Predictive Digital Twins as a Foundation for Improved Mission Readiness

Michael Kapteyn
Massachusetts Institute of Technology, Cambridge, MA 02139

Karen E. Willcox
University of Texas at Austin, Austin, TX 78712
UNITED STATES OF AMERICA

mkapteyn@mit.edu kwillcox@oden.utexas.edu

Keywords: Intelligent Solutions for Improved Mission Readiness, Digital Twin, Unmanned Aerial Vehicle, Dynamic Data-Driven Decisions.

ABSTRACT

A digital twin is an evolving virtual model of a specific system or physical asset, assimilating asset lifecycle data so that the digital twin becomes a dynamically updated asset-specific model that underpins intelligent automation and drives key decisions. Digital twins have potential impact across critical areas of national security, industrial development, and societal well-being. If made reliably predictive, digital twins could revolutionize key decision-making processes that depend on dynamically evolving estimates of the state of a complex system. This paper illustrates how a predictive digital twin – one that combines data-driven learning with predictive physics-based modelling – can contribute to improved mission readiness. The digital twin is represented mathematically as a probabilistic graphical model in which the key elements of state, control, observations, quantities of interest, and reward are modelled as random variables. The graphical model represents the relationships between these different elements, as well as their evolution in time and their uncertainties. The formulation is illustrated for the development of a structural digital twin for an unmanned aerial vehicle (UAV). The digital twin combines high-fidelity structural finite element models, computationally efficient reduced-order models, and observational data generated from onboard structural sensors. An illustrative example shows how the digital twin is updated as the UAV undergoes in-flight structural degradation and then used to optimally re-plan the mission trajectory.

1.0 INTRODUCTION

A digital twin is an evolving virtual model that mirrors an individual physical asset throughout its lifecycle (Grieves and Vickers, 2018; AIAA, 2020). An asset-specific model is a powerful tool to underpin intelligent automation and drive the key decision-making processes that contribute toward improved mission readiness of UxVs. For example, digital twins have the potential to facilitate predictive maintenance, optimized vehicle utilization and fleet management, and end-to-end integration of information throughout the asset lifecycle (Tuegel et al., 2011; Glaessgen and Stargel, 2012). In this work, we focus on the application of digital twins to enable dynamic mission reconfiguration in response to changes in the state of the UxV.

Key to the digital twin concept is the ability to sense, collect, analyze, and learn from the asset's data. Of particular importance is a synergistic multi-way coupling between the physical asset, data collection, computational models, and the decision-making process. This synergistic coupling has been explored extensively under the dynamic data-driven application systems (DDDAS) paradigm (Blasch et al., 2018; Darema, 2004). The tasks of data-driven modeling, inference, assimilation, prediction, control, and planning are all a part of enabling the digital twin paradigm. These tasks all can benefit from the formulations and methods of machine learning. However, when it comes to engineering assets such as UxVs, which serve high-consequence critical societal functions, it is not enough to rely on data-driven machine learning alone. The digital twins of these UxVs will be used to drive decisions that by their very nature are based on predictions that go beyond the available data.

The digital twins of these UxVs must account for their complex multi-scale multi-physics dynamics, must represent high-dimensional uncertain parameters that cannot be observed directly, and must accurately characterize rare events. These predictive digital twins can only be achieved through a synergistic combination of predictive physics-based modeling, data-driven machine learning, and uncertainty quantification. This paper summarizes an approach to create, update, and deploy data-driven physics-based digital twins. We demonstrate the approach through the development of a structural digital twin for an unmanned aerial vehicle (UAV).

2.0 PREDICTIVE DIGITAL TWIN

Our recent work (Kapteyn et al., 2021) has developed a probabilistic graphical model formulation to formalize the definition of a digital twin. The elements of the formulation are as follows:

- Physical state, S: the parameterized state of the physical asset (e.g., vehicle geometry, structural health, etc.)
- Digital state, D: the parameters (i.e., model inputs) that define the computational models comprising the digital twin (e.g., geometry, structural parameters, boundary conditions)
- Control inputs, U: actions or decisions that influence the physical asset (e.g., in-flight maneuvers, maintenance decisions, inspection decisions, sensor installation)
- Observational data, O: available information describing the state of the physical asset (e.g., measured strain or accelerometer data, flight logs)
- Quantities of interest, Q: quantities describing the asset estimate via model outputs (e.g., stress, strain, displacement fields, failure stress, remaining useful life)
- Reward, R: quantifies the overall performance of the asset-twin system (e.g., mission success, fuel burn, maintenance costs, inspection costs)

To build the digital twin, we define each of these elements and represent them using random variables. The digital twin mathematical model is then realized as a dynamic Bayesian network that relates the various quantities at different timesteps through conditional probabilities. This network then becomes a foundation on which we can conduct the tasks of asset monitoring (via Bayesian inference), digital twin updating (via data assimilation), prediction, uncertainty quantification, and control. See Kapteyn et al. (2021) for details.

At the heart of the predictive digital twin are physics-based models of the vehicle. In contrast with purely data-driven models, physics-based models offer a greater degree of interpretability, reliability, and predictive capability. The physics-based models play a key role in defining the digital state, $\boldsymbol{D}$, and they will be used to compute the quantities of interest (e.g., characteristics of the structural response or measures of flight capability) with quantified uncertainty. For example, in building a structural digital twin of a UAV, we employ a high-fidelity finite element model. The digital state then comprises parameters defining the geometry, materials properties, and mass properties of the finite element model.

While physics-based models are already ubiquitous in engineering, accurately modeling a complete UAV system requires complex computational models that entail significant computational resources to solve. Computational resource challenges are heightened when digital twins are required to provide near real-time insights in order for them to be used effectively for operational decision making. This requires the ability to rapidly adapt the digital twin in the face of changing model parameters, and rapidly evaluate the underlying models to provide analysis and prediction. Traditional large-scale physics-based models (e.g., finite element models) are typically too expensive to solve in this type of real-time, many-query context (Hartmann et al., 2018). To overcome this challenge, we build a library of component-based reduced-order models derived from the high-fidelity finite element model of the vehicle (Kapteyn et al., 2020a). The component-based approach is key to achieving scalability to high-dimensional parameters. The use of mathematically rigorous projection-based model reduction retains the physics-based grounding of the model, while enabling rapid model evaluations suitable for the real-time or many-query context.

Figure 1 illustrates how this approach is implemented. In an offline phase, we construct the library of component-based models, each representing different asset states (e.g., different states of vehicle structural health). In the offline phase we also use machine learning to build a computationally efficient map from expected observations to vehicle state. This map may be manifested computationally as a classifier (e.g., a decision tree) or a regression model. In the online phase, we use this map to dynamically update the digital twin and use it for dynamic in-flight decision-making. For example, acquired sensor data may be used to rapidly update the structural health parameters of the digital twin models. The updated digital twin is then used to issue predictions, enabling the UAV to replan a safe mission in response to vehicle structural damage.

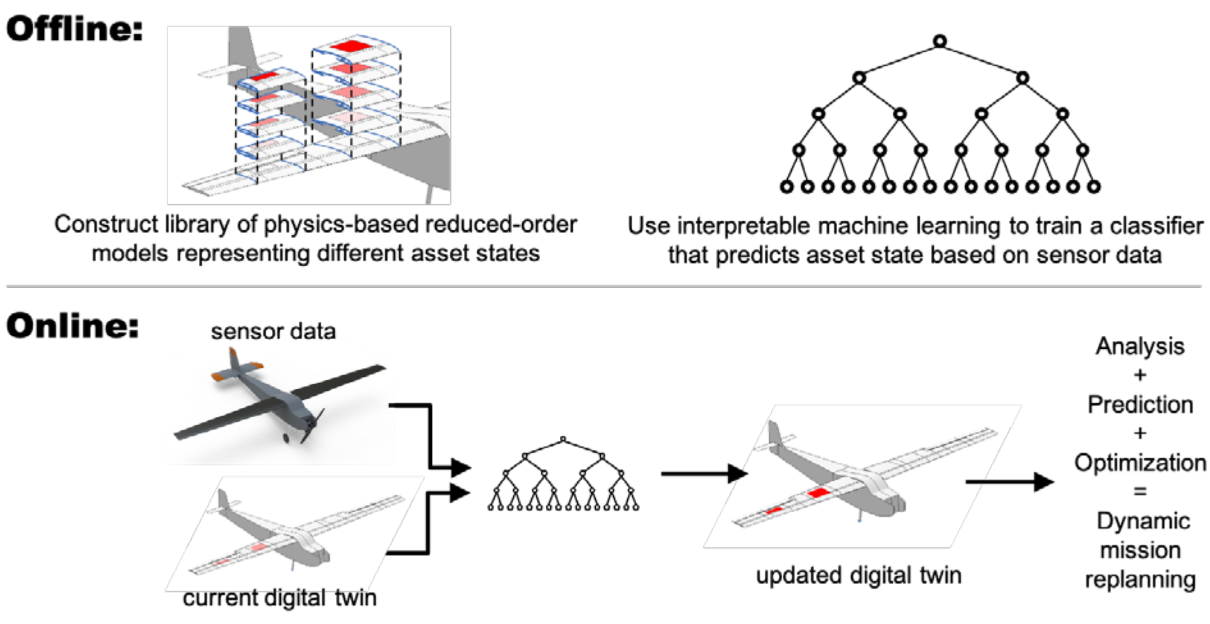

Figure 1: System architecture for a structural digital twin of a UAV. The digital twin combines physics-based models with interpretable machine learning to enable dynamic in-flight mission replanning.

3.0 DEMONSTRATION: DYNAMIC DATA-DRIVEN IN-FLIGHT HEALTH MONITORING AND MISSION ADAPTATION

In this section we demonstrate an operational phase in which the calibrated digital twin is deployed alongside the UAV. We formulate the problem of in-flight structural health monitoring and self-aware dynamic mission adaption. We then combine the component-based UAV model library within the probabilistic graphical model foundation in order to demonstrate the self-aware UAV capability.

3.1 UAV Asset and Digital Twin

Figure 2 shows the 12ft wingspan UAV asset developed as part of this research. More details on the experimental testbed are given in Salinger et al. (2020). A component-based reduced-order model of this UAV asset has been developed, and key model parameters have been experimentally calibrated to match the as-manufactured hardware. Figure 3 depicts a solution to the physics-based structural model comprising the digital twin, where the applied force corresponds to a 3g aerodynamic load.

Figure 2: Physical UAV asset developed for this work.

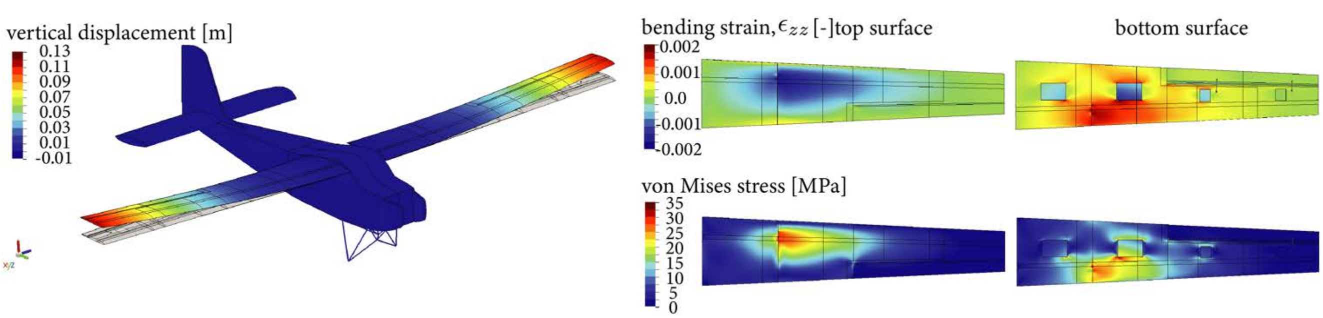

Figure 3: Static elasticity solution computed using UAV finite element model. Top: Displacement field of the full UAV under a 3g aerodynamic load. Middle: Bending strain field. Bottom: von Mises stress field.

3.2 Mathematical Formulation via Probabilistic Graphical Model

We consider an illustrative UAV mission in which the digital twin assimilates structural sensor data and dynamically estimates the structural health of the UAV. The estimated structural health is used to inform control inputs, which in this case determine how aggressively the UAV performs its mission.

Using onboard structural sensors, we dynamically estimate the structural health parameters, as the health of the UAV evolves over time. For illustrative purposes, we consider here a component library that includes two structural defect regions and 13 pristine components. The defect regions are located on the top surface of the second and fourth spanwise component on the right wing. We then define two structural health parameters within the digital state, z_1 and z_2, which represent the percentage reduction in material stiffness applied to each of the defect regions. Note that this defines only part of the full digital state space, as there can also be variability in the other digital state parameters (such as the parameters defining the UAV geometry and material properties).

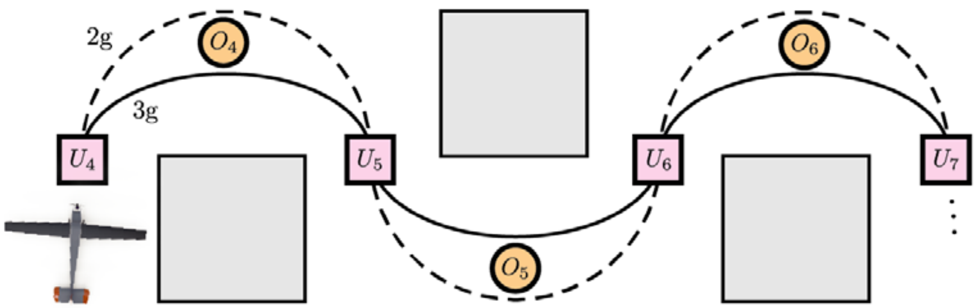

Figure 4: A schematic of the considered illustrative mission.

The illustrative UAV mission considered consists of successive level turns, as shown in Figure 3. In this phase, the timesteps $t = 4,5,6, \ldots$ with control inputs $U_4, U_5, U_6, \ldots$ and observational sensor data $O_4, O_5, O_6, \ldots$ depicted in Figure 4 correspond to successive turns executed during the UAV's mission. In this illustration, the UAV undergoes structural degradation throughout its mission. The digital twin responds adaptively to the evolving structural health by determining optimal updated maneuvers for the UAV. In particular, at each timestep t the digital twin issues a control input, $u_t \in \{2\text{g}, 3\text{g}\}$, which instructs the UAV to take the next turn at a bank angle corresponding to an aerodynamic load factor (the ratio of lift to weight) of either 2g or 3g. Taking a turn at a steeper bank angle makes the path shorter, but also subjects the UAV to an increased aerodynamic load, which has a greater chance of worsening the UAV structural health.

The digital twin is updated using dynamic estimates of the structural health of the UAV. This dynamic updating is achieved by using the digital twin's calibrated internal models (Figure 3) to assimilate observational data and adjust its predictions accordingly. The observational data at each timestep are noisy strain measurements, from each of 24 uniaxial strain gauges on the upper surface of the wing (positioned near the defect regions, as shown in Figure 2). Observational data are assumed to be acquired during a quasi-static section of each turn. This allows us to simplify the problem by ignoring the transient structural loading between successive turns.

3.3 Planning and Optimal Control for Dynamic Mission Reconfiguration

During flight, the digital twin uses its estimate of the current structural health parameters to select the appropriate reduced-order structural model from the library, which can then be evaluated to provide a deeper analysis of the UAV structural integrity and consequent flight capability. For this example, we define the quantities of interest to be computational estimates of the strain at strain gauge locations. This quantity is chosen to enable a posterior predictive check: the digital twin compares its posterior estimate of the strain with the observed strain in order to evaluate how well its models match reality. In practice this type of check can help validate other predictions made by the digital twin structural models, such as modal quantities or the full stress and strain fields. Here, we quantify the posterior predictive error via a reward function, R^{error}, that measures the difference between observed strain measurements and strains predicted by the digital twin, normalized by an estimate of the sensor standard deviation. We define two additional reward functions targeted at different aspects of the mission. The second reward function, R^{health}, measures how far the UAV is from structural failure, as defined by a maximum allowable strain level. This term rewards the UAV for remaining in good structural health, as indicated by low predicted strain. The third reward function, $R^{control}$, is assigned to each applied control input. In this illustration the faster 3g turn is assigned a higher reward, indicating a preference for the more aggressive flight path.

The UAV digital twin determines which health-aware control inputs to issue by solving a planning problem induced by the graphical model structure. In this demonstration we conduct offline planning: prior to the mission, the digital twin uses its internal models to predict how the structural health will evolve over the mission and uses this predictive capability to compute a health-dependent control policy. During planning, the digital twin seeks a control policy that optimizes a weighted sum of the reward functions over the mission, namely,

$$R^{policy} = R^{health} + \eta \, R^{control}$$

where η is a tradeoff parameter between UAV aggression and self-preservation. In this way, the optimal control strategy recommended by the digital twin balances structural preservation with mission aggressiveness, while enabling dynamic mission replanning in response to realized in-flight damage or degradation.

We make the planning problem tractable by approximating it as a fully-observable Markov decision process (MDP). This can be viewed as a way of decoupling the sensing and control problems: we design a controller assuming that in-flight strain measurements will provide an accurate and certain estimate of the UAV structural state. Such a simplification may be applicable to many engineering systems, provided they are equipped with sufficient sensing capability. We also adopt an infinite planning horizon (effectively assuming that in practice the UAV would be flying successive missions indefinitely). The resulting policy is suboptimal in general, but performs well if the system is indeed well observed (i.e., the point estimates are accurate). However, one noteworthy limitation of adopting this assumption is that the asset will never perform actions that are purely for the purpose of information-gathering or improving observability over the state. This simplified MDP planning problem is solved offline using the classical value iteration algorithm (Russell and Norvig, 2002). The control input U_t at time t is then defined by the policy, conditioned on the current estimated state of the digital twin.

3.4 Results

Figure 5 illustrates the control policies resulting from different values of η, the tradeoff parameter between UAV aggression and self-preservation within the planning reward function. As η increases, the reward function favors mission aggressiveness over structural preservation. Figure 5 illustrates the computed control policies as a function of the structural health parameters z_1, z_2, where the policy defines the decision boundary between the 2g and the 3g maneuver. The policy recommends that the UAV fly the more aggressive 3g maneuver until the maximum a posteriori estimate of the structural health parameters progresses beyond a certain decision boundary. Figure 5 shows that, as expected, the decision boundary becomes more aggressive as the value of η is increased.

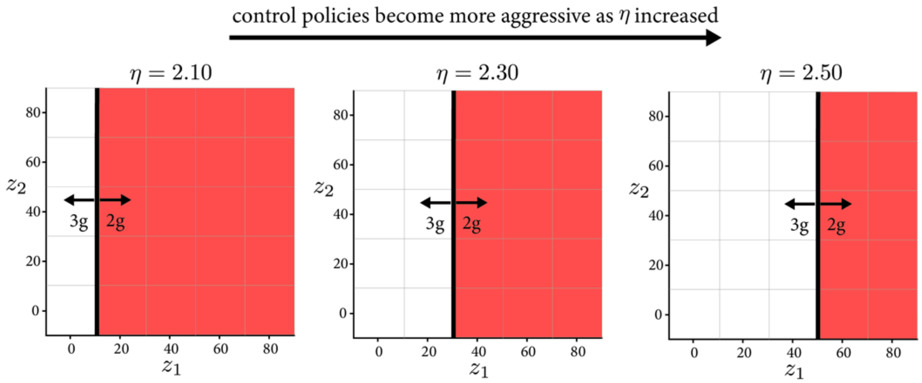

Figure 5: UAV control policies computed for the simulated mission, under various values of the tradeoff parameter η

Figure 5 also illustrates that in this example the control policies depend predominantly on z_1, the structural health parameter of the inboard panel highlighted in Figure 1. In fact, for $\eta = 2.10, 2.30$ and 2.50 the control policies shown are independent of z_2. This result highlights the benefit of utilizing the digital twin physics-based structural models within the planning process. The structural analysis afforded by the digital twin models reveals that z_1 has a greater influence than z_1 on the structural integrity of the UAV, and this is naturally reflected in the computed control policies.

Figure 6 depicts three snapshots in time for a simulation of the illustrative UAV mission we consider. Here we use a classification tree to estimate the structural health parameters, z_1 and z_2, based on incoming strain measurements (Kapteyn et al., 2020b). The UAV uses these estimates to update its internal structural model of the airframe, which can then be evaluated for rapid, up-to-date, structural analysis. In this example the UAV responds to worsening structural health by employing its control policy (in this case with $\eta = 2.30$), to switch from the more aggressive flight path (3g maneuvers) to the more conservative flight path (2g maneuvers). This illustrates how the structural digital twin enables dynamic mission reconfiguration through data assimilation, prediction, and principled decision making.

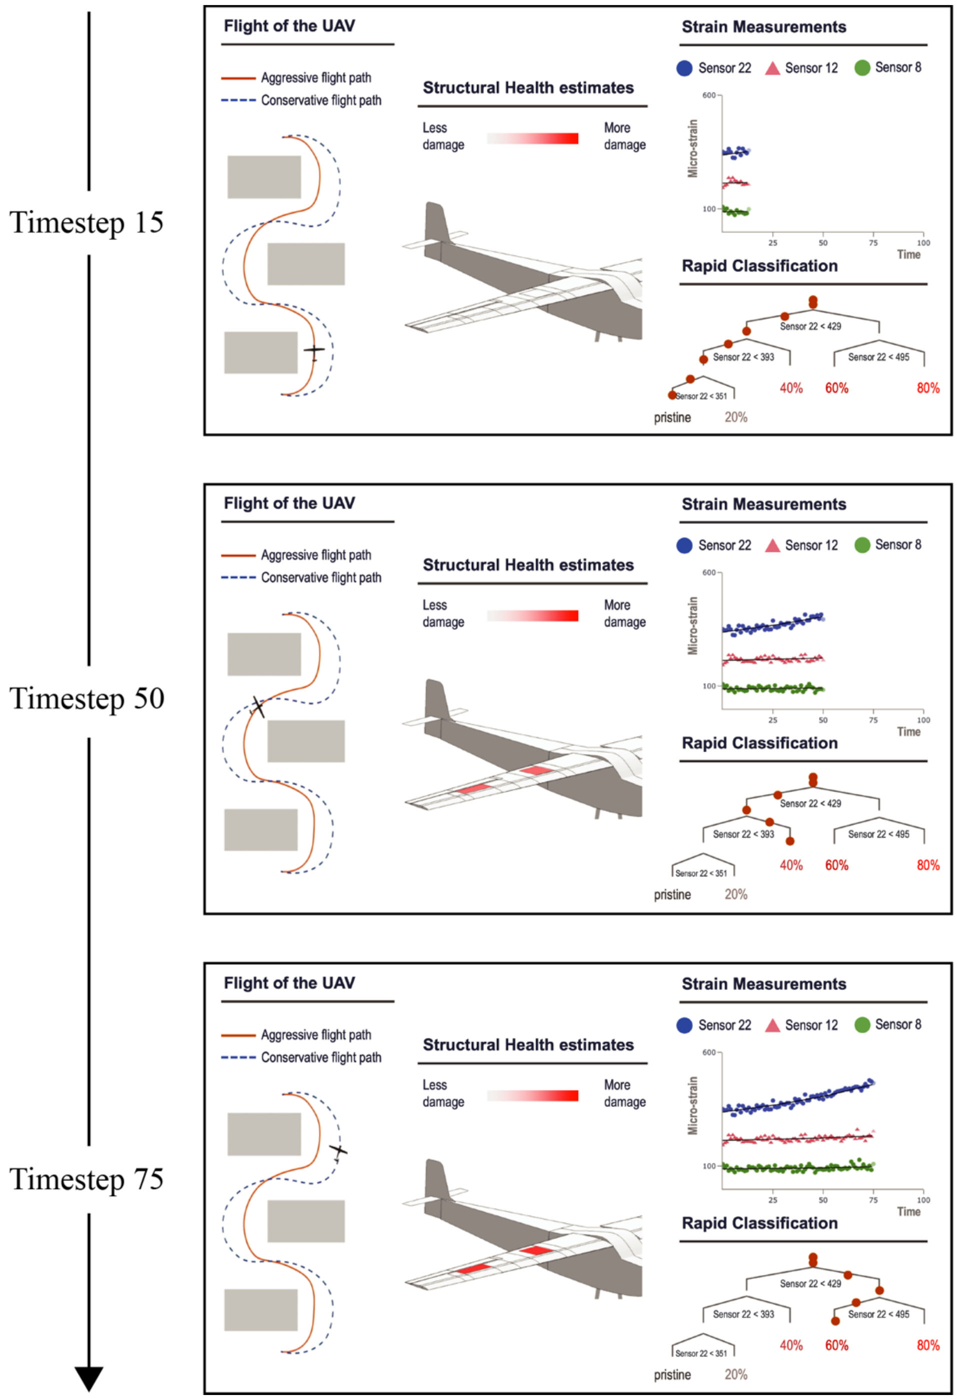

Figure 6: Snapshots of the simulated UAV mission. The digital twin acquires strain measurements from wing-mounted strain gauges, and assimilates these data using a rapid and interpretable classification tree. The UAV employs the precomputed control policy to respond dynamically to worsening structural health estimates by switching from the aggressive flight path to the conservative flight path.

4.0 CONCLUSION

Digital twins are being developed and deployed across a broad range of industries and disciplines. Early successes, particularly in structural health monitoring, point to the high-value potential of digital twins in achieving improved mission readiness. However, a number of open challenges remain in order for digital twins to achieve the levels of reliability and robustness needed to support mission-critical decisions. One key challenge is comprehensive quantification of uncertainty. This paper utilizes a Bayesian formulation of the digital twin that explicitly accounts for uncertainty, including time-dependent system uncertainty. A second key challenge is computational tractability, especially in real-time and resource-constrained settings. This paper has employed reduced order modeling to reduce computational cost, but more work is needed to create certified surrogate models that enable rapid model updating and tight feedback loops between models and data. A third challenge is addressing the interaction between the digital twin and a human decision-maker, including ensuring the interpretability of digital twin recommendations and information flows.

REFERENCES

[1] AIAA Digital Engineering Integration Committee, 2020. Digital Twin: Definition & Value. An AIAA and AIA position paper, American Institute of Aeronautics and Astronautics (AIAA) and Aerospace Industries Association (AIA).

[2] Blasch, E., Ravela, S. and Aved, A. eds., 2018. Handbook of dynamic data driven applications systems. Springer.

[3] Darema, F., 2004, June. Dynamic data driven applications systems: A new paradigm for application simulations and measurements. In International Conference on Computational Science (pp. 662-669). Springer, Berlin, Heidelberg.

[4] Glaessgen, E. and Stargel, D., 2012, April. The digital twin paradigm for future NASA and US Air Force vehicles. In 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, Honolulu, Hawaii.

[5] Grieves, M. and Vickers, J., 2017. Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In Transdisciplinary perspectives on complex systems (pp. 85-113). Springer, Cham.

[6] Hartmann, D., Herz, M. and Wever, U., 2018. Model order reduction a key technology for digital twins. In Reduced-order modeling (ROM) for simulation and optimization (pp. 167-179). Springer, Cham.

[7] Kapteyn, M., Knezevic, D., Huynh, D.B.P., Tran, Minh and Willcox, K., 2020a. Data-driven physics-based digital twins via a library of component-based reduced-order models. International Journal for Numerical Methods in Engineering, https://doi.org/10.1002/nme.6423.

[8] Kapteyn, M., Knezevic, D., and Willcox, K., 2020b. Toward predictive digital twins via component-based reduced-order models and interpretable machine learning. In Proceedings of 2020 AIAA SciTech Forum & Exhibition, Orlando FL.

[9] Kapteyn, M., Pretorius, J. and Willcox, K., 2021. A Probabilistic Graphical Model Foundation for Enabling Predictive Digital Twins at Scale. Nature Computational Science, Vol 1.

[10] Russell, S. and Norvig, P., 2002. Artificial Intelligence: A Modern Approach. Prentice Hall.

[11] Salinger S., Kapteyn M., Kays C., Pretorius J., and Willcox K., 2020. A Hardware Testbed for Dynamic Data-Driven Aerospace Digital Twins. In: Darema F., Blasch E., Ravela S., Aved A. (eds) Dynamic Data Driven Application Systems. DDDAS 2020. Lecture Notes in Computer Science, Vol 12312. Springer, Cham.

[12] Tuegel, E.J., Ingraffea, A.R., Eason, T.G. and Spottswood, S.M., 2011. Reengineering aircraft structural life prediction using a digital twin. International Journal of Aerospace Engineering, 2011.