Session 2

Digital Twins

S2-A1	Sajad Salavatidezfouli — Data-driven generative digital twins for wind-farm flows from highly sparse measurements
S2-A2	Sebastiano Randino — A Self-Calibrating Digital Twin for Rotor-Effective Wind-Speed Estimation and Aerodynamic-Map Identification
S2-A3	Yves-Marie Lefebvre (Philipp Bekemeyer) — Bridging the Gap between Numerical and Experimental Analysis through Machine Learning Techniques Employed in Real-Time
S2-A5	Yannick Lecomte — Bayesian Reinforcement Twinning: A Multi-Fidelity Framework for Reciprocal Learning Between Digital Twins and Control
S2-P4	Vincent Seyfried — From Neural Surrogates to Neural Twins

Data-driven generative digital twins for wind-farm flows from highly sparse measurements

[1]Sajad Salavatidezfouli; [2]Christoffer Hansen; [2]Henrik Karstoft; [3]Alexandros Iosifidis; [2]Mahdi Abkar*

[1]*Department of Electrical and Computer Engineering, Aarhus University, Aarhus, Denmark*

[2]*Department of Mechanical and Production Engineering, Aarhus University, Aarhus, Denmark*

[3]*Data Science Research Centre, Tampere University, Tampere, Finland*

Abstract

Reconstructing high-fidelity flow fields from sparse measurements remains a central challenge in applied sensing, monitoring, and real-time flow control. This work addresses the problem in a real-world-scale setting, focusing on highly advective, turbulent wind-farm flows observed through extremely sparse measurements consistent with practical sensing limitations. To span a representative range of reconstruction paradigms, we compare three complementary approaches: gappy proper orthogonal decomposition (GPOD), a linear reduced-order baseline with negligible computational overhead; the shallow recurrent decoder (SHRED), a state-of-the-art deep learning method for sparse flow reconstruction; and a proposed observation-guided generative framework (GGenAI). The methods are systematically evaluated on a fully developed wind-farm dataset across sensor densities ranging from 0.12% to 6.65% of the spatial domain. The results reveal a clear performance hierarchy governed by sensing coverage and model expressivity. In the extremely sparse regime (<0.49% coverage), GPOD yields relatively low global reconstruction error; however, qualitative analysis indicates that this performance primarily reflects recovery of a mean-flow-dominated field rather than faithful instantaneous dynamics. In this same low-coverage regime, SHRED exhibits behavior similar to GPOD, producing largely time-averaged reconstructions. As sensor coverage increases to intermediate levels (around 0.49%), SHRED begins to recover meaningful coarse flow structures. However, even at this coverage, the reconstructions remain overly smooth and fail to capture instantaneous turbulent features. In contrast, GGenAI undergoes an earlier transition toward instantaneous-like reconstruction, and at higher sensor densities (beyond approximately 0.49% coverage) it consistently outperforms both baselines, achieving lower reconstruction errors and substantially higher perceptual fidelity while recovering multiscale wake structures characteristic of instantaneous flow dynamics. To further assess physical consistency, we introduce a POD-based reduced-order reference field representing the dominant energy-containing flow structures. GGenAI exhibits improved agreement with this reduced-order reference, emphasizing that reconstruction quality must be interpreted in relation to the spatial scales that can be reliably inferred from sparse measurements. Overall, the study demonstrates that observation-conditioned diffusion models can be extended to complex, advection-dominated turbulent flows, offering a promising pathway toward data-driven digital twins.

Keywords: Generative AI, Digital Twin, Data-driven Flow Modeling, Sparse Measurement, Wind-farm Wakes

Reinforcement Twinning for Hybrid Control of Floating Wind Turbines in High-Fidelity Simulation

[1,2]Sebastiano Randino*; [1,3]Lorenzo Schena; [2]Emanuele Garone; [1,2,4]Miguel Alfonso Mendez;

[1]*von Karman Institute for Fluid Dynamics, Waterloosesteenweg 72, Sint-Genesius-Rode, Belgium*

[2]*Université Libre de Bruxelles (ULB), Avenue Franklin Roosevelt 50, 1050 Brussels, Belgium*

[3]*Vrije Universiteit Brussel (VUB), Pleinlaan 2,1050 Brussels, Belgium*

[4]*Experimental Aerodynamics and Propulsion Lab, Universidad Carlos III de Madrid, Leganés, 28911, Spain*

Abstract

Floating offshore wind turbines exhibit strongly coupled aero-hydro-servo-elastic dynamics, where platform motion can introduce destabilizing effects such as negative damping. In contrast to bottom-fixed turbines, floating systems involve additional degrees of freedom and coupling mechanisms that fundamentally alter the system dynamics, making control design significantly more challenging. State-of-the-art control strategies for floating wind turbines are often based on baseline controllers developed for bottom-fixed systems. Existing approaches include detuning, platform-motion feedback, and multivariable control laws. However, their effectiveness remains limited by model inaccuracies, operating-point dependence, and the increasing complexity of the coupled dynamics under realistic conditions [1].

In this work, we propose a hybrid model-based/model-free control framework based on the Reinforcement Twinning (RT) framework [2]. The approach combines a reduced-order, physics-based model acting as a digital twin with a learning-based control policy. The digital twin is used to support model-based control design and policy optimization, while the model-free component enables adaptation to complex dynamics. A referee mechanism is introduced to arbitrate between model-based and model-free control actions based on performance and consistency with the high-fidelity environment. The framework is implemented using OpenFAST as a high-fidelity reference environment, acting as a ground-truth simulator against which the reduced-order digital twin and the control strategies are evaluated under realistic aero-hydro-elastic conditions. The proposed approach aims to address the limitations of purely model-based or model-free strategies by leveraging their complementary strengths. Attention is given to model uncertainty, system coupling, and the role of reduced-order models in enabling efficient and robust control design for floating wind turbines.

Keywords:

Floating wind turbines, Reinforcement Twinning, Model-based control, Model-free control, Hybrid control, Digital twin, OpenFAST, Reduced-order modeling, Model-based Reinforcement Learning, Adjoint-based Assimilation

References:

[1] Stockhouse, D., Phadnis, M., Henry, A., Abbas, N., Sinner, M., Pusch, M., & Pao, L. Y. (2023, May). Sink or swim: A tutorial on the control of floating wind turbines. In 2023 American Control Conference (ACC) (pp. 2512-2529). IEEE

[2] Schena, L., Marques, P. A., Poletti, R., Ahizi, S., Van den Berghe, J., & Mendez, M. A. (2024). Reinforcement twinning: From digital twins to model-based reinforcement learning. Journal of Computational Science, 82, 102421.

Bridging the Gap between Numerical and Experimental Analysis through Machine Learning Techniques Employed in Real-Time

[1]Alexander Barklage*; [1]Philipp Bekemeyer; [2]Yves-Marie Lefebvre; [2]John Goetz

[1]German Aerospace Center, Institute of Aerodynamics and Flow Technology, Lilienthalplatz 7, 38108 Braunschweig, Germany

[2]Tecplot, Inc., Bellevue, WA, 98006, United States

Abstract

Wind tunnel testing is an essential element in the aerodynamic assessment of various technical systems as it offers high accuracy at moderate cost compared to full scale testing. However, wind tunnel measurements usually provide information for specific parts of the flow field and may not quantify all quantities of interest. Hence, data fusion methods can be used to fill these gaps using information from simulation data. These data fusion methods are usually applied as a post-processing step. In this work, we propose an infrastructure for online monitoring of distributed quantities and apply it to wind tunnel experiments with a full aircraft configuration under transonic, cryogenic conditions including aero-elastic effects and featuring various challenging aerodynamic conditions. We discuss how data-driven models, leveraging data fusion concepts, can be utilized for this purpose. Specifically, we demonstrate i) the online pressure and displacement field inference via Gappy POD, ii) the online prediction of aerodynamic loads using the inferred pressure field and iii) the online identification of contextual outliers via either a Bayesian technique or through variational autoencoders. The online monitoring aims at complementing the measurement data by providing a deeper insight into the flow field. The inferred distributed fields and scalar quantities are visualized and continuously updated via an automated Python interface, enabling immediate spatial interpretation of flow features and aero-elastic effects during testing. The visualization layer forms an integral part of the monitoring loop, providing near-real-time feedback to support experimental decision-making and rapid identification of anomalies. The presented framework achieves refresh rates, which are suitable for quasi-steady wind tunnel measurements and is capable of detecting outliers during wind-tunnel testing in near-real-time.

Keywords:

Data-Fusion, Outlier Detection, Real-Time ML Operations, Online Visualization, Wind-Tunnel Testing, Aircraft Aerodynamics

Bayesian Reinforcement Twinning: A Multi-Fidelity Framework for Reciprocal Learning Between Digital Twins and Control

[1,2]Yannick Lecomte*; [1,3]Alessandro Grava; [2]Emanuele Garone; [1,3,4]Miguel A. Mendez

[1] *Environmental and Applied Fluid Dynamics, von Karman Institute for Fluid Dynamics, Belgium.*

[2] *Department of Control Engineering and System Analysis, Université Libre de Bruxelles, Belgium.*

[3] *Aero-Thermo-Mechanics Laboratory, Université Libre de Bruxelles, Belgium.*

[4] *Experimental Aerodynamics and Propulsion Lab, Universidad Carlos III de Madrid, Spain.*

Abstract

Unmanned Aerial Vehicles (UAVs) are increasingly employed in applications requiring high levels of autonomy and reliability. However, their operation often involves strongly nonlinear dynamics, actuator couplings, and environmental disturbances that challenge the design of robust control policies. Purely model-based strategies can suffer from modeling inaccuracies, while model-free learning approaches typically require large amounts of experimental data and may explore unsafe or inefficient regions of the control space.

In this work, we propose a framework for **Bayesian Reinforcement Twinning**, a hybrid methodology that combines bidirectional links between digital twin modeling and model-free control. The approach leverages the digital twin, consider as low-fidelity, to guide the exploration of the model-free policy search, consider as high-fidelity, within a multi-fidelity Gaussian framework. This effectively constrains the learning process toward promising regions improving data efficiency while avoiding unsafe exploration. In contrast to standard multi-fidelity approaches, the proposed framework introduces a reciprocal interaction between the physical system and its digital counterpart.

Specifically, observations collected from the real system are used to update the digital twin online. This bidirectional exchange enables continuous refinement of the digital twin across the policy search space, leading to progressively improved predictions and more informed exploration–exploitation trade-offs. A demonstration of the proposed approach is carried out on a dedicated test bench replicating the attitude dynamics of a UAV. This setup features two propellers mounted on the ends of a beam pivoting on a central fulcrum, targeting a set point in terms of attitude, using the motor torque as control actuation and relying only on measurements of motor rpm and attitude angle.

Keywords: Reinforcement Twinning, Bayesian Optimization, Multi-Fidelity Gaussian Process, UAVs, hybrid control.

From Neural Surrogates to Neural Twins

[1,2]Vincent Seyfried*; [3] Andreas Radler; [3,4] Johannes Brandstetter; [1,2,4] Thomas Lichtenegger

[1] *Department of Particulate Flow Modelling, JKU Linz, Austria*

[2] *Christian Doppler Laboratory for Data-Assisted Simulations of Complex Flows, Linz, Austria*

[3] *LIT AI Lab, Institute for Machine Learning, JKU Linz, Austria*

[4] *Emmi AI GmbH, Linz, Austria*

Abstract

Neural surrogates have emerged as powerful reduced-order models for complex dynamical systems, delivering high accuracy and real-time performance. Expanding upon traditional surrogates, we define neural twins as models that ingest real-time measurements at inference time and are capable of reconstructing a high-dimensional state. How to effectively integrate such measurements into neural twins, however, remains an open challenge.

In recent work [1], we identified a critical limitation in standard autoregressive implementations of complex, transient dynamics, which would predict a system state given past predictions as well as sparse real-world measurements. Such models tend to over-rely on the autoregressive state rather than the incoming measurements when using standard conditioning, as successive states are in many cases significantly correlated and measurements offer only sparse information. As errors accumulate over the rollout, the model struggles to remain close to the true system trajectory. An ideal neural twin should recognize autoregressive error accumulation and actively correct itself using real-world measurements.

We propose an alternative approach, termed parallel-in-time neural twins (PAINT), which reconstructs the system state purely based on a temporal series of sparse measurements using a sliding window formulation. In the absence of autoregressive state propagation, this approach avoids error accumulation and exhibits a bounded error during inference. We tested this methodology on a challenging two-dimensional turbulent single jet of varying Reynolds number using a generative transformer-based neural network. Both theoretical analysis as well as empirical results demonstrate that this approach is able to stay close to the true trajectory over time and accurately recover system states from sparse measurements with high fidelity.

Another viable approach leverages physics guidance within the framework of diffusion models. In this setting, a prior model is trained exclusively to learn the underlying dynamics, allowing it to be flexibly adapted to different measurements at test time. As the model is not exposed to any measurements during training, it avoids the aforementioned issue of over-relying on the autoregressive state. As we demonstrate, this allows classical autoregressive models to be utilized for general-purpose applications.

Keywords: digital twin, real-time simulation, dynamic modelling, neural surrogate, turbulence

References

[1] Radler, Andreas, et al. "PAINT: Parallel-in-time Neural Twins for Dynamical System Reconstruction." arXiv preprint arXiv:2510.16004 (2025).