EuroGNC CEAS Navigation Guidance Control
Madrid, Spain May 5^{th}-7^{th} 2026
uc3m Universidad Carlos III de Madrid AIAA

Inferring the Flow: AI-Based State Estimation from Sparse Sensing for Aerodynamic Control

Carlos Sanmiguel Vila
Tenured Scientist, Sub-directorate general of aeronautical systems, Spanish National Institute for Aerospace Technology (INTA). csanvil@inta.es
Adjunct Professor, Department of Aerospace Engineering, Universidad Carlos III de Madrid, Leganés, Spain. csanmigu@ing.uc3m.es

ABSTRACT

The increasing frequency of extreme and clear-air turbulence (CAT) in aviation poses a significant challenge, as many of these events are invisible to current onboard sensing and forecasting tools. Effective mitigation strategies, such as gust load alleviation or buffet suppression, require anticipatory control systems. These systems are critically dependent on robust, real-time estimation of the high-dimensional flow state. This estimation, however, is operationally constrained to using sparse, non-intrusive, wall-mounted sensors. Recent advances in data-driven reduced-order modelling (ROM) offer a path to solving this problem. Latent-space models, such as observable-augmented autoencoders, can create physically interpretable, low-dimensional representations of complex flow phenomena. AI-driven estimators, including recurrent neural networks, can then map sparse sensor histories to these latent states. This presentation will explore how these components—nonlinear ROMs, AI-driven estimators, and interpretable sparse sensing—converge to enable the robust state estimation required for next-generation control systems, including Model Predictive Control and Model-Based Reinforcement Learning.

Keywords: Turbulence prediction, State estimation, Machine learning, Reduced-order modelling, Model-based control

1 Motivation and Background

Episodes of clear-air turbulence (CAT) and short, energetic gusts are a growing operational concern, with airlines already reporting tens of injuries per year and significant maintenance and rerouting costs [1]. Forecasting based on weather products, satellite imagery or terrain-induced models is effective for large-scale, slowly evolving disturbances, but it remains blind to locally generated shear layers and to the fine-scale structures that actually load the wing. For future data-driven controllers to be useful in this context, the aircraft must first know the instantaneous and near-future flow state around the lifting surfaces [2]. The difficulty lies in the fact that aerodynamic flows are high-dimensional and multi-scale, so sampling thousands of sensors at Nyquist with conventional sensors would result in prohibitive data rates. This is exactly the bottleneck identified in recent data-driven flow-control studies: control is feasible, but only if sensing is sparse, interpretable and placed on the body rather than in the flow [3].

2 Wall-Sensor-Informed State Estimation: Opportunities

Despite the high-dimensional and multi-scale nature of turbulent flows, there is a growing consensus that their dynamics are often dominated by recurrent flow patterns, or coherent structures, which evolve on a low-dimensional manifold. This "hope" for a low-dimensional representation is the central opportunity for data-driven control. Machine learning provides a powerful framework for first identifying this manifold and then building a model that can be "queried" in real-time [4, 5]. A primary challenge is the compression of high-dimensional flow-field data into a compact, meaningful representation. Recent studies have demonstrated the power of physics-augmented autoencoders (PA-AE) [6]. By training a deep neural network to compress full-field simulation snapshots while simultaneously reconstructing a key physical variable, such as the lift coefficient or power coefficient, these models can discover a very low-dimensional latent space (e.g., 3D) that is not just a "black box" but is physically interpretable. This representation effectively captures the essential physics, such as shock-wave movement or flow separation [3].

This latent space provides a viable, low-dimensional "state" for a control system. However, in a real-world scenario, we cannot access the full flow field; we are restricted to sparse, non-intrusive, surface-mounted sensors. The next opportunity, therefore, is to create an "estimator" that maps the information from these practical sensors to the discovered latent state. This can be accomplished by training a temporal encoder, often a recurrent neural network like an LSTM, which processes a short history of sensor measurements to infer the current latent state of the flow [7]. Furthermore, to ensure this estimator is efficient for control applications, the number of sensors must be minimized. Instead of relying on intuition, interpretable AI methods like SHAP can rigorously rank sensor importance from a densely-instrumented model [8]. This enables the design of a lightweight "slim" encoder that operates on only a few, physically-critical sensors. This efficient, sparse estimation model opens the path to implementing predictive models of the latent state dynamics, which can incorporate control actions [7]. This provides the necessary components for applying advanced, anticipatory control strategies.

3 Outlook and Research Vision

The application of the AI-driven state estimation techniques to advanced control systems can significantly mitigate the computational and sensing challenges inherent in real-time flow control. By building an estimator upon a sparse sensor set and an interpretable latent model, we can create a low-latency "digital twin" of the real-time flow physics. This component is the critical enabler for practical, high-performance control, reducing the high-dimensional problem to a computationally tractable one. This real-time latent-state predictor is the ideal "plant model" for advanced control strategies, such as Model Predictive Control and Model-Based Reinforcement Learning.

Acknowledgments

Work produced with the support of a 2024 Leonardo Grant for Researchers and Cultural Creators, BBVA Foundation, project PREVENT grant n. LEO24-2-15988. The Foundation takes no responsibility for the opinions, statements and contents of this project, which are entirely the responsibility of its authors.

Declaration of Use of Artificial Intelligence

AI was employed for proofreading and improve writing clarity. The author remains ultimately responsible for the correctness of the content.

References

[1] Mark C Prosser, Paul D Williams, Graeme J Marlton, and R Giles Harrison. Evidence for large increases in clear-air turbulence over the past four decades. *Geophysical Research Letters*, 50(11):e2023GL103814, 2023.

[2] Zhecheng Liu, Diederik Beckers, and Jeff D Eldredge. Model-based reinforcement learning for control of strongly disturbed unsteady aerodynamic flows. *AIAA Journal*, pages 1–21, 2025.

[3] Kai Fukami, Yuta Iwatani, Soju Maejima, Hiroyuki Asada, and Soshi Kawai. Compact representation of transonic airfoil buffet flows with observable-augmented machine learning. *Journal of Fluid Mechanics*, 1021:A39, 2025.

[4] Alberto Solera-Rico, Carlos Sanmiguel Vila, Miguel Gómez-López, Yuning Wang, Abdulrahman Almashjary, Scott TM Dawson, and Ricardo Vinuesa. β-variational autoencoders and transformers for reduced-order modelling of fluid flows. *Nature Communications*, 15(1):1361, 2024.

[5] Koji Fukagata and Kai Fukami. Compressing fluid flows with nonlinear machine learning: mode decomposition, latent modeling, and flow control. *Fluid Dynamics Research*, 2025.

[6] Kai Fukami and Kunihiko Taira. Observable-augmented manifold learning for multi-source turbulent flow data. *Journal of Fluid Mechanics*, 1010:R4, 2025.

[7] Alberto Solera-Rico, Carlos Sanmiguel Vila, and Stefano Discetti. A framework for realisable data-driven active flow control using model predictive control applied to a simplified truck wake. *arXiv preprint arXiv:2510.11600*, 2025.

[8] Abhijeet Vishwasrao, Sai Bharath Chandra Gutha, Andres Cremades, Klas Wijk, Aakash Patil, Catherine Gorle, Beverley J McKeon, Hossein Azizpour, and Ricardo Vinuesa. Diff-sport: Diffusion-based sensor placement optimization and reconstruction of turbulent flows in urban environments. *arXiv preprint arXiv:2506.00214*, 2025.