PARAMETRIC VALIDATION OF THE LIGHT-CURVE-BASED ATTITUDE DETERMINATION CAPABILITIES OF SPOOK

M. Kuhn, D. Vallverdú Cabrera, and J. Utzmann

Airbus Defence & Space, Germany, Email: {moritz.kuhn, david.cabrera_vallverdu.external, jens.utzmann}@airbus.com

ABSTRACT

The Special Perturbations Orbit determination and Orbit analysis toolKit (SPOOK), developed by Airbus Defence and Space, has recently been expanded with attitude determination capabilities based on the light curves of space objects. A sequential estimation approach is taken, using an Adaptive Gaussian Mixture Unscented Kalman Filter (AGMUKF) to cope with the nonlinear nature of the attitude estimation problem. This paper investigates the performance and validity of this method under various conditions and parameters using simulated light curves. Performance is found to be generally satisfactory across the parameters tested, though certain cases such as undersampled light curves expectedly remain challenging. Future avenues of research as well as prospects for further improvements to the AGMUKF in SPOOK are explored.

Keywords: Object Characterization; Gaussian Mixtures Filter; Unscented Kalman Filter; Attitude Determination; Light Curve.

1. INTRODUCTION

The number of man-made Resident Space Objects (RSOs) in orbit around the Earth is greater than ever before and continues increasing at an unprecedented pace. This development poses new challenges in managing active satellites and mitigating the danger posed by space debris. Space Surveillance and Tracking (SST) refers to a set of activities aimed at gathering the information needed to overcome these challenges. This information includes the attitude state of RSOs, which is critical for endeavours such as active debris removal or precise orbit determination. Due to their relatively small size and large distance from observers, it is usually not possible to obtain resolved optical images of RSOs for use in attitude determination. However, RSOs' light curves – their brightness as a function of time – can be used to reconstruct their attitude states (e.g. [8, 11, 17]).

The Special Perturbations Orbit determination and Orbit analysis toolKit (SPOOK) is a software suite for SST developed at Airbus Defence and Space GmbH, Germany [10]. It offers various functionalities for SST, including sensor simulation, observation planning, orbit determination and propagation. SPOOK has recently been upgraded with light curve simulation and light curve-based attitude determination capabilities. The latter comprise two distinct approaches. The first is a shape-independent approach based on the asteroid community's epoch method, which reconstructs an object's spin axis and period from the variation of its apparent period; this method has been shown to work in observation geometries that present sufficient variation of the apparent period [15]. The second is a sequential estimation method employing an Adaptive Gaussian Mixture Unscented Kalman Filter (AGMUKF) [16]. This filter aims to cope with the considerable nonlinearity of the attitude determination problem by representing arbitrarily complex and generally non-Gaussian attitude state probability density functions (PDFs) as Gaussian mixtures. The number of Gaussian kernels in the mixture is dynamically adjusted in order to maintain a balance between faithful representation of the PDF and computational efficiency. This paper aims to explore the validity and performance of the AGMUKF implementation for attitude determination in SPOOK over a wide parameter space.

Light curves for various test cases are generated using SPOOK's light curve simulation capabilities as input for the AGMKUKF attitude determination algorithm. In this way, this paper assesses the impact of different parameters on the convergence and success of the attitude determination. The parameters taken into consideration include an RSO's orbital regime, the observation geometry (phase angle), Signal-to-Noise Ratio (SNR) of the measurements, length and sampling density of the analysed light curve, as well as the object's attitude state itself. Additionally, the impact of an object's shape on attitude determination is examined by simulating different simple shapes representing different classes of objects.

2. THE ADAPTIVE GAUSSIAN MIXTURE UNSCENTED KALMAN FILTER

This section aims to briefly recapitulate the principal ideas of the AGMUKF as applied to the attitude determination problem by Vallverdú Cabrera et al. 2022 [16]

in contrast to the "classical" Unscented Kalman Filters (UKF) and Particle Filter (PF); an in-depth treatment of the problem and description of the filtering algorithm can be found there.

The use of sequential nonlinear estimation approaches to attitude determination is well documented in the literature, with UKFs [17] and PFs [1, 5, 8] having seen considerable use. However, the computational requirements of the Particle Filter are significant, while the UKF fails to capture the complex, generally non-Gaussian and possibly multimodal state PDFs that arise due to the nonlinearity of the attitude determination problem. The AGMUKF aims to address both of these challenges: a sufficiently fine Gaussian mixture approximates any PDF arbitrarily closely, while the adaptive variation of the number of kernels in the mixture may lower the computational load compared to Particle Filter variants.

The AGMUKF represents the estimated state PDF as a mixture of Gaussian distributions. Their number is adaptively adjusted during the filter run. The decision to split kernels if the filter becomes too nonlinear, refining the mixture, is made based on a nonlinearity index developed for this purpose [16], while kernels are merged based on their Kullback-Leibler distance [12], coarsening the mixture. Each of the Gaussian kernels in essence follows the steps of the normal UKF, of which thorough descriptions are given elsewhere in the literature [6]. As a derivative of the Kalman filter, the UKF also consists of a propagation step and an update step, which are carried out alternatingly while the filter runs. The propagation step encompasses propagating the filter state from the last update forward to the time of the next measurement; the update step updates the state with the information from a new measurement. In the case of attitude determination, the measurement model is the primary source of nonlinearity [16]. Therefore, the AGMUKF refines the Gaussian mixture before the measurement update to capture the nonlinearity of the measurement model and merges kernels after the measurement update to preserve computational resources.

3. VALIDATION METHODOLOGY

In this work, the SPOOK implementation of an AGMUKF for attitude determination is validated using simulated light curves generated by SPOOK's own light curve simulation functionality [15], using a simulated observer modeled on the Airbus Robotic Telescope (ART) [14] in Extremadura, Spain.

The high dimensionality of the parameter space of the attitude determination problem and the significant computational requirements render it impractical to exhaustively explore this parameter space. Instead, parameter combinations that are likely to be of interest in real-world use cases were chosen for validation. For of these cases, Monte Carlo studies were performed. Due to the large number of cases and the computational burden of nonlinear filtering, only 10 iterations could be performed for each case; this should suffice to identify problematic scenarios, but with more time and computational resources, the results' reliability could certainly be improved further. The performance measure used in this work is the root mean square error (RMSE) of the filter's estimated final state with respect to the true state:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N_{\mathrm{MC}}} \sum_{i=1}^{N_{\mathrm{MC}}} (\hat{x} - x)^T (\hat{x} - x)} \qquad (1)$$

where N_{MC} is the number of Monte Carlo iterations and x and $\hat{x}$ are the true and estimated state, respectively. The filter state consists of an object's instantaneous orientation and angular velocity; the RMSE is calculated separately for both components.

The cases considered include objects in different orbital regimes: low Earth orbit (LEO), medium Earth orbit (MEO) and geostationary orbit (GEO). To get as close to realistic observation scenarios as possible, the orbital properties of the simulated objects for these tests were based on three real RSOs that were observable from Extremadura, Spain between 4:00 and 5:00 UTC on 2022/11/12. They are ESA's ENVISAT (SatCat ID: 27386) for LEO, GPS BIIF-10 (40730) for MEO and the commercial communications satellite ASTRA 2F (38778) for GEO. The impact of the object's initial orientation, rate of rotation and the direction of the rotation axis was also explored. Additionally, the performance of the algorithm for different signal-to-noise ratios (SNRs) and phase angles of the simulated observations was investigated. Finally, the influence of the light curve duration and sampling density was explored. The latter was done only for the GEO case because of that case's straightforward observation geometry and to keep the computational effort manageable. The initial attitude state of each MC run was drawn from a Gaussian distribution of a covariance chosen as a middle ground between values used in other works [8, 17]:

$$P_0 = \mathrm{diag}[0.1\ 0.1\ 0.1\ 2.5 \times 10^{-5}\ 2.5 \cdot 10^{-5}\ 2.5 \cdot 10^{-5}]$$

In order to investigate the impact of an object's shape on the filter's performance, three different shape models were used. They are intended to be representative, at least in a qualitative sense, of different common classes of RSOs. The cube, which measures 4 m in each dimension, is intended to serve as a baseline without any large features and a relatively high degree of symmetry, but can also represent satellites or even debris of a similarly simple shape. The winged cuboid, consisting of a cube with 4 m sides and wings of 2 m by 8 m, represents the large class of boxy satellites with solar arrays. Finally, the rocket body consists of a cylinder 9 m in length, with a radius of 1.5 m, capped by two cones of 1.5 m height. The shapes are depicted in Figure 1. As is fairly common in space object characterisation efforts (e.g. [17, 18]), all three shape models use a Cook-Torrance reflection model as a compromise between physical realism

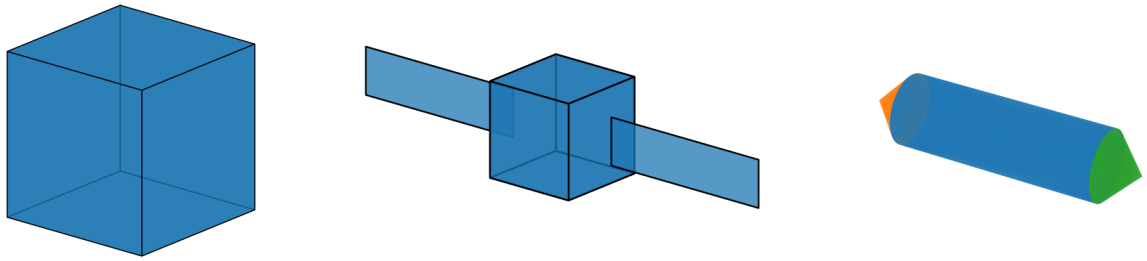

(a) (b) (c)

Figure 1: Object shapes used in this work (not to scale): a) cube; b) winged cuboid with specularly reflecting solar panels; c) cylindrical rocket body based on the shape used by Wetterer et al. [17] and Vallverdú Cabrera et al. [16]. See text for more information on dimensions and reflective properties.

and computational tractability [2]; however, the model's parameters are different for each shape. The cube has a diffuse reflection fraction $d = 0.8$ and diffuse and specular reflectances $\rho = F_0 = 0.5$, on all sides. The winged cuboid shares the same properties on all faces of the main body, but the wings, representing solar panel arrays, reflect more specularly with $d = 0.2$, $\rho = 0.3$, $F_0 = 0.9$. The rocket body is based on the shape model used by Vallverdú Cabrera et al. 2022 [16], which in turn is based on the one used by Wetterer et al. 2009 [17]. All shapes share a slope parameter $m = 0.17$.

4. RESULTS AND DISCUSSION

In this section, the results of applying the AGMUKF to the attitude determination scenarios described in Section 3 are presented in the form of the RMSE of the filter's estimated attitude and angular velocity in the final state, i.e. at the end of the filter run. It should be noted that the angular velocity RMSE given is the absolute error.

Figure 2 shows the results for a simulated LEO object, for each of the three shapes described above and rotation periods from just under 10 s to over 2500 s, corresponding to logarithmically equispaced initial angular velocities from 0.0025 rad/s to 0.64 rad/s. The angular velocities are aligned with the z-axis of the Earth-Centered Inertial (ECI) reference frame, i.e. they point North; where applicable, the objects rotate around one of their short axes. RSOs in LEO are generally only observable for short periods at a time; the same is true for ENVISAT, whose orbital properties the simulated objects inherit. In the early morning of 2022/11/12, the time of all simulations in this paper, ENVISAT was observable by ART for just over three minutes, so the simulated light curve in this case is limited to 180 s. Indeed, the RMS deviations in Figure 2 rise sharply for rotation periods longer than that. In addition, the errors for the winged cuboid are very large, on the order of tens of degrees, indicating filter divergence for most of the winged cuboid runs

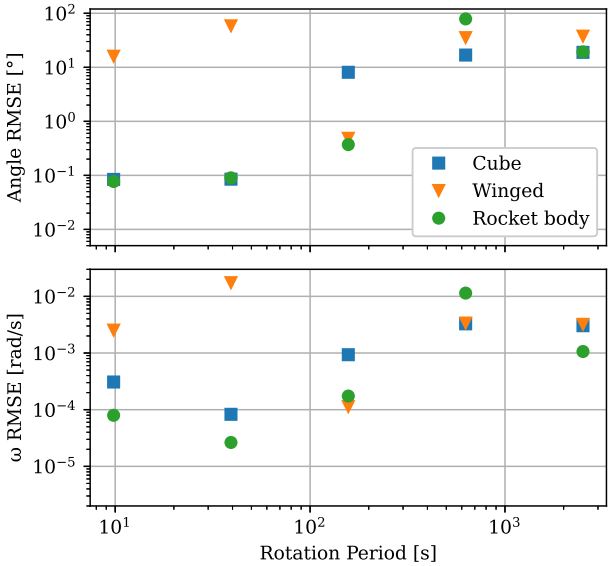

Figure 2: Final state RMSE for a LEO object with different shapes, as a function of the object's rotation period.

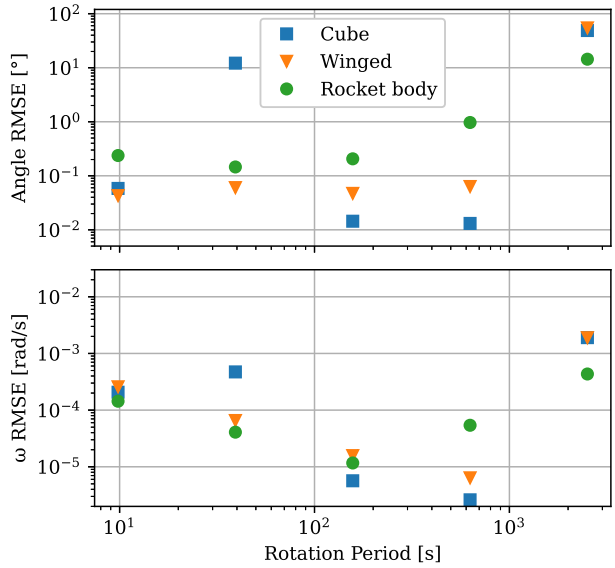

Figure 3: Final state RMSE for a MEO object with different shapes, as a function of the object's rotation period.

in LEO. This may be due to the pronounced specular reflections from the wings saturating the simulated detector when the object is in a low orbit and therefore close to the observer.

The results for the MEO scenario are shown in Figure 3, which is otherwise analogous to Figure 2. The orbital properties in this case were taken from the GPS BIIF-10 satellite, which was observable for several hours on 2022/11/12. For computational reasons, the simulated light curve is however 900 s long, as are all others used in this work, unless otherwise stated. Once again, it can be seen that the AGMUKF results are reasonably close to the simulated truth, with the exception of the longest rotation period, which is longer than the light curve duration. One of the errors for the cube shape, with a rotation period of 40 s, is also extremely high, indicating the divergence of at least one of the Monte Carlo iterations for that scenario. One possible reason for this is that for phase angles greater than 90°, observation geometries exist in which all sunlit faces of a cube-shaped object are hidden from the observer, making the object unobservable for a time and possibly causing filter divergence. In this case, the phase angle is just over 90°.

Figure 4 shows the analogous graph for an object in GEO. Unlike the previous cases, this orbit - based on that of the ASTRA 2F satellite - is equatorial, so the object's angular velocity, aligned with the ECI z-axis, is orthogonal to the orbital plane. As before, the winged cuboid and rocket body rotate around one of their short axes. This is a very favourable set of conditions, and indeed the RMS errors are very small here. Once again, the filter's failure to converge for the longest rotation period can be seen.

The scenario displayed in Figure 5 is identical to the pre-

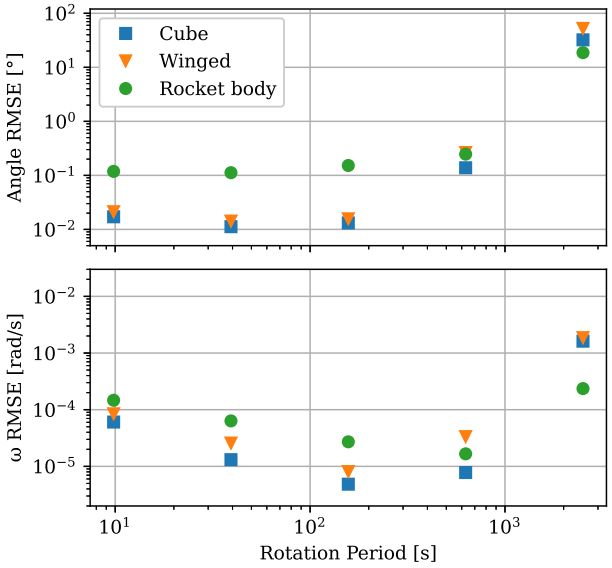

Figure 4: Final state RMSE for a GEO object with different shapes, as a function of the object's rotation period. In this case, the object's long axis(for the winged cube and rocket body) lies in the orbital plane, while the angular velocity vector is orthogonal to it.

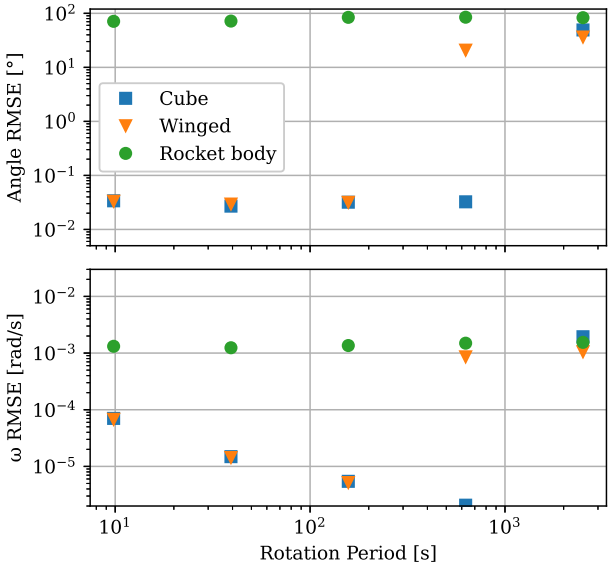

Figure 5: Final state RMSE for a GEO object with different shapes, as a function of the object's rotation period. The object's long axis lies in the orbital plane and is also the rotation axis.

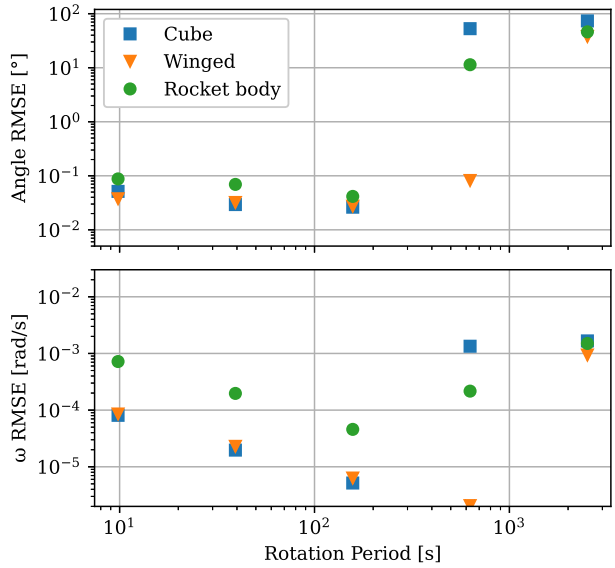

Figure 6: Final state RMSE for a GEO object with different shapes, as a function of the object's rotation period. The angular velocity vector lies in the orbital plane; the objects are rotating around one of their shorter axes.

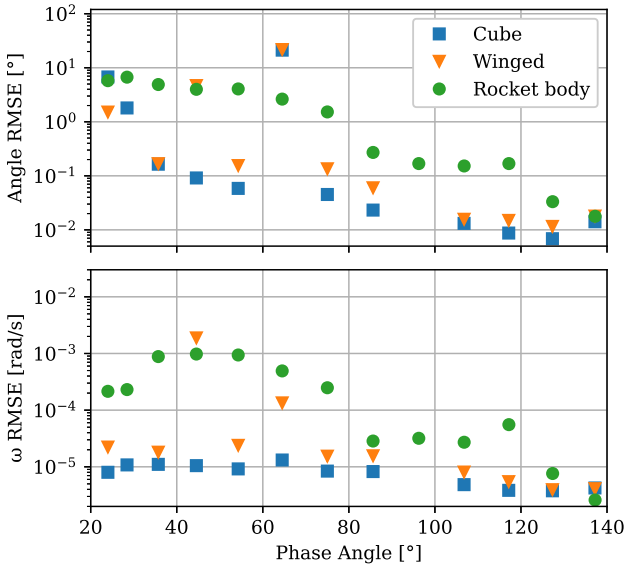

Figure 7: Final state RMSE as a function of phase angle for an object in GEO, rotating in the orbital plane with a rotation period of 150 s. The phase angle given is the average over the 900 s simulation period.

vious one, except the objects now rotate around their longest axis. The rocket body is axially symmetric around this axis, producing an entirely featureless light curve. Indeed the algorithm is entirely unable to ascertain the rocket body's state in this configuration. Another variation of this scenario is shown in Figure 6, but now the angular velocity vector lies in the orbital plane; the objects are again rotating around one of their shorter axes. This is again a rather favourable scenario, but the inability to resolve states with rotation periods longer than the light curve duration remains.

The first GEO scenario presented, shown in Figure 4, was then taken as a baseline for further tests because of its favourable observation conditions. In Figure 7, the impact of the simulated object's phase angle on the RMS error was investigated. The phase angle naturally varies over time for an object in GEO, but over the 900 s time span considered here, the deviation from the average phase angle shown in Figure 4 amounts to only a few degrees. In this case, the object's rotation period was chosen to be 150 s. The figure shows a clear trend: the RMS error becomes smaller, i.e. the filter estimate becomes better, the larger the phase angle is. In addition, both the cube and the winged cube exhibit some diverged filter runs around a phase angle of 65°. The reason for this could not be ascertained, but the fact that both cube-like shapes show this phenomenon at least lends credibility to it not being coincidence, but indeed related to the shapes, phase angle and/or observation geometry.

The influence of the simulated measurement's SNR on the filter performance was also investigated, taking the same scenario as above as a base. The SNR was var-

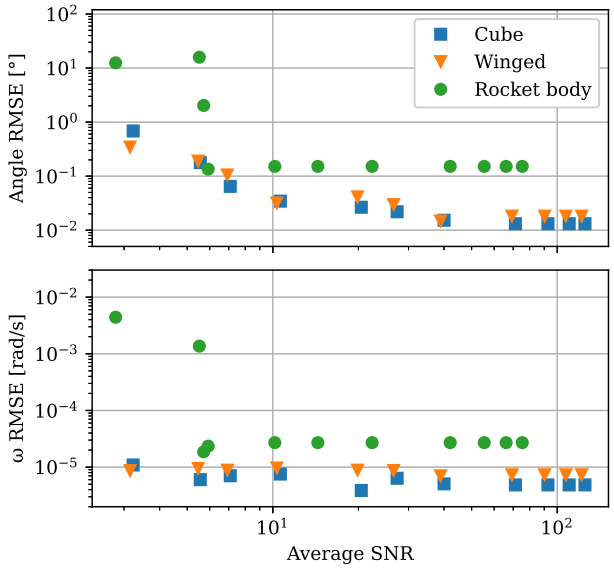

Figure 8: Final state RMSE as a function of average signal-to-noise ratio of the simulated observations for an object in GEO, rotating in the orbital plane with a rotation period of 150 s.

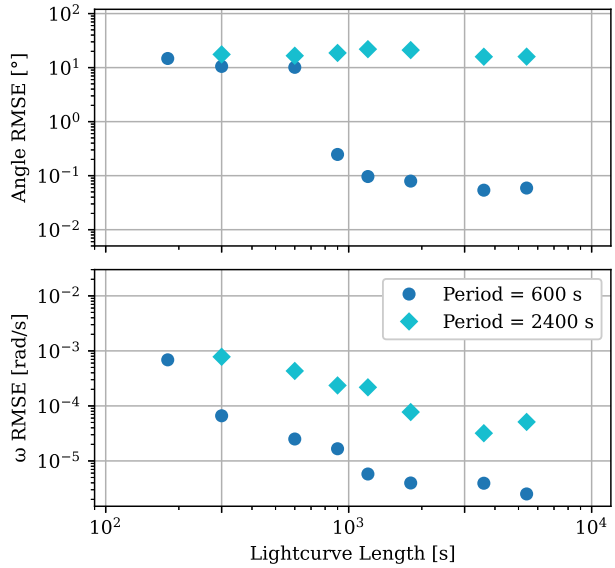

Figure 9: Final state RMSE as a function of light curve duration for objects in GEO, using the "rocket body" shape and rotating in the orbital plane with rotation periods of 600 s and 2400 s, respectively.

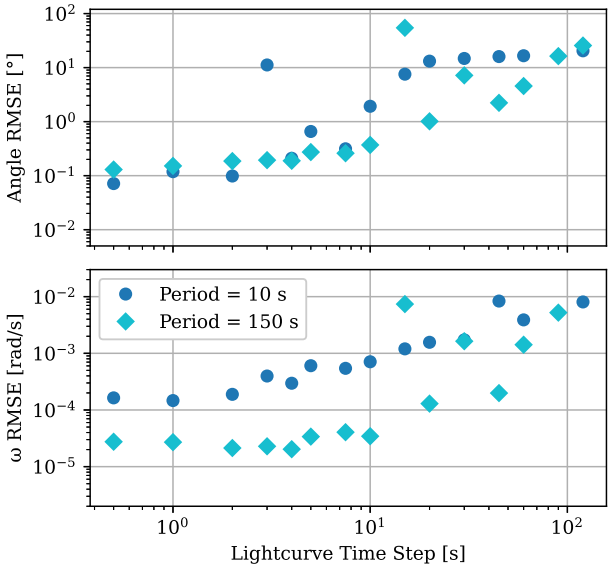

Figure 10: Final state RMSE as a function of light curve sampling density for objects in GEO, using the "rocket body" shape and rotating in the orbital plane with rotation periods of 10 s and 150 s, respectively.

ied by varying the total signal received by the simulated telescope. The results are shown in Figure 8. Note that depending on the object shape, attitude state and observation geometry, the SNR varies considerably over the course of a light curve; plotted here is the average over the light curve length of 900 s. The RMS errors remain quite small down to fairly low SNR, but increase considerably towards very low SNRs. For the rocket body shape in particular, they rapidly become very large below a certain average SNR. This is essentially an artifact of considering the *average* SNR. The rocket body exhibits very large changes between maximum and minimum brightness, so at an average SNR at which the other shapes are still well observed, the rocket body is not observable for a considerable fraction of the time.

The final two parameters investigated in this paper are the length and sampling density of the light curve. The results of the former are shown in Figure 9. The conditions are identical to those of the other GEO scenarios described above, but with only the rocket body shape model considered. Two rotation periods of 600 s and 2400 s were tested. The results confirm the previous observation that rotation periods longer than the lightcurve cause the filter to diverge, signified by final state RMS deviations of tens of degrees. Notably, these large errors persist for the longer rotation period even for light curve lengths of more than 5000 s. It is possible that for these slow rotators and short time steps, each new measurement provides less new information than is lost in each filter step due to added noise, numerical inaccuracies and even inaccuracies in the AGMUKF's repeated merging and splitting operations. Figure 10 shows the impact of the light curve sampling density on the filter results. The scenario is otherwise the same as above, but the two rotation periods considered are now 10 s and 150 s, respectively. Similar to above, time steps longer than the rotation period result in filter divergence.

5. FUTURE PROSPECTS

This work focused on exploring the performance of SPOOK's AGMUKF for attitude determination in various observation scenarios. A different set of parameters that was not taken into account here are those of the filter itself: the maximum number of kernels, splitting and merging thresholds, or the tunable parameters of the Unscented Transform [6]. Thorough investigation of these parameters and their effect on the filter's function may enable further gains in performance and reliability.

Airbus Defence & Space also operates ART, a 0.4 m aperture optical telescope with a wide field of view that is also capable of measuring RSO light curves [15]. The extension of the present investigation to real-world measured data is therefore a natural next step. Unlike simulations, the true attitude state of real objects is usually not known; this holds especially true in the case of uncontrolled objects or debris. One class of objects for which the attitude is known at any point in time are scientific telescopes. ESA's International Gamma-Ray Astrophysics Laboratory (INTEGRAL) might be one possible target for such an investigation, as its varied pointing profile and highly elliptic orbit may enable some of the observation geometries explored in this paper (or qualitatively similar ones) to be recreated.

To put the RMSE values obtained in this work into context, it would be useful to be able to compare them against a measure of the observability of the attitude state [3] or the Cramér-Rao lower bound for the error, as has been done for other filtering problems [13]. Work on the observability analysis of the attitude determination problem is documented in the literature [4], including on simplified systems of lower dimension [9]. However, analysing the observability of the full, nonlinear attitude determination problem is not trivial and the calculation of the Cramér-Rao lower bound for the AGMUKF presents both implementational and computational challenges. However, recent advances in nonlinear observability may enable a more in-depth analysis than this one in the future [7].

6. CONCLUSIONS

The performance of the Adaptive Gaussian Mixture Unscented Kalman Filter (AGMUKF) for attitude determination in SPOOK was explored under a wide variety of parameters. These include three different shape models, three orbital regions, SNR, phase angle, light curve sampling and duration and the attitude state itself. While performance is satisfactory over large parts of the parameter space, the algorithm struggles with cases that have historically been challenging for light curve-based methods, such as pronounced specular flares and short or undersampled light curve data. Future work may include the application to real-world measured data, further characterisation of the filter and its application to attitude determination and an analysis of the observability of the full, highly nonlinear attitude determination problem.

ACKNOWLEDGMENTS

The first author thanks Giovanni Cirillo and Martin Michel for their support in working with the SPOOK software suite.

REFERENCES

1. Coder, R. D., Holzinger, M. J., and Linares, R. (2018). "Three-Degree-of-Freedom Estimation of Agile Space Objects Using Marginalized Particle Filters". *Journal of Guidance, Control, and Dynamics* 41.2, pp. 388–400.
2. Cook, R. L. and Torrance, K. E. (1982). "A Reflectance Model for Computer Graphics". *ACM Trans. Graph.* 1.1, pp. 7–24.
3. Friedman, A. M. and Frueh, C. (2022). "Observability of Light Curve Inversion for Shape and Feature Determination Exemplified by a Case Analysis". *The Journal of the Astronautical Sciences* 69.2, pp. 537–569.
4. Hinks, J. C., Linares, R., and Crassidis, J. L. (2013). "Attitude Observability from Light Curve Measurements". AIAA Guidance, Navigation, and Control (GNC) Conference. Boston, MA.
5. Holzinger, M. et al. (2012). "Attitude Estimation for Unresolved Agile Space Objects with Shape Model Uncertainty". *Advanced Maui Optical and Space Surveillance Technologies Conference*. Ed. by S. Ryan, 25, p. 25.
6. Julier, S. and Uhlmann, J. (2004). "Unscented Filtering and Nonlinear Estimation". *Proceedings of the IEEE* 92.3, pp. 401–422.
7. Kim, J. W. and Mehta, P. G. (2021). "A Dual Characterization of Observability for Stochastic Systems". *IFAC-PapersOnLine* 54.9, pp. 659–664.
8. Linares, R., Crassidis, J., and Jah, M. (2014). "Particle filtering light curve based attitude estimation for non-resolved space objects". *Advances in the Astronautical Sciences* 152, pp. 119–130.
9. Linares, R. and Crassidis, J. L. (2016). "Dynamic Observability Analysis for Attitude, Angular Velocity, Shape, and Surface Parameters". 26th AAS/AIAA Space Flight Mechanics Meeting. Napa, United States.
10. Pedone, G. et al. (2021). "SPOOK: A tool for space objects catalogue creation and maintenance supporting space safety and sustainability". *Acta Astronautica* 188, pp. 89–98.
11. Piergentili, F., Santoni, F., and Seitzer, P. (2017). "Attitude Determination of Orbiting Objects from Lightcurve Measurements". *IEEE Transactions on Aerospace and Electronic Systems* 53.1, pp. 81–90.
12. Runnalls, A. (2007). "Kullback-Leibler Approach to Gaussian Mixture Reduction". *Aerospace and Electronic Systems, IEEE Transactions on* 43, pp. 989–999.
13. Tulsyan, A. et al. (2012). "Performance Assessment of Nonlinear State Filters". *IFAC Proceedings Volumes* 45.15, pp. 371–376.
14. Utzmann, J., Vesselinova Dimitrova, M. G., and Rodriguez Fernandez, O. (2019). "Airbus Robotic Telescope". 1st NEO and Debris Detection Conference. Darmstadt, Germany.
15. Vallverdú Cabrera, D. et al. (2021). "Integration of Attitude Characterization in a Space Debris Catalogue Using Light Curves". 8th European Conference on Space Debris (virtual). Darmstadt, Germany.
16. Vallverdú Cabrera, D., Utzmann, J., and Förstner, R. (2022). "The adaptive Gaussian mixtures unscented Kalman filter for attitude determination using light curves". *Advances in Space Research*.
17. Wetterer, C. J. and Jah, M. (2009). "Attitude Determination from Light Curves". *Journal of Guidance, Control, and Dynamics* 32.5, pp. 1648–1651.
18. Wetterer, C. J. et al. (2014). "Refining Space Object Radiation Pressure Modeling with Bidirectional Reflectance Distribution Functions". *Journal of Guidance, Control, and Dynamics* 37.1, pp. 185–196.