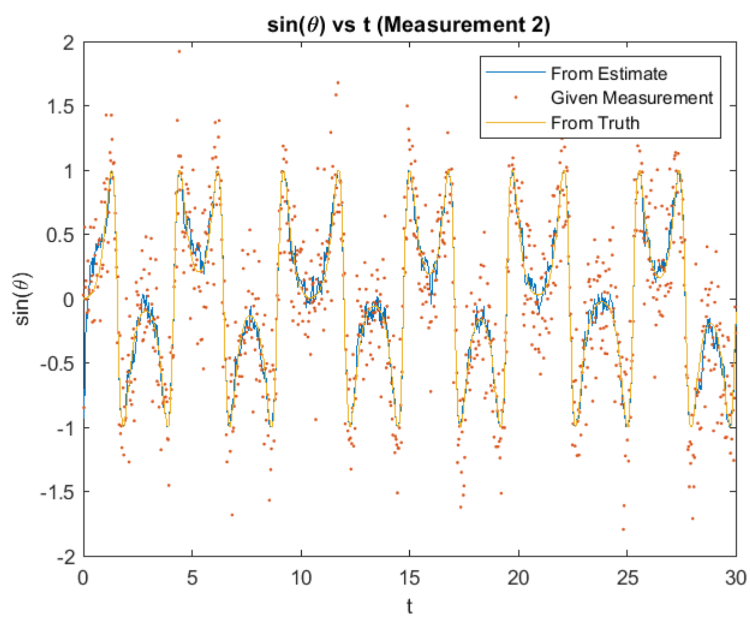

Figure 11: The sin of the pendulum angle based on the estimated state, true state, and measurement from the **augmented UKF**

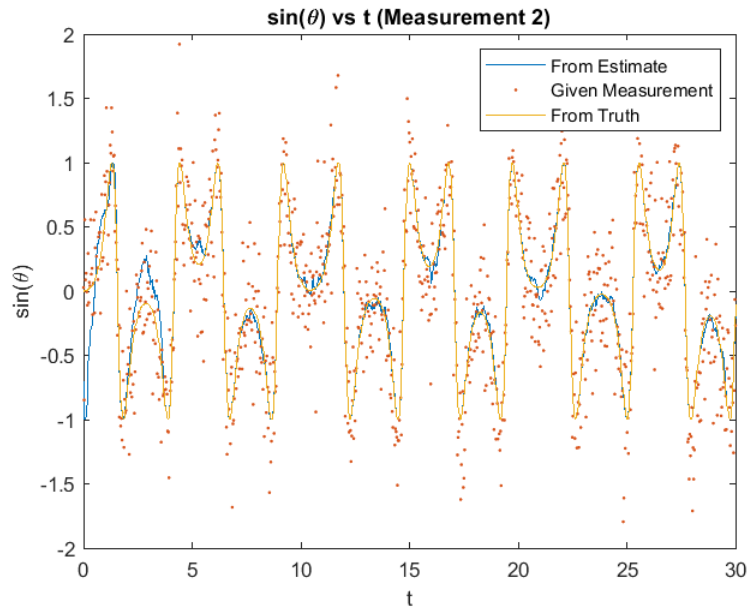

Figure 12: The sin of the pendulum angle based on the estimated state, true state, and measurement from the **higher moment UKF**

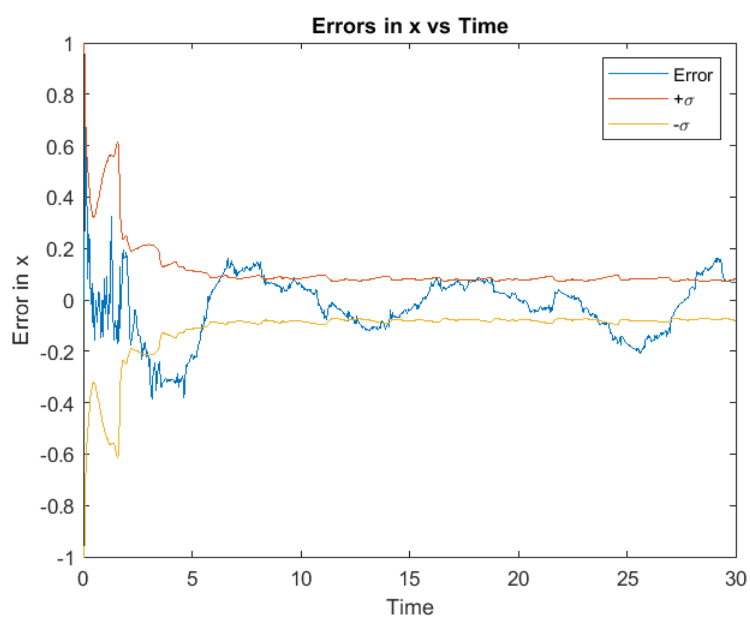

Figure 13: The errors in the cart position of the estimate of the **non-augmented UKF**

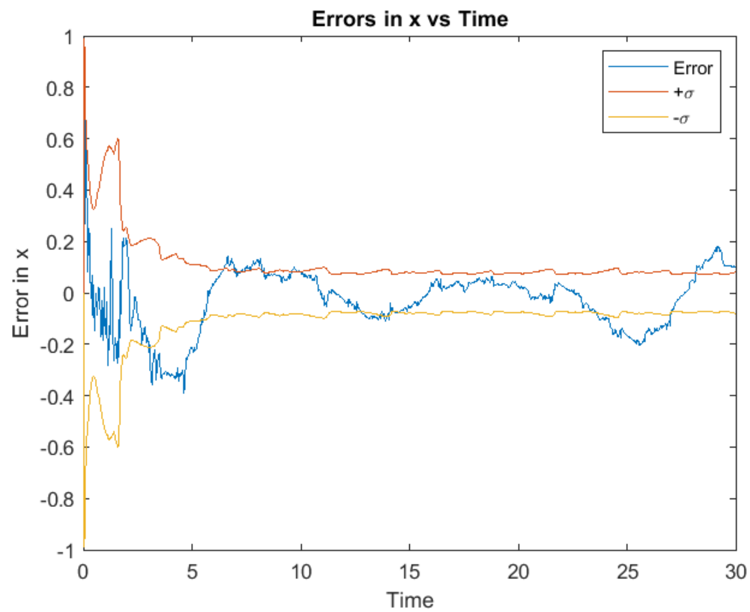

Figure 14: The errors in the cart position of the estimate of the **augmented UKF**

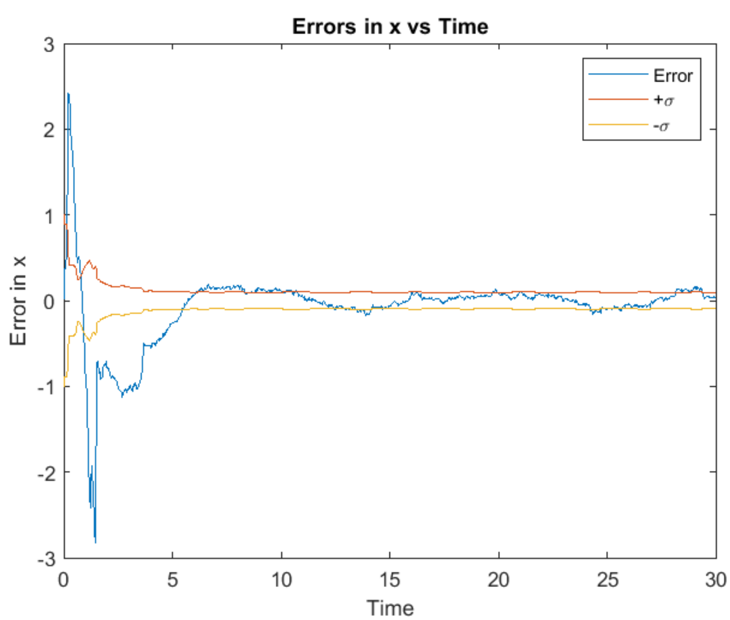

Figure 15: The errors in the cart position of the estimate of the **higher moment UKF**

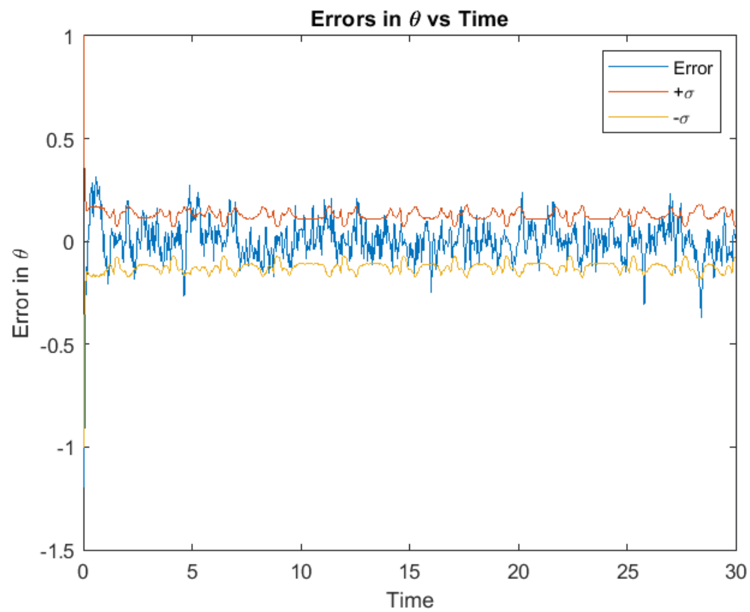

Figure 16: The errors in the pendulum angle of the estimate of the **non-augmented UKF**

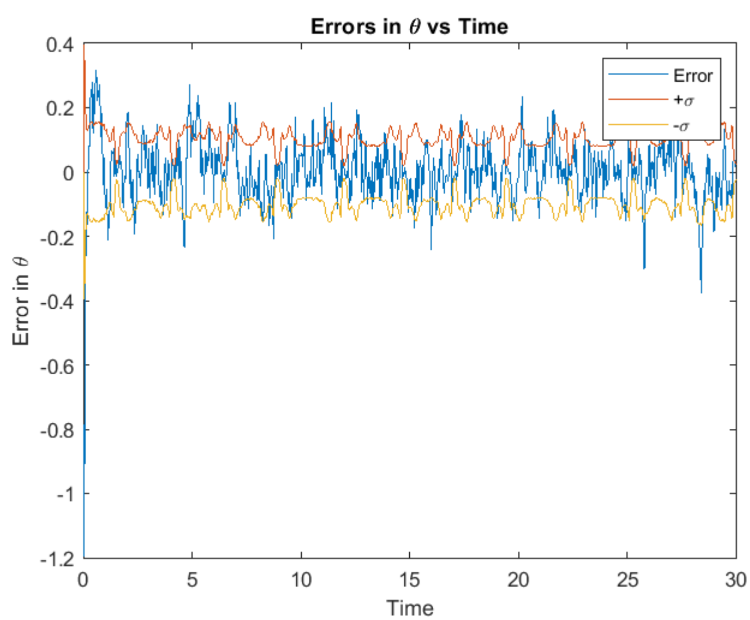

Figure 17: The errors in the pendulum angle of the estimate of the **augmented UKF**

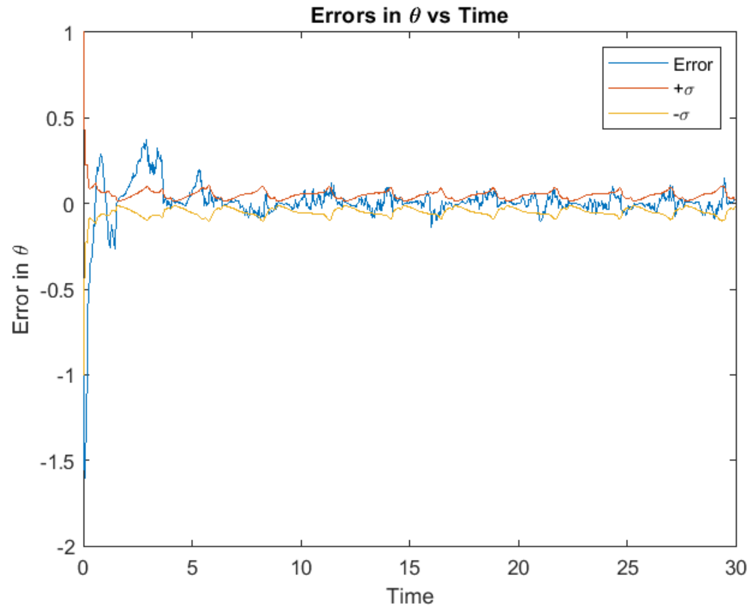

Figure 18: The errors in the pendulum angle of the estimate of the **higher moment UKF**

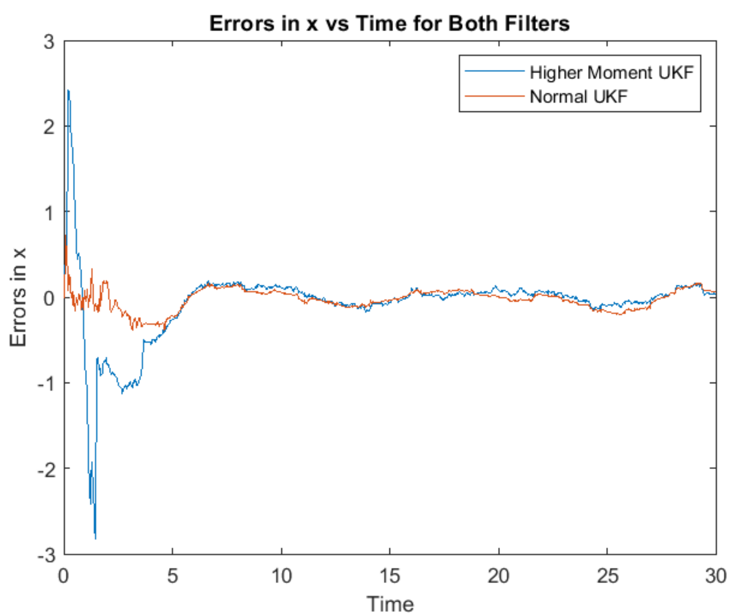

Figure 19: The errors in the cart position of both the **higher moment UKF** and **non-augmented UKF** overlayed.

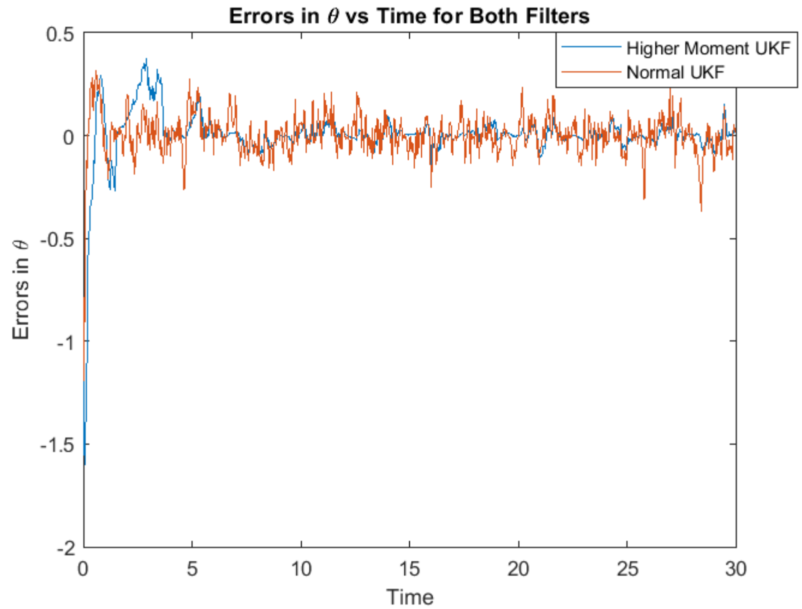

Figure 20: The errors in the pendulum angle of both the **higher moment UKF** and **non-augmented UKF** overlayed.

3.3 Comparison with Existing UKF

in addition to comparing my filters with an EKF that I had previously written, I also compared the normal UKF and higher moment UKF with a UKF provided as an example for a textbook exercise in Crassidis and Junkins [5]. My normal UKF lines up almost exactly with Crassidis and Junkins's UKF, which is to be expected since they should be doing identical operations. However, my higher moment UKF performed significantly worse than the provided UKF. After trying to tune the higher moment filter so that it would perform better, I concluded that there is something about the problem Crassidis and Junkins chose that made it harder to tune the higher moment UKF to perform better, so I moved on to test the filters on satellite dynamics.

3.4 Satellite Dynamics

Once I convinced myself that the filters I had written were working correctly, I applied the filters to the satellite system. The system has a 13 element state: 3 position components, 3 velocity components, 3 angular velocity components, and a 4 element quaternion representing the rotation of the spacecraft. The measurements of the system are the position, velocity, and rotation of the satellite. The filter does not have knowledge of the drag force on the system, so that force is acting as a disturbance. Overall, the satellite model considers a gravity torque and perturbations from the moon, the earth's non-spherical gravity field, and drag.

Even though the satellite system is much more complex than the pendulum on a cart system, integrating the filter with this system was mainly an exercise in tuning the filter correctly. The results of running the higher moment UKF and standard UKF for this problem are depicted in figures 22 through 33.

4 Filter Evaluation

As shown by figures 22 to 33, the difference between the errors of the two filters is very similar. This is expected, since the noise injected into the system for both of these filters was Gaussian. Had the error been non-Gaussian, I would expect the higher moment filter to be able to perform better. However, one difference

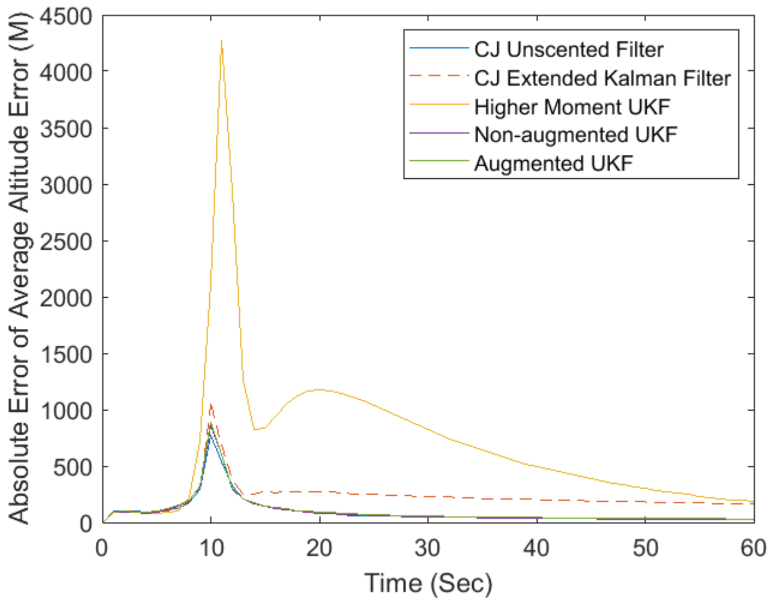

Figure 21: A comparison between my non-augmented UKF, augmented UKF, and higher moment UKF with Crassidis and Junkins's example UKF and EKF from [5]. As shown, my non-augmented and augmented UKFs perform identically to their UKF, but the higher moment UKF performs considerably worse.

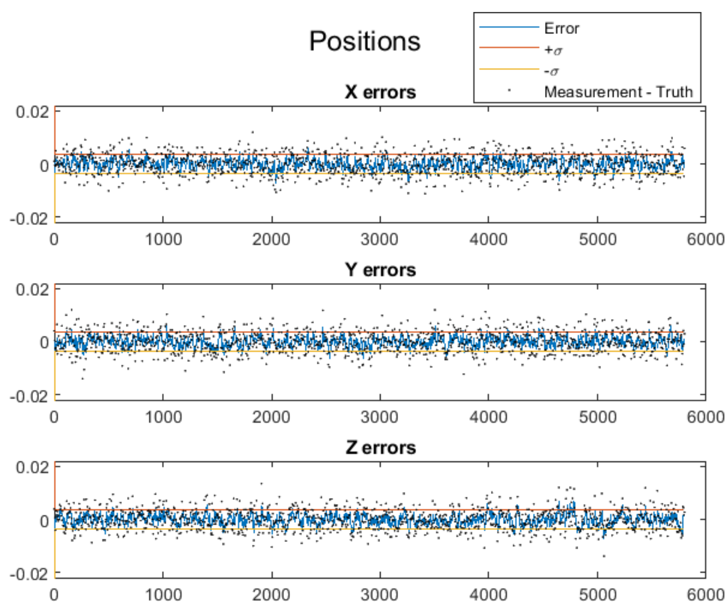

Figure 22: The errors of the satellite position components of the estimate of the **non-augmented UKF**

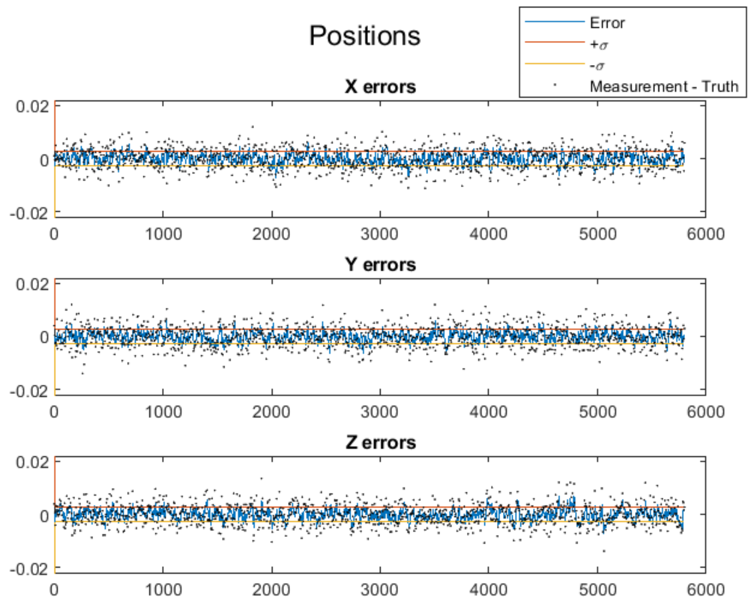

Figure 23: The errors of the satellite position components of the estimate of the **higher moment UKF**

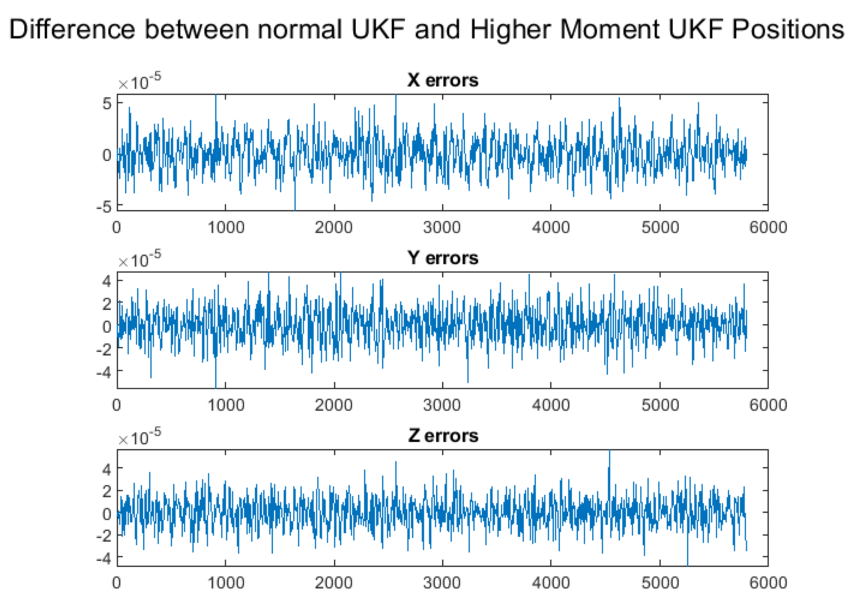

Figure 24: The difference in satellite position estimates of the **non-augmented UKF** and **higher moment UKF**

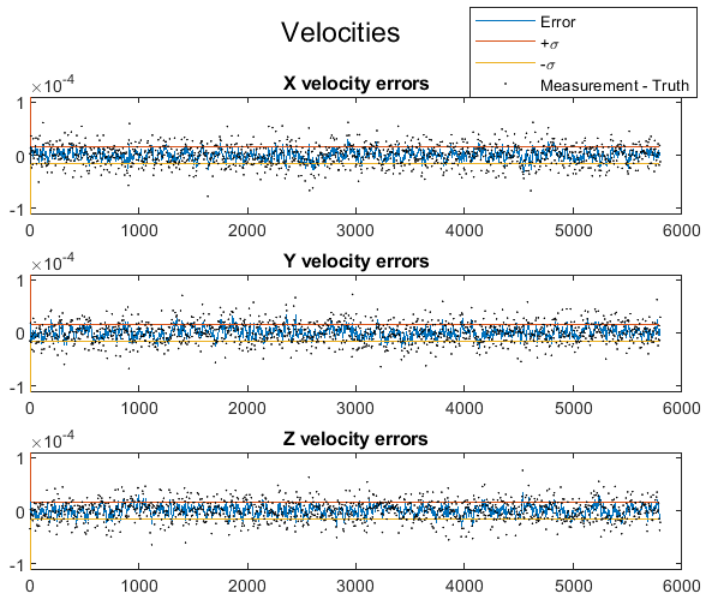

Figure 25: The errors of the satellite velocity components of the estimate of the **non-augmented UKF**

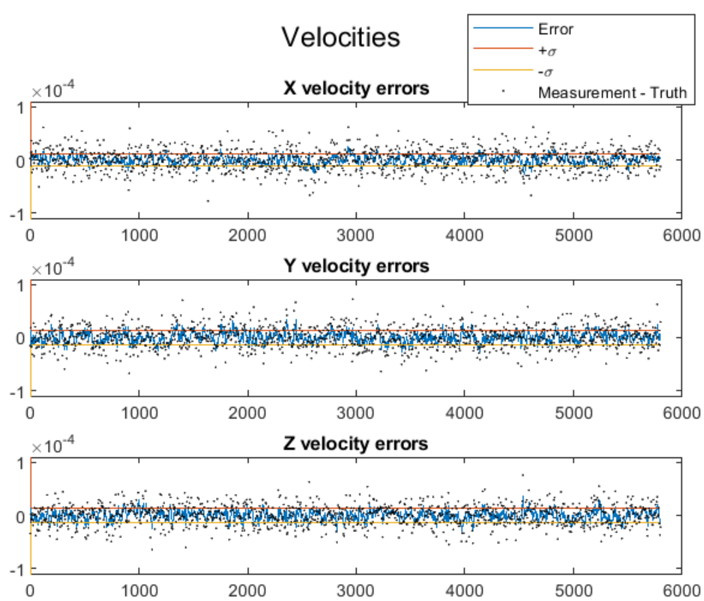

Figure 26: The errors of the satellite velocity components of the estimate of the **higher moment UKF**

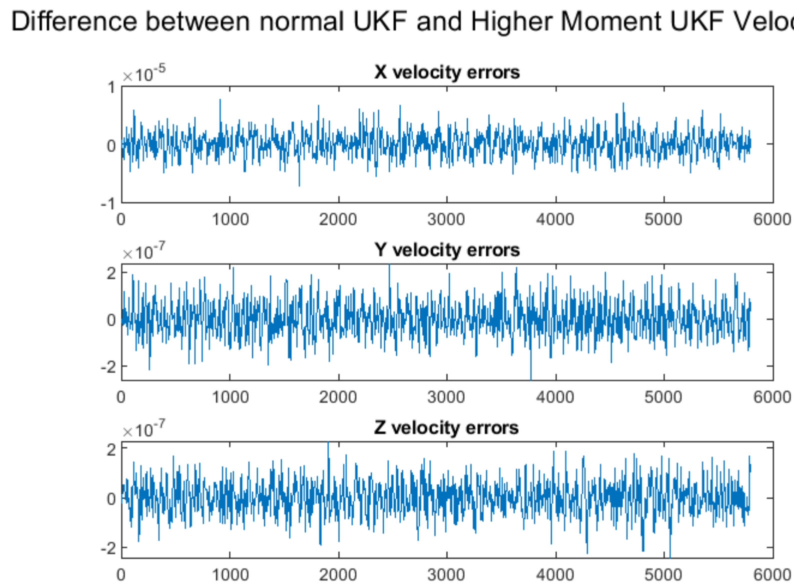

Figure 27: The difference in satellite velocity estimates of the **non-augmented UKF** and **higher moment UKF**

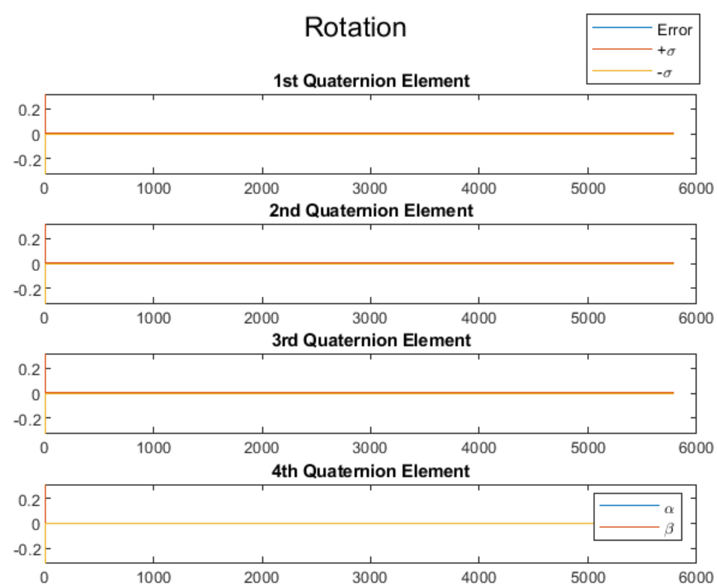

Figure 28: The errors of the satellite orientation components of the estimate of the **non-augmented UKF**

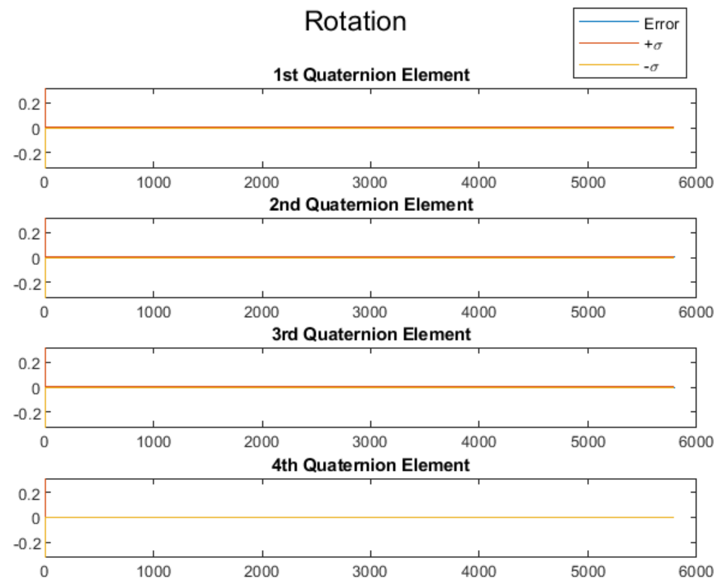

Figure 29: The errors of the satellite orientation components of the estimate of the **higher moment UKF**

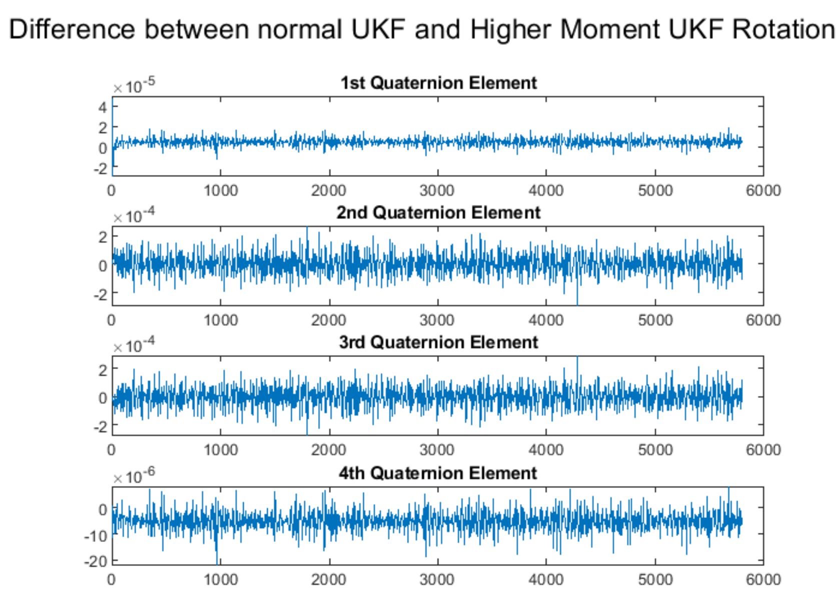

Figure 30: The difference in satellite orientation estimates of the **non-augmented UKF** and **higher moment UKF**

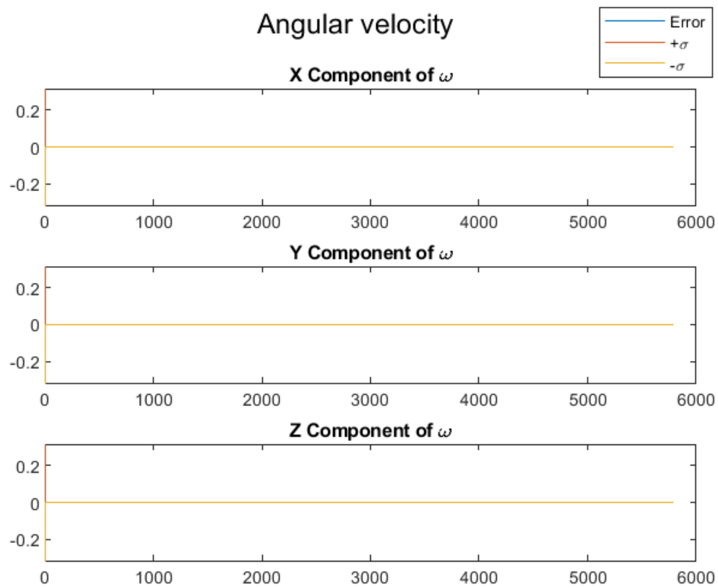

Figure 31: The errors of the satellite angular velocity components of the estimate of the **non-augmented UKF**

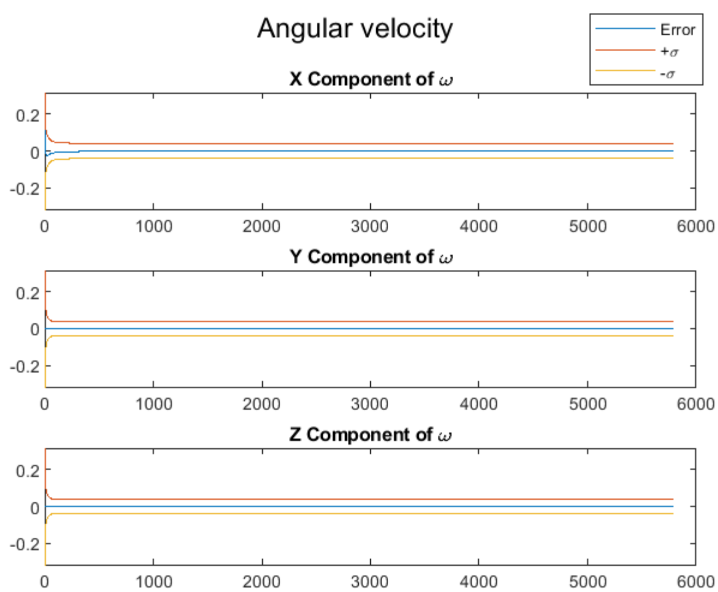

Figure 32: The errors of the satellite angular velocity components of the estimate of the **higher moment UKF**

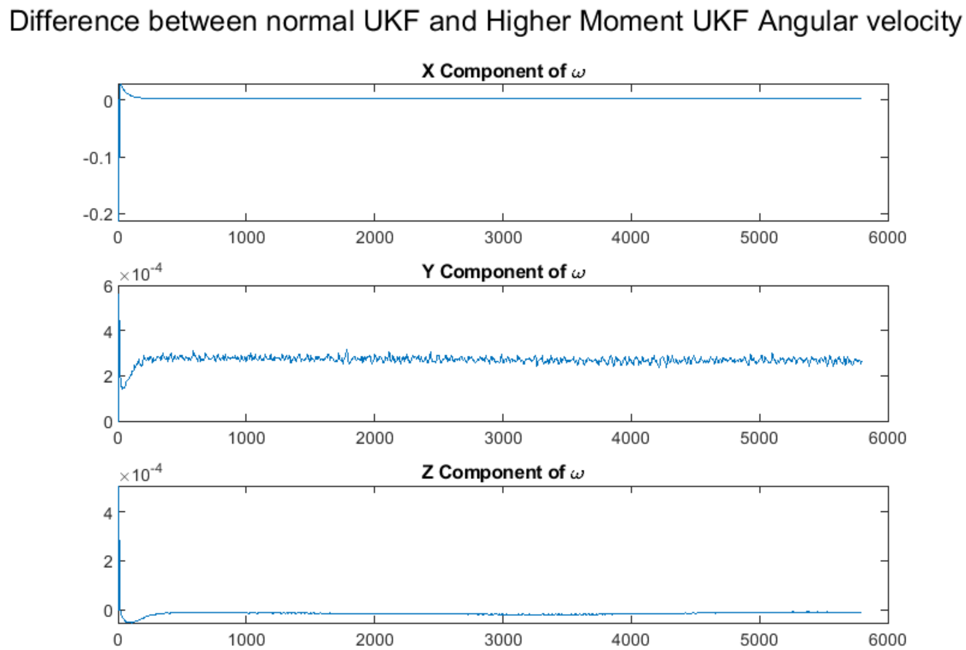

Figure 33: The difference in satellite angular velocity estimates of the **non-augmented UKF** and **higher moment UKF**

between two filters is their run time. The normal UKF takes about 4 minutes to run over a full orbit, and the higher moment UKF takes about 10 minutes to run over a full orbit. Therefore, for Gaussian noise,

References

[1] S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," in Proceedings of the IEEE, vol. 92, no. 3, pp. 401-422, March 2004, doi: 10.1109/JPROC.2003.823141.

[2] Yuanxin Wu, Dewen Hu, Meiping Wu and Xiaoping Hu, "Unscented Kalman filtering for additive noise case: augmented versus nonaugmented," in IEEE Signal Processing Letters, vol. 12, no. 5, pp. 357-360, May 2005, doi: 10.1109/LSP.2005.845592.

[3] Yi Cao (2020). Learning the Unscented Kalman Filter (https://www.mathworks.com/matlabcentral/fileexchange/18217-learning-the-unscented-kalman-filter), MATLAB Central File Exchange. Retrieved May 18, 2020.

[4] Ksenia Ponomareva et al, "A new unscented Kalman filter with higher order moment-matching" July 2010

[5] John L. Crassidis and John L. Junkins, *Optimal Estimation of Dynamic Systems*. 2012