CHALLENGES AND PROGRESS IN PREDICTIVE MAINTENANCE OF LONG-ENDURANCE & LONG-RANGE UNCREWED PLATFORMS

Kenneth Lu[*], Sanja Cvijic[†], Arjuna Balasuriya[‡]

Today's uncrewed platforms are typically operated by humans using remote control to guide every detailed aspect of a mission. However, as missions become more complex, there are many scenarios (particularly in the marine and ground domains) in which operators are unable to communicate with these uncrewed platforms in real time (due to adverse environmental conditions, regulatory restrictions on communications in ecologically sensitive areas, active interference by adversaries, or the desire to remain covert), making it challenging to know the health and status of the platform, and to recalibrate and update mission and control parameters on the fly. This becomes challenging especially when the uncrewed platforms are deployed for long-endurance and long-range missions. Fortunately, significant technical advances in onboard computing power and enhanced sensors offer a pathway to a level of autonomy that can overcome such communications limitations. Predictive maintenance algorithms and digital twins of health and status are becoming essential in preventing unexpected failure and extending the lifespan of uncrewed platforms.

In this paper, we present insights into the importance of predictive maintenance and provide examples of implementation of predictive maintenance in uncrewed platforms such as robotic combat vehicles (RCVs). Moreover, we will discuss how hybrid artificial intelligence (Hybrid-AI) techniques rooted in probabilistic models serve as a foundation for predicting health and status of uncrewed platforms.

INTRODUCTION

Uncrewed platforms provide a bevy of advantages for operators, Warfighters, and decision-makers across the Department of Defense (DoD) and industry abroad. For one, their ability to engage adversaries and navigate hazardous terrain (both underseas and on the ground) far beyond the forward line of troops offers protection to Warfighters operating remotely. Second, their payloads support the deployment and usage of a wide range of command, control, communications, computers, intelligence, surveillance, and reconnaissance (C4ISR) capabilities, providing intel on environments, adversaries, and other information far beyond the natural reach of humans. Finally, combining their combat and reconnaissance capabilities, as well as their overall mobility in rugged terrain, these uncrewed platforms can dramatically reduce the overall labor requirements on humans. However, as uncrewed platforms are increasingly used in hostile, complex, and dynamic environments, maintaining these vehicles to ensure that they are mission capable becomes increasingly difficult. Due to long-endurance and long-range missions in harsh environments, hardware system (sensors and actuators) performance of these uncrewed platforms can deteriorate and due to unreliable communication, the operators will not be aware of these performance changes in the platform affecting the mission objectives. Therefore, it is essential to have health monitoring/predictive maintenance

[*] Scientist, Charles River Analytics, 625 Mount Auburn St. Cambridge, MA 02138.
[†] Senior Scientist, Charles River Analytics, 625 Mount Auburn St. Cambridge, MA 02138.
[‡] Senior Scientist, Charles River Analytics, 625 Mount Auburn St. Cambridge, MA 02138.

https://doi.org/10.52202/075106-0009

and adaptive behavior capabilities on board uncrewed platforms to successfully complete missions in such environments.

Fortunately, thanks to advances in the modern sensing and computing capabilities on board these uncrewed platforms, these predictive analytics are attainable. However, there are several key technical challenges to address for such analytics to be feasible and useful for the end user. First, modern uncrewed systems like robotic combat vehicles (RCVs) will have modular mission payloads (MMPs) in which various capabilities can be swapped in and out. Furthermore, these systems will have other lowest replaceable units (LRUs) that can be changed based on the upgrade cycle/schedule of the system. As a result, an analytic designed for one set of MMPs and LRUs must be *adaptable* to various combinations of capabilities and components. Furthermore, the analytic must *learn* and *evolve* through use, meaning that it must incorporate any new information and data about the system without having to rely on a software engineer completely rewriting the analytic. Second, data sources on uncrewed systems will often be sparse and limited, meaning that the solution must be robust to partial or infrequent information. This requires that the solution be *robust* enough to reason well under uncertainty and avoid overfitting to spurious trends or red-herring correlations. Third, the solution must be *computationally compact* as edge-computing environments on uncrewed systems are limited in hardware and communication bandwidth. For instance, and RCV computing platform may have as little as 20 kilobits (20kb) of wireless bandwidth depending on signal strength, and much of that bandwidth is designated for remote control of the vehicle. Finally, the solution must build trust with the end user. This requires that the solution can *explain* its outputs to the end user in an intuitive way and is particularly important because an end user who does not understand or believe the output of an analytic will not use it, regardless of how accurate or performant its underlying computation.

CURRENT APPROACHES

In reliability engineering and condition-based maintenance (CBM+), expert models are useful for building digital twins that can provide highly accurate and faithful predictions of real-world systems and are useful for systems with known dynamics. We have seen examples of this when building physics-based models for assessing structural damages of generic structures to predicting the remaining life of lithium-ion batteries.[1,2] However, these approaches are often brittle and difficult to maintain as operating conditions and environmental dynamics keep changing. For instance, a physics model of an RCV with a combination of lidar and radar along with an aerial drone platform will differ wildly from a physics model of an RCV that has an almost identical suite of components but swapping out the drone platform for a 9 mm turret.

Data-driven algorithms based on techniques in artificial intelligence and machine learning (AI/ML), specifically in deep learning and reinforcement learning, are used for predictive maintenance. However, these methods require huge amounts of training data to learn the performance models of various components on board uncrewed platforms. Also, these data-driven models are often seen as "black boxes" and lack explainability, making it difficult for the end user to build trust or explain the outcome of these AI/ML models. A 2021 survey by Theissler et al. showed that most papers published in areas of predictive maintenance rely on deep learning approaches centered around a labeled dataset.[3] We have seen work from Li et al. that shows promise in detecting the remaining useful life for vehicle's power system and work from Manoharan et al. that yields similar results.[4,5] However, these methods do not work in the context of assessing vehicle control stability, as neural network models, while highly optimized, are difficult to generalize across domains. Ragone et al. use a physical model that simulates a vehicle powertrain to predict battery state of charge based on an electrochemical–thermal model of lithium-ion batteries.[6] While this method

https://doi.org/10.52202/075106-0009

works well for a well-scoped component like the battery, this approach would become enormously difficult to sustain and generalize across multiple components for a vehicle.

Furthermore, we have seen in the past few years the emergence of powerful large language models (LLMs) that have had success in generative AI.[7] However, these models are enormous in size, requiring 500 pflops/s of computation to train and involving thousands of GPUs. Even after training, these models frequently take up to 200 GB of memory to load, which far exceeds the computational capacity of any RCV platform.

Work that combines expertise and data-driven approaches has attempted to extend physics-based models to inform deep-learning and has seen success in determining the natural degradation of machine components in a vacuum.[8] However, this approach still has limitations. For one, failure modes that do not follow the natural degradation (i.e., sudden perturbations like environmental impact or battle damage) will not be captured. This is because physics-of-failure models cannot capture these sudden dynamics easily, and training data for deep-learning models will usually lack these scenario-specific datapoints. Furthermore, this approach is largely shoehorning a physics approach into a data-driven approach and lacks the compactness, adaptability, and ease of implementation of alternative Hybrid-AI approaches.

HYBRID-AI APPROACH USING PROBABILISTIC PROGRAMMING

In this paper, we describe a more holistic and natural approach to Hybrid-AI for predictive maintenance rooted in probabilistic programming. Probabilistic programming is a computer programming paradigm that raises uncertainty to a first-class programming language construct and automates the processes of parameter learning and statistical inference, including evaluation of complex queries against the model. This enables the modeler to combine domain expertise, in the form of hard-coded prior probability distributions, with data, in which a model can learn summary statistics and parameters directly from a dataset. PP builds on probabilistic graphical models (PGMs), greatly generalizing their capabilities in three ways: (1) Instead of specifying model structure in a domain-specific language, PP models are specified in an ordinary programming language that is augmented with two simple, intuitive additional *sample* and *observe* statements. Writing x = sample(a, p) is equivalent to sampling $x \sim p(x| \ldots)$ and recording the sampled value, along with its probability, in a maplike data structure at the address a, while writing observe(a, p, v) is equivalent to evaluating $p(x = v| \ldots)$ and similarly recording the observed value and its probability at a. (2) Instead of enforcing a static model structure, probabilistic programs can encode arbitrarily complex model structures, including those with stochastic choice (in which future model state depends on the value of random choices) and open-universe structure (in which the model may include arbitrarily many random variables).

This interplay between incorporating domain knowledge with data is crucial because it provides probabilistic programs with (1) *adaptability*¸ as probabilistic programs are naturally modular and composable, meaning that with minimal sets of changes, they can be repurposed to fit new operating contexts of an uncrewed platform; (2) *robustness*, as probabilistic programs can rely on prior domain knowledge in the absence of data and also encode that knowledge as probability distributions with confidence bounds, enabling reasoning under uncertainty; (3) *compactness*, as these models do not require detailed understanding like a rich physics simulator, nor terabytes of data to train; and finally (4) *explainability*, in which probabilistic programs are easily traceable, and an end user can gain an understanding into the logical underpinnings of a probabilistic model.

 https://doi.org/10.52202/075106-0009

CASE STUDY EXPERIMENTS AND RESULTS

Robotic Combat Vehicle

As part of our latest work with the Army Applications Lab, we conducted a series of experiments on a simulated dataset related to an RCV. This builds upon our previous work on this program, where we demonstrated diagnostics and health monitoring on a surrogate vehicle using probabilistic programming.[9] In this case study, we demonstrated that we could predict future health and status of a vehicle at least 100 minutes into the future with error bounds that are reasonably accurate so that the predictions are useful for the vehicle operator and decision-makers higher up the chain of command. In this experiment, we worked with the Army as well as domain experts on the vehicle to create a dataset representing the gradual, and then sudden, breakage, of a bogey-beam that supports the track and suspension system of an RCV, and we observe this gradual decay in the form of a changing angle of the bogey-beam relative to the ground.

For this experiment, we were able to quickly design and implement a probabilistic program representing a dynamic Bayesian network and we used a sequential Monte Carlo (SMC) inference algorithm to predicting with this model. In this model we encoded domain expertise representing the relationships between the observe bogey-beam angle and the health and status of the bogey-beam, along with a historical dataset of angles representing nominal operating conditions of the bogey-beam.

Our results showed that we were able to predict 100 minutes into the future on a non-noisy dataset with an average error of 6.1 minutes, and when predicting up to 30 minutes in advance, we had an average error of 2.7 minutes. Additionally, when we added noise to the data, equivalent to the hard-coded rate of change of the bogey-beam's angle, our error rate went up to 13.1 minutes when predicting up to 100 minutes, and 6.0 minutes when predicting up to 30 minutes. We visualize our prediction results in Figures 1 and 2.

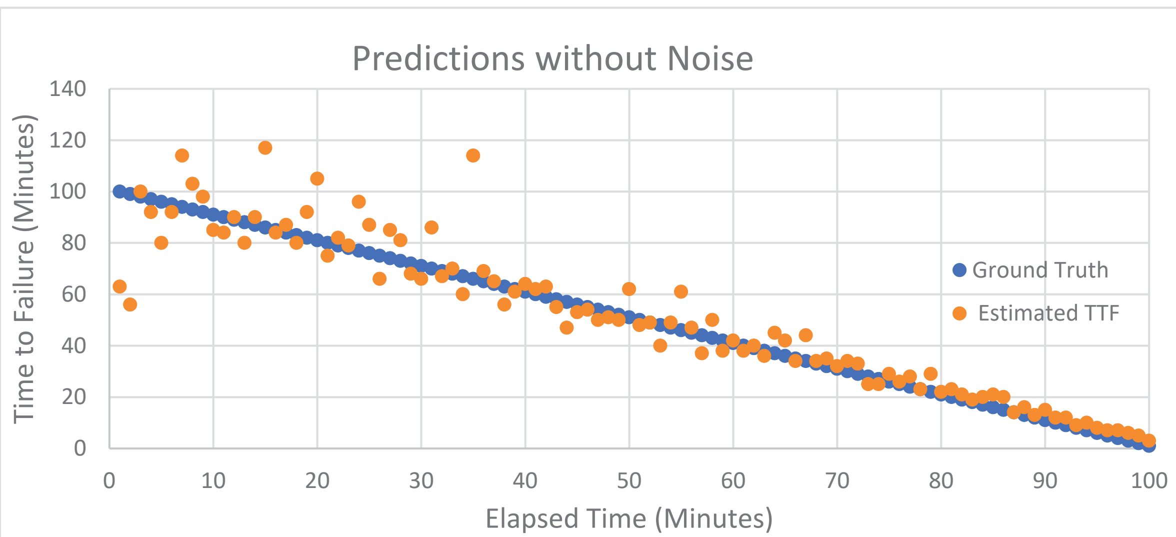

Figure 1. Results of Predictions without Noise. The blue line indicates the ground-truth remaining life of the bogey-beam, which degrades in perfect linear fashion as time elapses. Orange dots indicate the predicted remaining useful life of the bogey-beam at each time step.

https://doi.org/10.52202/075106-0009

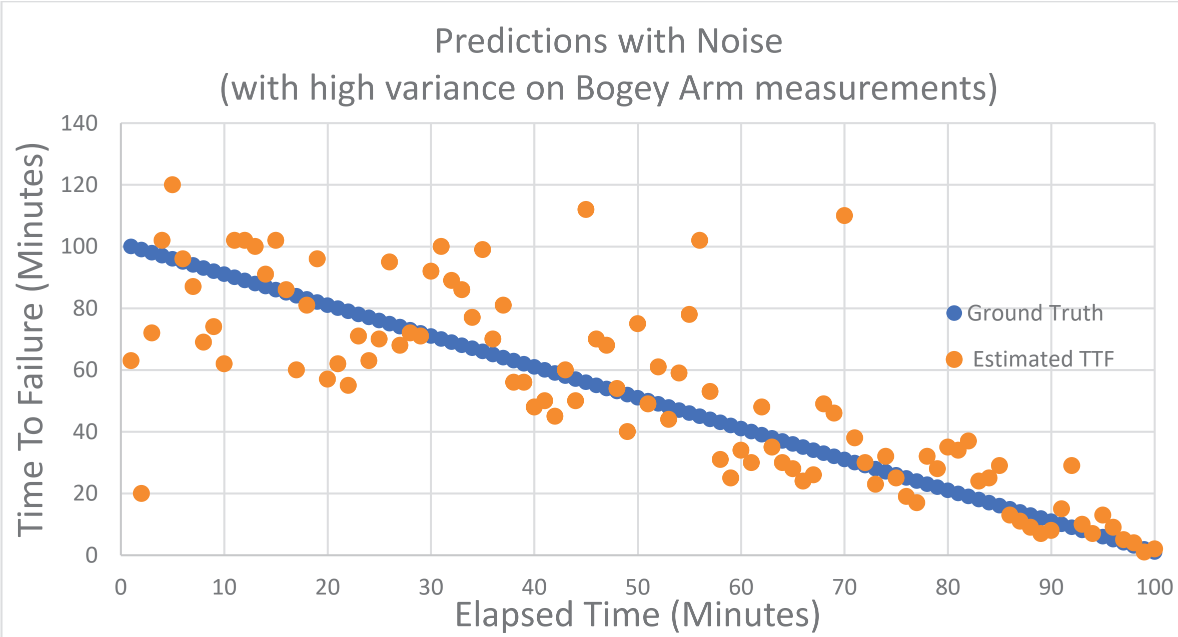

Figure 2. Results of Predictions with Noise. Compared with Figure 1, the results are more erratic initially, but within a 30-minute prediction window, we see that the error rate starts converging to be similar to that of the noiseless scenario. Also, the noise induced in this scenario was deemed to be much higher than that of a real-world sensor.

In terms of compactness and performance, this prognostics model takes up less than 1 MB of memory. Moreover, we have demonstrated previously that over 100 probabilistic models can run on less than 2 GB of memory on a Jetson Nano hardware system. Furthermore, this approach shows how a probabilistic model can compactify streams of sensor data into a single prediction of health and status, making its output (one number) suitable for transmission in a low-bandwidth signal environment.

CONCLUSION

In this work, we demonstrated the power of PP-based Hybrid-AI as a tool for predictive maintenance at the edge for uncrewed systems. We extended domain expertise to relate the observed angle of a bogey-beam for a track and suspension system to the likely health and status state of the vehicle. We also leveraged historical data to learn summary statistics representing the nominal operating conditions of the bogey-beam. This approach has been demonstrated on both a simulated dataset as well as real-world data, to show both health monitoring and predictive analytics, and can be easily extended to an RCV thereby providing Army Warfighters with a diagnostics tool that enables operational decision-making. Specifically, by notifying the RCV operator of a future failure in the bogey-beam, the operator can adjust their driving to compensate, by either reducing their speed in rough terrain, or by changing their steering to compensate for a steering angle bias. They can also inform a higher-level Commander on whether to bring the RCV back to the motor pool for repair. This in turn not only reduces the downtime for RCV repair, but also improves the ability to support longer range missions, as operators can drive RCV further beyond the forward line of troops, knowing the Hybrid-AI prediction capabilities can accurately gauge vehicle health and status. At the same time, commanders at higher level echelons can plan out tactical engagements and missions where RCVs are deployed across a wide geographical area, knowing that they will have awareness and insights into future fleet level availability of capabilities

https://doi.org/10.52202/075106-0009

REFERENCES

[1] M. Montaru, S. Fiette, J.L. Koné, and Y. Bultel, "Calendar ageing model of Li-ion battery combining physics-based and empirical approaches." *Journal of Energy Storage*. Vol. 51, 2022, 104544. https://doi.org/10.1016/j.est.2022.104544

[2] T.G. Ritto, and F.A. Rochinha, "Digital twin, physics-based model, and machine learning applied to damage detection in structures." *Mechanical Systems and Signal Processing*. Vol. 155, 2021, 107614. https://doi.org/10.1016/j.ymssp.2021.107614

[3] A. Theissler, J. Pérez-Velázquez, M. Kettelgerdes, and G. Elger, "Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry." *Reliability Engineering & System Safety*. Vol. 215, 2021, 107864. https://doi.org/10.1016/j.ress.2021.107864

[4] P. Li, Z. Zhang, R. Grosu, et al. "An end-to-end neural network framework for state-of-health estimation and remaining useful life prediction of electric vehicle lithium batteries." *Renewable and Sustainable Energy Reviews*. Vol. 156, 2022, 111843. https://doi.org/10.1016/j.rser.2021.111843

[5] A. Manoharan, K.M. Begam, V.R. Aparow, and D. Sooriamoorthy, "Artificial neural networks, gradient boosting and support vector machines for electric vehicle battery state estimation: A review." *Journal of Energy Storage*. 55, 2022, 105384. https://doi.org/10.1016/j.est.2022.105384

[6] M. Ragone, V. Yurkiv, A. Ramasubramanian, B. Kashir, and F. Mashayek, "Data driven estimation of electric vehicle battery state-of-charge informed by automotive simulations and multi-physics modeling." *Journal of Power Sources*. Vol. 483, 2021, 229108. https://doi.org/10.1016/j.jpowsour.2020.229108

[7] D. Narayanan, M. Shoeybi, J. Casper, et al. "Efficient large-scale language model training on GPU clusters using Megatron-LM." In: *SC21: International Conference for High Performance Computing, Networking, Storage and Analysis*. 2021, pp. 1–14. https://ieeexplore.ieee.org/document/9910057

[8] M.A. Chao, C. Kulkarni, K. Goebel, and O. Fink, "Fusing physics-based and deep learning models for prognostics." *Reliability Engineering & System Safety*. Vol. 217, 2022, 107961. https://doi.org/10.1016/j.ress.2021.107961

[9] K. Lu, S. Cvijic, D. Dewhurst, J. Gorman, R. Hyland and J. Templin, "Democratizing AI for condition-based maintenance leveraging probabilistic programming for symbolic reasoning. In: *2023 Annual Reliability and Maintainability Symposium (RAMS)*. 2023, 1–6. dhttps://doi.org/10.1109/RAMS51473.2023.10088272