
Towards Automated Satellite Conjunction Management with Bayesian Deep Learning

Francesco Pinto
University of Oxford
francesco.pinto@eng.ox.ac.uk

Giacomo Acciarini
University of Strathclyde
giacomo.acciarini@gmail.com

Sascha Metz
Technische Universität Darmstadt
smetz1405@gmail.com

Sarah Boufelja
IBM
boufelja.sarah@gmail.com

Sylvester Kaczmarek
Imperial College London
sylvester.kaczmarek@gmail.com

Klaus Merz
European Space Agency
klaus.merz@esa.int

José A. Martínez-Heras
European Space Agency
jose.antonio.martinez.heras@esa.int

Francesca Letizia
European Space Agency
francesca.letizia@esa.int

Christopher Bridges
University of Surrey
c.p.bridges@surrey.ac.uk

Atılım Güneş Baydin
University of Oxford
gunes@robots.ox.ac.uk

Abstract

After decades of space travel, low Earth orbit is a junkyard of discarded rocket bodies, dead satellites, and millions of pieces of debris from collisions and explosions. Objects in high enough altitudes do not re-enter and burn up in the atmosphere, but stay in orbit around Earth for a long time. With a speed of 28,000 km/h, collisions in these orbits can generate fragments and potentially trigger a cascade of more collisions known as the Kessler syndrome. This could pose a planetary challenge, because the phenomenon could escalate to the point of hindering future space operations and damaging satellite infrastructure critical for space and Earth science applications. As commercial entities place mega-constellations of satellites in orbit, the burden on operators conducting collision avoidance manoeuvres will increase. For this reason, development of automated tools that predict potential collision events (conjunctions) is critical. We introduce a Bayesian deep learning approach to this problem, and develop recurrent neural network architectures (LSTMs) that work with time series of conjunction data messages (CDMs), a standard data format used by the space community. We show that our method can be used to model all CDM features simultaneously, including the time of arrival of future CDMs, providing predictions of conjunction event evolution with associated uncertainties.

1 Introduction

The risk of collisions between man-made objects in space is growing. Moreover, due to the increased amount of space debris, the expected growth of the space sector, and the planned launch of megaconstellations, the situation is expected to worsen [1], [2], [3], [4], [5]. With the growth of space debris population, a few collisions between objects in orbit might trigger a chain reaction that could pollute

the orbit to the point of making it inaccessible, a scenario known as the Kessler syndrome [6], [7]. Unless appropriate techniques are developed to handle this problem, this could severely endanger the satellite infrastructure that is increasingly essential in Earth sciences applications. Indeed, besides being fundamental to space sciences, navigation, and telecommunication, satellite constellations are nowadays central in a wide range of geoscience-related applications in climate science, Earth observation, disaster detection, and monitoring of oceans, ice, groundwater, and landmasses [8].

The current space population in Earth’s orbit is monitored via the global Space Surveillance Network (SSN): a network of optical and radar sensors handled by the US Strategic Command (USSTRATCOM). This network observes and tracks man-made objects in space, while producing a publicly available running catalog of unclassified space objects.¹ For collision assessment purposes, SSN observations are fed to a propagator (i.e., a physics simulator that predicts the evolution of the state of the objects over a time period, typically seven days) and each satellite (usually called *target*) is screened against the rest of the catalogue population in order to detect close approaches, which are referred to as “conjunctions”. If a conjunction is predicted between the target and another object (usually called *chaser*), a conjunction data message (CDM) is automatically issued to the owner/operator (O/O) of the satellite, containing information about the event at the time of closest approach (TCA), such as the probability of collision and covariances of the state variables of the two objects. Furthermore, information relevant to how all the above data are determined can be also present (e.g., details of the orbit simulator and dynamical environment used). For the week leading up to TCA, more SNN observations are collected and screened and CDM updates are issued with updated information about TCA (usually around three CDMs per day) to the O/O as the conjunction event evolves. O/Os are generally alerted to the event and, up to one day prior to TCA, they must make a decision about whether to perform collision-avoidance manoeuvres, which are expensive due to resource constraints. The last CDM received is usually considered the best knowledge of the state of target and chaser at TCA. The collision risk assessment during this decision process is currently performed manually. Several methods exist to establish if pairs of objects lead to a conjunction event in the propagated time-interval [9], [10] and to obtain collision risk estimates [11, 12].

A major barrier to the application of machine learning techniques to aid the collision management process has been the lack of publicly available large-scale CDM datasets due to the presence of sensitive information about the assets of private companies and governments. The only publicly available CDM dataset, which we call the Kelvins dataset, has been provided by the European Space Agency (ESA) in the context of the Spacecraft Collision Avoidance Challenge,² where participants developed techniques to predict the final risk of collision for each event. The dataset contains CDMs collected by ESA from 2015 to 2019, with modifications to anonymise the data (e.g., absolute time information and some of the orbital state elements are not given—only relative states).

In this work we develop a recurrent neural-network architecture that works with sequences of CDMs representing conjunction events, where the network operates on the whole set of numerical features present in the CDM format. Specifically, our architecture is a stacked LSTM³ [13] trained in a Bayesian deep learning [14, 15, 16, 17] setting in order to provide uncertainty information at test time. We train this network using the Kelvins dataset. This model can be used in two modes: (1) predicting the contents of the next CDM given all previous CDMs in an event currently unfolding; (2) predicting the whole sequence of future CDMs until TCA, by feeding each predicted CDM to back to the network and predicting the next, until the last CDM is predicted by the network. These modes are analogous to the usage of probabilistic language models, where complete sentences can be sampled from a generative model primed with some initial symbols (characters or words) [18, 19].

2 Related Work

Conjunction Management Starting from the CDMs, current collision avoidance strategies include a collision risk assessment between pairs of objects, optionally augmenting the data with external information sources. Typical risk assessment procedures have two steps. First, the states of the two objects at the time of closest approach are predicted. This prediction is generally represented by means of a probability distribution over the state variables. Then, the information contained in such distributions is leveraged to evaluate the probability of collision. Sometimes, CDMs might contain information about probability of collision and methods used for computing it.

¹<https://www.space-track.org/> (September 2020).

²<https://kelvins.esa.int/collision-avoidance-challenge/> (August 2020)

³Long short-term memory




Figure 1: *Left*: At training time the LSTM is optimized to reproduce the time-shifted CDM sequences in training data. *Right*: At test time, given a sequence of CDMs for the current event, the trained LSTM is used to predict future CDMs that characterise the event evolution.

In the current practices, propagation is performed by either using linear/linearised methods (e.g., Kalman filter) [20], or by using simulation-based techniques (e.g., Monte Carlo simulation) [21]. Semi-analytic techniques have been also developed in order to trade-off the accuracy of simulation-based techniques with the speed of linearised methods. In this context, some of the most popular semi-analytic techniques are: differential algebra with Gaussian mixture model [22], polynomial chaos expansion [23], generalized polynomial algebra [24], unscented transform [25].

Given the output of the propagation, the collision risk (i.e., the probability that the two objects will collide at time of closest approach) can be evaluated using several different algorithms [26, 27, 10, 28, 29]. The operators then use this risk information combined to their own risk assessment and evaluation to decide whether to maneuver the satellite for avoiding conjunctions.

Bayesian Deep Learning and Monte Carlo Dropout The problem of quantifying the uncertainty associated with the estimates produced by machine learning methods is paramount for their usage in safety-critical applications [30, 31, 15] like satellite conjunction management. Typical deep learning approaches produce only point estimates, without providing a quantitative assessment of the reliability of the prediction being made. Given the lack of interpretability of deep learning approaches, it is difficult to ascertain whether the space operator should trust the outputs of the neural network. For this reason, Bayesian deep learning methods output distributions of predictions, which allow to quantify this uncertainty (e.g., in terms of entropy or variances of the output [15]). This is attained by placing a prior distribution $p(\omega)$ over the parameters ω of a differentiable function implemented with a neural network f , and by finding the posterior distribution $p(\omega|X, Y)$ (where X, Y represent the training dataset input and output respectively) that capture the most likely functions given data, which can be used to predict the output y^* for a new input x^* . The predictive distribution is: $p(y^*|x^*, X, Y) = \int p(y^*, f^*)p(f^*|x^*, \omega)p(\omega|X, Y)df^*d\omega$. Unfortunately, $p(\omega|X, Y)$ cannot be evaluated analytically. To tackle this issue, variational inference is applied: an easy-to-evaluate variational distribution $q(\omega)$ is defined, and used to approximate $p(\omega|X, Y)$ by minimizing the KL divergence $KL(q(\omega)||p(\omega|X, Y))$ resulting in the approximate predictive distribution: $p(y^*|x^*, X, Y) = \int p(y^*, f^*)p(f^*|x^*, \omega)q(\omega)df^*d\omega$.

Many variational approximation schemes are possible. In this paper, we choose an approach called Monte Carlo dropout [16, 32] which has been popular in many fields (e.g. autonomous driving [15], medical diagnosis [33]) thanks to its implementation simplicity. In this setting applying dropout at training time is equivalent to performing variational inference using a Bernoulli approximating distribution. In order to obtain distributions over new data points y^* , it is sufficient to run the network multiple times on the same input with the dropout at test time.

3 Conjunction Event Evolution Analysis with Bayesian Deep Learning

While the contestants of the Spacecraft Collision Avoidance Challenge were required to predict the final risk and perform binary classification (i.e., high- versus low-risk event) [34], this approach does not reflect the two-step nature of the established risk assessment procedures. In particular, the risk prediction and the decision of whether to classify a CDM sequence as either high or low risk is made by a neural network. This decision cannot be explained in terms of physical variables, and hence cannot be trusted by operators. Indeed, the probability of collision is a quantity obtained by integrating physical quantities contained in CDMs (e.g., using [26, 27, 10, 28, 29]) and its evaluation is carefully designed to capture the uncertainty contained in the distributions of the states of the two objects at TCA. Another problem with a classification approach in this setting is the severe data imbalance [35] between the high- and low-risk classes, since high-risk events and collisions are very rare in actual real data. For these reasons, we propose to use generative modeling to model

distributions of future CDMs conditioned on previous CDMs in an event, and include all the numerical features within the CDM data format in model inputs and outputs. This allows us to fully make use of the available training data to aid representation learning [36], and also grounds our approach in the already established CDM data format for which there are tools and techniques developed by the space community. In particular, we employ a stack of LSTMs trained by minimising the mean squared error (MSE) between predicted next CDMs and the ground-truth, using each event’s CDM sequence as input and the same sequence shifted by one time step as the target, so that in each time step the network learns to predict the next CDM conditioned on all previous CDMs (Figure 1, left).

In order to produce a distribution over predicted CDMs, representing uncertainty, we apply Monte Carlo dropout [16] and sample multiple predictions at test time. Referring to a conjunction event (a CDM sequence) as $\gamma_{1:t} = \{\gamma_1, \dots, \gamma_t\}$, where γ_t is the CDM at time t , the network at test time allows us to sample the next predicted CDM $\tilde{\gamma}_{t+1} \sim p(\tilde{\gamma}_{t+1}|\gamma_{1:t})$ conditioned on $\gamma_{1:t}$, the CDMs since the beginning of the event up to current time t . In addition, the predictions can be extended further into the future by sampling subsequent $\tilde{\gamma}_{t+2} \sim p(\tilde{\gamma}_{t+2}|\gamma_{1:t}, \tilde{\gamma}_{t+1})$, and then $\tilde{\gamma}_{t+3} \sim p(\tilde{\gamma}_{t+3}|\gamma_{1:t}, \tilde{\gamma}_{(t+1):(t+2)})$ and so on, until a complete event prediction until TCA is obtained as $\tilde{\gamma}_{1:T} = \{\gamma_1, \dots, \gamma_t, \tilde{\gamma}_{t+1}, \dots, \tilde{\gamma}_T\}$ (Figure 1, right).

4 Experiments

For our experiments, we use the Kelvins dataset,⁴ and discard the features `c_rcs_estimate`, `t_rcs_estimate` (as these contain many NaN values and we used `c_span`, `t_span` to leverage the size of the objects), the parameters related to the solar flux (e.g., F10, F3M, SSN, AP) because they also contain many NaN values and are not normally included in CDMs (i.e., they have been augmented by the organizers of the competition). Furthermore, we drop `mission_id`, `c_object_type` because they remain constant through the whole event sequence; moreover, the `mission_id` cannot generalize when applied to test data with other mission ids. Then we drop all the CDMs with NaN values, and CDMs with abnormal values in the variances of the state variables (i.e., `t_sigma_r` > 20, `c_sigma_r` > 1000, `t_sigma_t` > 2000, `c_sigma_t` > 100000, `t_sigma_n` > 10, `c_sigma_n` > 450). After these steps, given the wide scale variation among the variables involved, to stabilise the training, we normalize each feature (except the categorical ones) by subtracting the mean and dividing the result by the standard deviation. For the next CDM prediction approach, we consider only event sequences with at least two CDMs. Starting from the 199,082 CDMs in the original dataset, after the cleaning procedure 156,668 CDMs remain, which are grouped into 11,387 events. We randomly split the full Kelvins dataset into a training and test split with the test split amounting for the 15% of the total data (9,367 events in the training set, with 133,435 CDMs). With this data, we train a network that handles 52 CDM features in total, including miss distance, relative velocity and position, and target and chaser covariance matrices.⁵ The model used is a stack of two 256-dimensional LSTMs, whose hidden state is passed through a ReLU activation and fed to a linear layer to obtain the final prediction. Dropout is applied to all but the output layers of the network with 0.2 dropout rate. The overall model has 857,140 parameters, and is trained for 500 epochs with the Adam optimizer [37], learning rate 10^{-4} and batch size 128. We will make the code of our implementation publicly available.

We report the mean squared error (MSE) of a baseline that predicts the variables at the next time step to be the same of the previous time step. While using the risk contained in the last CDM received was the core strategy leveraged by the participants to the Space Collision Avoidance challenge [34] to discriminate between high and low risk events, in order to predict the physical quantities involved, it performs poorly ($MSE = 0.2433$). Even only with one Monte Carlo dropout sample we obtain better performance ($MSE = 0.1753$). With $n = 50$ samples, we obtain further improvements ($MSE = 0.1419$). An important remark is that, since the arrival time of the next CDM is modelled as one of the variables predicted, our model can give a prediction also in terms of when the next CDM will arrive, further informing the operators about the decision to be made. Figure 2 gives an example event evolution from the validation set, showing that the model has learned non-trivial patterns about the behaviour of the physical variables involved.

⁴Description of data features: <https://kelvins.esa.int/collision-avoidance-challenge/data/>

⁵Note that when training with CDM datasets other than Kelvins, there can be more features which were not available in Kelvins data. We are planning to show examples of this in an upcoming paper.




Figure 2: *Top row*: the evolution of the values of a selection of CDM features through a whole event; the prediction is made using only the first CDM. *Bottom row*: as new CDMs arrive, the predictions made using only the first CDM are confirmed (blue dots), showing that the LSTM can predict non-trivial patterns describing event evolution.

5 Conclusion

In this work we introduced a new recurrent neural network architecture to model the time evolution of conjunction events, using sequences of the standard CDM data format in use by the space community. We demonstrated that non-trivial time evolution dynamics of conjunction events can be learned, and the trained models can be used to predict event outcomes at test time. Furthermore we applied Monte Carlo dropout in order to produce distributions over the CDM features describing uncertainties associated with the two monitored objects at the time of closest approach. This method opens up a new way of working with CDM data sequences using machine learning, which we believe is an essential step in the automation and scaling of conjunction management operations in the near future.

6 Broader Impact

Given the increasing amount of space debris and the plans for new mega-constellations of satellites in orbit, the problem of automated collision management will become paramount to preserve the global satellite infrastructure network that currently provides critical services (e.g., weather data monitoring, global communications, Earth observation). Furthermore, collisions in the increasingly crowded low Earth orbit might trigger a cascade of collisions (Kessler Syndrome) that could hinder future space operations and trap us on Earth for years to come. Hence it is crucial to develop tools that will help maintain a managed low Earth orbit population. We do not foresee our research to have negative ethical impacts on society.

Acknowledgments

This work has been enabled by Frontier Development Lab (FDL) Europe, a public–private partnership between the European Space Agency (ESA), Trillium Technologies and the University of Oxford, and supported by Google Cloud. We would like to thank Dario Izzo and Moriba Jah for sharing their technical expertise and James Parr, Jodie Hughes, Leo Silverberg, Alessandro Donati for their support. AGB is supported by EPSRC/MURI grant EP/N019474/1 and by Lawrence Berkeley National Lab.

References

- [1] A Rossi, A Cordelli, P Farinella, L Anselmo, and C Pardini. Long term evolution of the space debris population. *Advances in Space Research*, 19(2):331–340, 1997.
- [2] Donald J Kessler. Collisional cascading: The limits of population growth in low earth orbit. *Advances in Space Research*, 11(12):63–66, 1991.
- [3] L Anselmo, A Rossi, and C Pardini. Updated results on the long-term evolution of the space debris environment. *Advances in Space Research*, 23(1):201–211, 1999.
- [4] S Le May, S Gehly, BA Carter, and S Flegel. Space debris collision probability analysis for proposed global broadband constellations. *Acta Astronautica*, 151:445–455, 2018.
- [5] B Bastida Virgili, JC Dolado, HG Lewis, J Radtke, H Krag, B Revelin, C Cazaux, CAMILLA Colombo, R Crowther, and M Metz. Risk to space sustainability from large constellations of satellites. *Acta Astronautica*, 126:154–162, 2016.

- [6] Donald J Kessler and Burton G Cour-Palais. Collision frequency of artificial satellites: The creation of a debris belt. *Journal of Geophysical Research: Space Physics*, 83(A6):2637–2646, 1978.
- [7] J-C Liou and Nicholas L Johnson. Instability of the present leo satellite populations. *Advances in Space Research*, 41(7):1046–1053, 2008.
- [8] National Academies of Sciences, Engineering, and Medicine. *Thriving on our changing planet: A decadal strategy for Earth observation from space*. National Academies Press, 2019.
- [9] Felix R Hoots, Linda L Crawford, and Ronald L Roehrich. An analytic method to determine future close approaches between satellites. *Celestial Mechanics*, 33(2):143–158, 1984.
- [10] Heiner Klinkrad. Space debris. *Encyclopedia of Aerospace Engineering*, 2010.
- [11] K Merz, B Bastida Virgili, V Braun, T Flohrer, Q Funke, H Krag, S Lemmens, and J Siminski. Current collision avoidance service by esa’s space debris office. In *Proceedings 7th European Conference on Space Debris, Darmstadt, Germany*, pages 18–21, 2017.
- [12] V Braun, T Flohrer, H Krag, K Merz, S Lemmens, B Bastida Virgili, and Q Funke. Operational support to collision avoidance activities by esa’s space debris office. *CEAS Space Journal*, 8(3):177–189, 2016.
- [13] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [14] David JC MacKay. A practical Bayesian framework for backpropagation networks. *Neural Computation*, 4(3):448–472, 1992.
- [15] Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5574–5584. Curran Associates, Inc., 2017.
- [16] Yarin Gal and Zoubin Ghahramani. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1050–1059, New York, New York, USA, 20–22 Jun 2016. PMLR.
- [17] Andrew Gordon Wilson. The case for Bayesian deep learning. *arXiv preprint arXiv:2001.10995*, 2020.
- [18] Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. A neural probabilistic language model. *Journal of Machine Learning Research*, 3(Feb):1137–1155, 2003.
- [19] Andriy Mnih and Yee Whye Teh. A fast and simple algorithm for training neural probabilistic language models. *arXiv preprint arXiv:1206.6426*, 2012.
- [20] David A Vallado. *Fundamentals of astrodynamics and applications*, volume 12. Springer Science & Business Media, 2001.
- [21] Eric B Ford. Quantifying the uncertainty in the orbits of extrasolar planets. *The Astronomical Journal*, 129(3):1706, 2005.
- [22] Zhen-Jiang Sun, Ya-Zhong Luo, Pierluigi Di Lizia, and Franco Bernelli Zazzera. Nonlinear orbital uncertainty propagation with differential algebra and Gaussian mixture model. *SCIENCE CHINA Physics, Mechanics & Astronomy*, 62(3):34511, 2019.
- [23] Brandon A Jones, Alireza Doostan, and George H Born. Nonlinear propagation of orbit uncertainty using non-intrusive polynomial chaos. *Journal of Guidance, Control, and Dynamics*, 36(2):430–444, 2013.
- [24] Massimiliano Vasile, Carlos Ortega Absil, and Annalisa Riccardi. Set propagation in dynamical systems with generalised polynomial algebra and its computational complexity. *Communications in Nonlinear Science and Numerical Simulation*, 75:22–49, 2019.
- [25] Nagavenkat Adurthi and Puneet Singla. Conjugate unscented transformation-based approach for accurate conjunction analysis. *Journal of Guidance, Control, and Dynamics*, 38(9):1642–1658, 2015.
- [26] Kyle T Alfriend, Maruthi R Akella, Joseph Frisbee, James L Foster, Deok-Jin Lee, and Matthew Wilkins. Probability of collision error analysis. *Space Debris*, 1(1):21–35, 1999.

- [27] Russell P Patera. General method for calculating satellite collision probability. *Journal of Guidance, Control, and Dynamics*, 24(4):716–722, 2001.
- [28] Russell P Patera. Satellite collision probability for nonlinear relative motion. *Journal of Guidance, Control, and Dynamics*, 26(5):728–733, 2003.
- [29] F Kenneth Chan et al. *Spacecraft Collision Probability*. Aerospace Press El Segundo, CA, 2008.
- [30] Yihui He, Chenchen Zhu, Jianren Wang, Marios Savvides, and Xiangyu Zhang. Bounding box regression with uncertainty for accurate object detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [31] Gregory P. Meyer, Ankit Laddha, Eric Kee, Carlos Vallespi-Gonzalez, and Carl K. Wellington. Lasernet: An efficient probabilistic 3d object detector for autonomous driving. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12669–12678, 2019.
- [32] Yarin Gal and Zoubin Ghahramani. Bayesian convolutional neural networks with Bernoulli approximate variational inference. *arXiv:1506.02158*, 2015.
- [33] Terrance DeVries and Graham W. Taylor. Leveraging uncertainty estimates for predicting segmentation quality. *CoRR*, abs/1807.00502, 2018.
- [34] Thomas Uriot, Dario Izzo, Luis Simoes, Rasit Abay, Nils Einecke, Sven Rebhan, Jose Martinez-Heras, Francesca Letizia, Jan Siminski, and Klaus Merz. Spacecraft collision avoidance challenge: design and results of a machine learning competition. *arXiv preprint arXiv:2008.03069*, 2020.
- [35] S. Wang and X. Yao. Multiclass imbalance problems: Analysis and potential solutions. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 42(4):1119–1130, 2012.
- [36] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- [37] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.