
Spacecraft Collision Risk Assessment with Probabilistic Programming

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

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

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

Over 34,000 objects bigger than 10 cm in length are known to orbit Earth. Among them, only a small percentage are active satellites, while the rest of the population is made of dead satellites, rocket bodies, and debris that pose a collision threat to operational spacecraft. Furthermore, the predicted growth of the space sector and the planned launch of megaconstellations will add even more complexity, therefore causing the collision risk and the burden on space operators to increase. Managing this complex framework with internationally agreed methods is pivotal and urgent. In this context, we build a novel physics-based probabilistic generative model for synthetically generating conjunction data messages, calibrated using real data. By conditioning on observations, we use the model to obtain posterior distributions via Bayesian inference. We show that the probabilistic programming approach to conjunction assessment can help in making predictions and in finding the parameters that explain the observed data in conjunction data messages, thus shedding more light on key variables and orbital characteristics that more likely lead to conjunction events. Moreover, our technique enables the generation of physically accurate synthetic datasets of collisions, answering a fundamental need of the space and machine learning communities working in this area.

1 Introduction

The low Earth orbit (LEO) environment is becoming increasingly crowded [1], with an estimate of more than half a million uncontrolled orbiting objects bigger than 1 cm in length.¹ Furthermore, the space population is expected to increase more in the future, due to the expansion of the space sector and the advent of megaconstellations [2], [3], [4], [5]. This scenario poses a danger to operational satellites, scientific missions, and the future of human access to space since the risk of collision among spacecraft will increase, with an increased risk of a cascade of collisions that could affect the whole space population—the Kessler syndrome [6]. Moreover, the burden on human operators that have to maneuver these satellites can become complex and unmanageable [7]. Several guidelines have been developed by the Inter-Agency Space Debris Coordination Committee (IADC) in recent years to mitigate collision risk in the future [8]. Meanwhile, space agencies, private operators and the research community have been investigating ways to tackle this problem [9], [10], [11], [12], [13], [14]. Current collision avoidance strategies are based on techniques that assess the risk of collision between pairs of objects, starting from a series of conjunction data messages (CDMs) associated with a potential collision event.

A CDM is automatically issued by the US-based Combined Space Operations Center (CSpOC), which uses the data provided by the US Strategic Command (USSTRATCOM) that routinely observes and tracks objects in space via a global Space Surveillance Network (SSN). The CDMs are issued to the owner/operator (O/O) of the involved spacecraft. Usually, they contain the following information about a potential collision event: (1) the predicted orbital position and velocity of the two objects at their time of closest approach (TCA); (2) information related to the physical simulation environment used for making the predictions; (3) the propagated uncertainty information about the two objects’ physical states, in the form of covariances. Additionally, they might also optionally contain other details about the probability of collision and the method used to compute it. The data contained in CDMs are calculated by CSpOC through a range of different algorithms, involving the propagation of uncertainties from current observational data up until the time of closest approach [15], [16], [17], [18], [19]. Once uncertainties are propagated to TCA, several algorithms can be used for the evaluation of collision risk (i.e., the probability that the two objects will collide at TCA) [20], [21], [22], [23], [24]. The operators then use this information combined with their own risk assessment and data (e.g., telemetry) to decide whether to undertake a maneuver to avoid the conjunction.

Usually, for each collision event a time series of CDMs covering one week is released. At the European Space Operations Center (ESOC) of the European Space Agency (ESA), CDM updates are received roughly every eight hours. However, they do not necessarily contain updated information for both the objects, since very often only one of the two objects (the *target*, often an operational satellite) is well observed, whereas the other one (the *chaser*, often a piece of debris) is not [7]. Target and chaser objects have different frequency of observation updates and uncertainty characteristics, creating a complex setting where the tools of probabilistic machine learning (ML) and simulation-based inference [25] are applicable. A major obstacle in applying ML in this setting is the scarcity of publicly available data, where space agencies are reluctant to share data due to security concerns, and a severe class imbalance in any conjunction dataset as collisions are rare events.

In this work, we build a physics-based probabilistic generative model to simulate spacecraft orbits, observations, and the CDM generation process, using probabilistic programming [26]. We calibrate the model using real data, and use it to perform Bayesian inference. We show that this approach could play a pivotal role in helping operators determine the key variables that lead to conjunction events and to make predictions that are reliable (i.e., with associated uncertainties) and explainable (in the context of model-based Bayesian inference). The approach also enables the sampling of large datasets of synthetic but realistic conjunction events that can deliver much-needed publicly available data for the ML and space communities to develop and evaluate new techniques.

2 Probabilistic Programming

Probabilistic programming is a paradigm that allows to specify probabilistic models for data generation and to perform inference in the model conditioned on observed data [27, 26]. Using a universal programming language enhanced by two constructs (`sample` and `observe`), a domain expert can describe a system’s behavior using general-purpose (Turing-complete) programming languages,

¹<https://sdup.esoc.esa.int/discosweb/statistics/> (August 2020)

essentially creating a stochastic simulator in which one can perform automated Bayesian inference [28]. We denote the observable data we want to perform inference on as \mathbf{y} , and assume that \mathbf{y} can be generated by executing a generative program that depends on the values sampled from some latent variables \mathbf{x} . A probabilistic program specifies the joint probability distribution $p(\mathbf{x}, \mathbf{y}) = p(\mathbf{y}|\mathbf{x})p(\mathbf{x})$ in terms of the prior distributions $p(\mathbf{x})$ over the latent variables \mathbf{x} (using the `sample` statements), an interpretable procedure that computes \mathbf{y} given \mathbf{x} , and a likelihood distribution $p(\mathbf{y}|\mathbf{x})$ (using the `observe` statements). The model can be run forward by sampling \mathbf{x} from the priors and applying the generative procedure to generate some data \mathbf{y} . However, probabilistic programming also allows to invert the generative process to perform inference, i.e., given some observed data \mathbf{y} , infer distributions over the latent variables \mathbf{x} that could have generated \mathbf{y} . In other words, inference aims to compute the posterior distribution $p(\mathbf{x}|\mathbf{y})$, which represents our updated belief about the latents \mathbf{x} after incorporating the information from observed data \mathbf{y} .

The order in which the execution of a probabilistic program encounters `sample` statements may be different from one execution to another. Hence, a way to uniquely identify each of their executions is to represent the execution trace of a probabilistic program as a sequence $\mathcal{T} = \{(x_t, a_t, i_t)\}_{t=1}^T$, where x_t , a_t and i_t respectively represent the sampled value, the lexical address of the statement in the program and the instance (i.e. a counter recording the number of times the `sample` statement at address a_t has been called until the time step t) of the t -th entry in the trace. The trace length T and ordering can change at each execution. For each trace, we can collect the T sampled values in a sequence $\mathbf{x} := \{x_t\}_{t=1}^T$. Assuming, for notation simplicity, that the order in which the `observe` statements are encountered is fixed, we can denote the observed values as $\mathbf{y} := \{(y_n)\}_{n=1}^N$.

The joint probability density of an execution trace is therefore $p(\mathbf{x}, \mathbf{y}) := \prod_{t=1}^T f_{a_t}(x_t|x_{1:t-1}) \prod_{n=1}^N g_n(y_n|x_{1:\tau(n)})$, where f_{a_t} is the probability distribution specified by the `sample` statement d_t at address a_t (i.e., the prior conditional density given the sample values $x_{1:t-1}$ before encountering the t -th `sample` statement) and g_n is the probability distribution specified by the n -th `observe` statement (i.e., the likelihood density given the sample values $x_{1:\tau(n)}$ obtained before encountering the n -th `observe` statement, where τ maps the index n of the `observe` statement to the index of the last `sample` statement encountered before the n -th `observe` statement during the execution of the program). In this framework, performing inference means computing an approximation of the posterior $p(\mathbf{x}|\mathbf{y})$ and the expectations over chosen functions ξ under the posterior, i.e., $I_\xi = \int \xi(\mathbf{x})p(\mathbf{x}|\mathbf{y})d\mathbf{x}$. Many inference procedures are available for this purpose [29, 30, 31, 32, 33, 34]. For the preliminary results presented in this paper we use importance sampling (likelihood weighting) [26].

3 Probabilistic Model of Orbits, Conjunctions, and CDM Generation

The proposed stochastic generative model includes the orbital mechanics of the target and chaser and a CDM generation pipeline (incorporating domain knowledge) that mimics the real-world CDM issuing process. An essential part of our probabilistic model are the distributions characterizing the objects and their orbits. We first generate ground truth trajectories of chaser and target by sampling the initial conditions (i.e., initial *two-line elements* [35]) from prior distributions of the LEO population that we constructed based on real data from the public USSTRATCOM catalog and ESA DISCOS database² (Figure 1), and instantiating the objects. We then simulate the objects for seven days using physics-based orbit propagators, which model the object’s state (position and velocity) taking factors such as drag, radiation, and gravitation effects from planetary bodies into account.

Once this is achieved, a check is made to inspect if a conjunction alert is to be issued (by checking if the distance between the two objects is less than a threshold of 5 km). If that is true then the event is flagged as a conjunction and further monitoring is required leading to a series of generated CDMs. In this work, we calibrated the satellite measurement errors using anonymized real conjunction data released by ESA as the Kelvins challenge.³ For doing this, we tuned the initial observation measurement noise up until the propagated covariance values at TCA were matching the data from Kelvins. This calibration was done by inspection (i.e., checking whether the distributions of covariances overlapped) and we plan to assess this in a more systematic fashion in future work. Besides the covariances, one could think of calibrating the probabilistic generative model with other real CDM data. However, the main hindrance for a full calibration of the model is caused by the lack

²<https://discosweb.esoc.esa.int/>

³<https://kelvins.esa.int/collision-avoidance-challenge/data/> (September 2020)

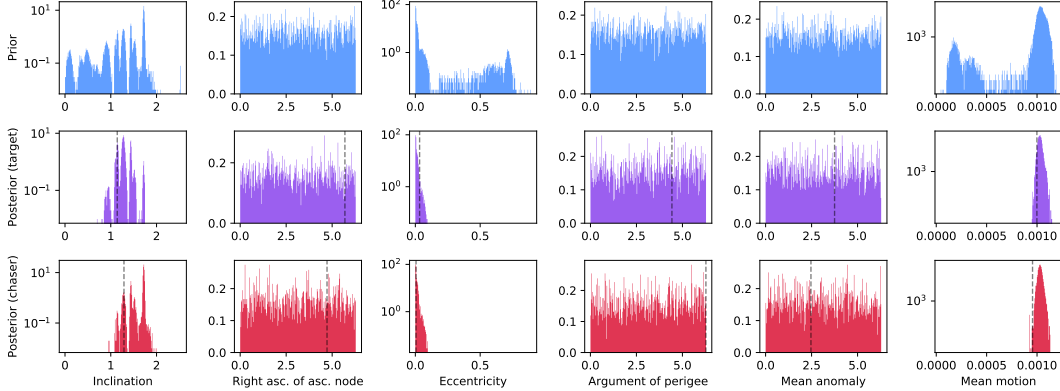


Figure 1: *Top*: a part of the prior distributions for orbital elements of LEO population derived from real object catalogs. *Middle and bottom*: posterior distributions of a subset of latent variables for target and chaser. The posterior is conditioned on a ground truth test event, here highlighted with a dashed vertical line, which was generated through a forward run of the simulator.

of publicly available data, which can only be addressed by collaborating with an owner/operator of satellites.

For issuing CDMs, the state of target and chaser are observed with some associated uncertainty (that vary depending on the sensors used to observe the objects, such as GPS, radar, etc.), then a series of orbit simulation are performed to propagate the state and the associated uncertainties to a future time range. Not only the nominal state, but also the associated initial uncertainties (due to observation errors) are propagated at future times by using a sample-based Monte Carlo approach [36].

Therefore, the essential elements for a probabilistic generative model of CDMs are: the initial state of both chaser and target, the dynamical environment characteristics (that depend on the simulator used for orbit propagation), the associated observation uncertainties (that are propagated into the future) and the method used for propagating them until time of closest approach. For our preliminary experiments we use SGP4 for orbit propagation, a fast low-precision propagator [37]: this is a well known propagator often used for preliminary orbit analysis in LEO, when two-line elements data is available.

4 Experiments

We run experiments on synthetic CDMs to evaluate our pipeline: first we sample a synthetic event from our model, generating a series of CDMs (for a single simulated conjunction) that we designate as observed data for subsequent inference; next, we perform posterior inference over model latents conditioned on the observed CDMs and compare resulting distributions with the known ground truth values of latents that produced the test observation. We use Gaussian likelihoods $\prod_{y \in \mathbf{y}} \mathcal{N}(\mu_y, \sigma_y)$, where μ_y is the value of each simulated observable and σ_y is an empirically-tuned tolerance factor, taking \mathbf{y} to include the TCA and observed object states transformed to Keplerian elements (eccentricity, inclination, semi-major axis). Figure 1 shows a result using importance sampling with 70,000 samples, where we pick a selected subset of latent variables \mathbf{x} , namely the six initial orbital elements of target and chaser (i.e., mean motion, eccentricity, inclination, mean anomaly, argument of perigee and right ascension of ascending node). The presented results were conditioned on the first CDM of a test collision event with 16 CDMs (Figure 2). The ground truth values are shown with a dashed vertical line in Figure 1. As we observe, some distributions are of multimodal nature, suggesting that the estimated posterior can provide insight into multiple different explanations of the observed values. It is interesting to notice how the initial right ascension of the ascending node, argument of perigee, and mean anomaly, do not seem to play a relevant role in explaining the conjunction event. On the other hand, the mean motion, inclination, and eccentricity have posterior distributions that differ significantly from the priors, underlying their pivotal roles in explaining the observed conjunction event. The less relevant role of the first three orbital elements is to be expected since these orbits are almost circular (i.e., both chaser and target have very low eccentricity values). The probabilis-




Figure 2: Orbits of the target and chaser in the conjunction ground truth event that was used as a test case. “X” marks the conjunction.

tic model has been implemented in Python and integrated with Pyprob,⁴ a universal probabilistic programming framework [38]. We are currently investigating a range of likelihood formulations motivated by the physics of the problem and running experiments with Markov chain Monte Carlo [39] inference engines in PyProb that provide rich diagnostics tools to assess the convergence of the approximate posteriors. Moreover, due to the unavailability of public data, we are also working with space operators on evaluating this technique with real CDMs. This will enable us to further investigate and quantify the benefits of such methods on real-world applications.

5 Conclusions

In this preliminary work, we have introduced a novel probabilistic model to simulate conjunction events and the issuing of CDMs associated with each event. The model is calibrated with real data and the prior distributions have been selected to represent the current population of resident space objects in the LEO environment. Our work constitutes the first probabilistic model of synthetic CDMs for collision events by sampling from realistic prior distributions of resident space objects in LEO.

This model can be leveraged for several purposes, on which we are currently working on: (1) Generating the first large dataset of synthetic CDMs calibrated with real data, contributing much needed, large-scale, public data to foster further work by the space and ML communities; (2) Investigating the joint probability space to understand which key variables and configurations lead to, and are useful to predict and prevent, conjunction events; (3) Creating a new generation of risk assessment tools using a Bayesian model-based framework, providing posterior distributions over event characteristics and future uncertainty evolution conditioned on real data observations that space operators have access to during an ongoing event.

Broader Impact

Research that addresses the risk assessment and prevention of collisions in space is of significant short- and long-term benefit to humanity. A Kessler syndrome scenario where a chain reaction of collisions render low Earth orbit inaccessible would be a major environmental catastrophe that would significantly hinder space access for many generations. We believe that simulation-based inference and machine learning will be essential additions to the tool set for addressing these issues in a principled way. We do not foresee any negative societal impacts or harmful use of this research given the current circumstances.

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.

⁴<https://github.com/pyprob/pyprob> (September 2020)

References

- [1] Stijn Lemmens and Francesca Letizia. ESA’s annual space environment report. Technical Report GEN-DB-LOG-00288-OPS-SD, ESA Space Debris Office, 2020.
- [2] John A Kennewell and Ba-Ngu Vo. An overview of space situational awareness. In *Proceedings of the 16th International Conference on Information Fusion*, pages 1029–1036. IEEE, 2013.
- [3] Theodore J Muelhaupt, Marlon E Sorge, Jamie Morin, and Robert S Wilson. Space traffic management in the new space era. *Journal of Space Safety Engineering*, 6(2):80–87, 2019.
- [4] Jonas Radtke, Christopher Kebschull, and Enrico Stoll. Interactions of the space debris environment with mega constellations—using the example of the oneweb constellation. *Acta Astronautica*, 131:55–68, 2017.
- [5] A Rossi, EM Alessi, GB Valsecchi, H Lewis, J Radtke, C Bombardelli, and B Bastida Virgili. A quantitative evaluation of the environmental impact of the mega constellations. In *7th European Conference on Space Debris*. The European Space Agency (ESA), 2017.
- [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] 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.
- [8] Inter-Agency Space Debris Coordination Committee et al. IADC space debris mitigation guidelines. *IADC, September*, 2007.
- [9] R Walker and CE Martin. Cost-effective and robust mitigation of space debris in low earth orbit. *Advances in Space Research*, 34(5):1233–1240, 2004.
- [10] Holger Krag. Consideration of space debris mitigation requirements in the operation of leo missions. In *SpaceOps 2012*, page 1257086. American Institute of Aeronautics and Astronautics, 2012.
- [11] Massimiliano Vasile, Víctor Rodríguez-Fernández, Romain Serra, David Camacho, and Annalisa Riccardi. Artificial intelligence in support to space traffic management. In *Proceedings of the International Astronautical Congress, IAC*. International Astronautical Federation (IAF), 2018.
- [12] Federico Zuiani and Massimiliano Vasile. Preliminary design of debris removal missions by means of simplified models for low-thrust, many-revolution transfers. *International Journal of Aerospace Engineering*, 2012, 2012.
- [13] Kartik Kumar, N Ortiz Gómez, Marko Jankovic, JM Romero Martín, Francesco Topputo, SJ Walker, Frank Kirchner, and Massimiliano Vasile. Agora: Mission to demonstrate technologies to actively remove ariane rocket bodies. In *Proceedings of the International Astronautical Congress, Jerusalem, Israel*, volume 6, page 6, 2015.
- [14] Lourens Visagie, Vaios Lappas, and Sven Erb. Drag sails for space debris mitigation. *Acta Astronautica*, 109:65–75, 2015.
- [15] David A Vallado. *Fundamentals of Astrodynamics and Applications*, volume 12. Springer Science & Business Media, 2001.
- [16] 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.
- [17] 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.

- [18] 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.
- [19] 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.
- [20] 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.
- [21] Russell P Patera. General method for calculating satellite collision probability. *Journal of Guidance, Control, and Dynamics*, 24(4):716–722, 2001.
- [22] Heiner Klinkrad. Space debris. *Encyclopedia of Aerospace Engineering*, 2010.
- [23] Russell P Patera. Satellite collision probability for nonlinear relative motion. *Journal of Guidance, Control, and Dynamics*, 26(5):728–733, 2003.
- [24] F Kenneth Chan et al. *Spacecraft collision probability*. Aerospace Press El Segundo, CA, 2008.
- [25] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 2020.
- [26] Jan-Willem van de Meent, Brooks Paige, Hongseok Yang, and Frank Wood. An introduction to probabilistic programming. *ArXiv*, abs/1809.10756, 2018.
- [27] Frank Wood, Jan Willem Meent, and Vikash Mansinghka. A new approach to probabilistic programming inference. In *Artificial Intelligence and Statistics*, pages 1024–1032, 2014.
- [28] Atılım Güneş Baydin, Lei Shao, Wahid Bhimji, Lukas Heinrich, Lawrence Meadows, Jialin Liu, Andreas Munk, Saeid Naderiparizi, Bradley Gram-Hansen, Gilles Louppe, Mingfei Ma, Xiaohui Zhao, Philip Torr, Victor Lee, Kyle Cranmer, Prabhat, and Frank Wood. Etalumis: Bringing probabilistic programming to scientific simulators at scale. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '19*, New York, NY, USA, 2019. Association for Computing Machinery.
- [29] David Wingate, Andreas Stuhlmüller, and Noah Goodman. Lightweight implementations of probabilistic programming languages via transformational compilation. In Geoffrey Gordon, David Dunson, and Miroslav Dudík, editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 770–778, Fort Lauderdale, FL, USA, 11–13 Apr 2011. PMLR.
- [30] Daniel Ritchie, Anna Thomas, Pat Hanrahan, and Noah Goodman. Neurally-guided procedural models: Amortized inference for procedural graphics programs using neural networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 622–630. Curran Associates, Inc., 2016.
- [31] Frank Wood, Jan Willem Meent, and Vikash Mansinghka. A New Approach to Probabilistic Programming Inference. In Samuel Kaski and Jukka Corander, editors, *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 33 of *Proceedings of Machine Learning Research*, pages 1024–1032, Reykjavik, Iceland, 22–25 Apr 2014. PMLR.
- [32] Tom Rainforth, Christian A. Naesseth, Fredrik Lindsten, Brooks Paige, Jan-Willem van de Meent, Arnaud Doucet, and Frank D. Wood. Interacting particle Markov chain Monte Carlo. In *ICML*, pages 2616–2625, 2016.
- [33] Tuan Anh Le, Atılım Gunes Baydin, and Frank Wood. Inference Compilation and Universal Probabilistic Programming. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1338–1348, Fort Lauderdale, FL, USA, 20–22 Apr 2017. PMLR.

- [34] William Harvey, Andreas Munk, Atilim Gunes Baydin, Alexander Bergholm, and Frank Wood. Attention for inference compilation. *ArXiv*, abs/1910.11961, 2019.
- [35] NASA. Definition of two-line element set coordinate system, 2011. https://spaceflight.nasa.gov/realdata/sightings/SSapplications/Post/JavaSSOP/SSOP_Help/tle_def.html.
- [36] C Yanez, M Gupta, V Morand, and JC Dolado. On the gaussianity validity time for orbital uncertainty propagation. In *ESA NEO and Debris Detection Conference, Darmstadt, Germany*, 2019.
- [37] FR Hoots and RL Roehrich. Models for propagation of norad element sets, spacetrack report no. 3. *Department of Commerce, National Technical Information Service*, 9, 1980.
- [38] Atilim Gunes Baydin, Lei Shao, Wahid Bhimji, Lukas Heinrich, Saeid Naderiparizi, Andreas Munk, Jialin Liu, Bradley Gram-Hansen, Gilles Louppe, Lawrence Meadows, et al. Efficient probabilistic inference in the quest for physics beyond the standard model. In *Advances in Neural Information Processing Systems*, pages 5459–5472, 2019.
- [39] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of Markov Chain Monte Carlo*. CRC press, 2011.