

Compact representation of transonic airfoil buffet flows with observable-augmented machine learning

Kai Fukami¹ , Yuta Iwatani¹ , Soju Maejima¹ , Hiroyuki Asada¹  and Soshi Kawai¹ 

¹Department of Aerospace Engineering, Graduate School of Engineering, Tohoku University, Sendai 980-8579, Japan

Corresponding author: Kai Fukami, kfukami1@tohoku.ac.jp

(Received 13 March 2025; revised 25 July 2025; accepted 19 September 2025)

Transonic buffet presents time-dependent aerodynamic characteristics associated with shock, turbulent boundary layer and their interactions. Despite strong nonlinearities and a large degree of freedom, there exists a dominant dynamic pattern of a buffet cycle, suggesting the low dimensionality of transonic buffet phenomena. This study seeks a low-dimensional representation of transonic airfoil buffet at a high Reynolds number with machine learning. Wall-modelled large-eddy simulations of flow over the OAT15A supercritical airfoil at two Mach numbers, $M_\infty = 0.715$ and 0.730 , respectively producing non-buffet and buffet conditions, at a chord-based Reynolds number of $Re = 3 \times 10^6$ are performed to generate the present datasets. We find that the low-dimensional nature of transonic airfoil buffet can be extracted as a sole three-dimensional latent representation through lift-augmented autoencoder compression. The current low-order representation not only describes the shock movement but also captures the moment when the separation occurs near the trailing edge in a low-order manner. We further show that it is possible to perform sensor-based reconstruction through the present low-dimensional expression while identifying the sensitivity with respect to aerodynamic responses. The present model trained at $Re = 3 \times 10^6$ is lastly evaluated at the level of a real aircraft operation of $Re = 3 \times 10^7$, exhibiting that the phase dynamics of lift is reasonably estimated from sparse sensors. The current study may provide a foundation towards data-driven real-time analysis of transonic buffet conditions under aircraft operation.

Key words: machine learning, low-dimensional models, shock waves

1. Introduction

Transonic buffet phenomena determine the high-speed limit of a flight envelope. To extend the flight envelope towards the high-speed side with a better and safer design of modern commercial aircraft that includes asymmetric supercritical wings, the transonic buffet, specifically referred to as Type II buffet (Giannelis, Vio & Levinski 2017), needs to be tamed, rooted in profound understandings of its physics and practically useful models to describe the buffet. While extensive analyses using simulations and experiments have been performed providing a variety of posits to describe complex behaviours of transonic airfoil buffet, what is commonly believed is that there exists a self-sustained shock buffet cycle (Giannelis *et al.* 2017). We pose a question of whether such a seemingly complex, but cyclic dynamics of transonic buffet phenomena can be described in a low-order manner with nonlinear machine learning.

The aerodynamic instability known as transonic buffet, characterised by self-sustained shock-wave oscillations on aircraft wings, needs to be taken into account during transonic or high-subsonic flight. This phenomenon arises because shock waves can form when the wing geometry accelerates the flow along the leading edge of the suction side, generating a localised supersonic region (Tijdeman & Seebass 1980). The occurrence of transonic buffet depends on a flow condition characterised by a combination of parameters such as Mach number, Reynolds number and angle of attack.

To facilitate characterisation of transonic buffet phenomena, a range of numerical and experimental endeavours have been carried out. Such studies on the transonic buffet are classified based on their focus on dimension in phenomena, namely two-dimensional and three-dimensional. In the two-dimensional airfoil buffet, chordwise large-scale shock oscillations occur, which are numerically and experimentally reproduced by confining a flow field in a narrow spanwise domain (Lusher, Sansica & Hashimoto 2024). The chordwise shock oscillations result in a distinct spectral peak at a low frequency generally smaller than 0.1, for example, a Strouhal number $St \approx 0.06$ for the OAT15A supercritical airfoil (Deck 2005; Jacquin *et al.* 2009; Fukushima & Kawai 2018; Cuong Nguyen, Terrana & Peraire 2022).

On the other hand, the three-dimensional buffet is caused due to characteristics associated with the three-dimensionality of the wing, such as swept and taper effects. One notable feature of the three-dimensional buffet, absent in the two-dimensional buffet, is the occurrence of buffet cells (Iovnovich & Raveh 2015). Buffet cells refer to a cellular flow structure propagating outboard. A range of numerical (Ohmichi, Ishida & Hashimoto 2018; Tamaki & Kawai 2024) and experimental (Meneveau & Katz 2000; Dandois 2016; Sugioka *et al.* 2018, 2021; Masini, Timme & Peace 2020) studies have reported the occurrence of buffet cells.

It has widely been observed that the power spectrum density of relevant quantities, such as the pressure coefficient fluctuation, typically presents a broadband spectrum peak with a Strouhal number ranging from 0.2 to 0.6 (Dandois 2016; Koike *et al.* 2016), 10 times higher in frequency than that of the two-dimensional buffet counterpart, depending on the sweep angle (Plante, Dandois & Laurendeau 2020; Sugioka, Kouchi & Koike 2022; Lusher, Sansica & Hashimoto 2025). Particularly considering a full-aircraft configuration of the NASA Common Research Model, understanding of the buffet cell structure has been deepened with modal analysis, including tri-global stability analysis (Timme 2020; Sansica & Hashimoto 2023), tri-resolvent analysis (Houtman, Timme & Sharma 2023), dynamic mode decomposition (Ohmichi *et al.* 2018) and its Hankel variant (Asada & Kawai 2025). Based on them, the buffet cell has been recognised as a key player in the self-sustaining instability mechanism of a three-dimensional buffet. However, there

is still no widely accepted physical model that explains the self-sustaining mechanism of a three-dimensional buffet.

While acknowledging the significance of buffet cells, this study focuses on the two-dimensional airfoil buffet mechanism, which remains active and critical even under three-dimensional buffet conditions. Sugioka *et al.* (2018) experimentally demonstrated that shock-wave oscillations over the NASA Common Research Model at high angles of attack exhibit behaviour similar to that of a two-dimensional buffet. Paladini *et al.* (2019) showed that a two-dimensional global instability mode, akin to that observed in airfoil buffet (Crouch *et al.* 2009), can coexist with a spanwise-varying three-dimensional mode associated with buffet cells. Similar modal structures have been reported by Crouch *et al.* (2018, 2019). Paladini *et al.* (2019) performed a wavemaker analysis to reveal that the two-dimensional mode is primarily linked to the shock-wave dynamics, whereas the spanwise-varying mode originates from the separated shear layer. These findings highlight the importance of considering not only the three-dimensional buffet cells but also the underlying two-dimensional instability mechanisms that remain fundamental to understanding buffet phenomena.

For these reasons, the mechanism of self-sustained large-scale shock oscillations is of particular interest in the community (Lee 2001; Iwatani *et al.* 2023). While a Reynolds-averaged formulation had been considered for numerical investigations (Crouch *et al.* 2009; Iovnovich & Raveh 2012; Sartor, Mettot & Sipp 2015), recent advancements in computational resources along with wall-modelling approaches enable performing large-eddy simulations (LES) (Fukushima & Kawai 2018; Tamaki & Kawai 2024; Goc *et al.* 2025). This offers further reliable assessments of transonic buffet flows by accurately capturing the interaction between the shock wave and the turbulent boundary layer. Along with spatiotemporal high-resolution measuring techniques such as laser Doppler velocimetry (Jacquin *et al.* 2009), particle image velocimetry (D'Aguanno *et al.* 2021) and schlieren visualisation (Schauerte & Schreyer 2023), experimental studies have not only provided a simplified model of transonic buffet supporting the understanding of buffet phenomena (Lee 1990; Crouch, Garbaruk & Magidov 2007) but also suggested passive control devices to suppress buffet-associated instabilities (Lagemann *et al.* 2024). However, the self-sustaining mechanisms of the transonic airfoil buffet still require further clarification.

In analysing transonic buffet flows with a large degree of freedom in the direction of space, time and flow parameters, one can consider applying data-driven order-reduction techniques to flow-field snapshots made available through simulations and experiments. For example, proper orthogonal decomposition (POD) (Lumley 1967) has been considered to obtain a low-order representation of transonic buffet phenomena (Ohmichi *et al.* 2018; Poplingher, Raveh & Dowell 2019; Iwatani, Asada & Kawai 2022; Sansica *et al.* 2022). However, seeking a minimal representation of unsteady flows with such a linear technique is generally challenging because given data are linearly projected onto a flat manifold (Graham & Floryan 2021).

To extract a low-order representation that best captures the underlying characteristics of transonic buffet flows from data, this study considers a nonlinear autoencoder-based compression (Hinton & Salakhutdinov 2006). Nonlinear activation functions inside an autoencoder enable better compression of unsteady flow data compared with linear techniques, which has been discussed with wake shedding (Omata & Shirayama 2019; Murata, Fukami & Fukagata 2020), channel flow (Fukami *et al.* 2019; Yousif, Yu & Lim 2022), Kolmogorov turbulence (Page *et al.* 2024) and aerodynamic flows under gusty environments (Mousavi & Eldredge 2025). Compressed representations obtained from the autoencoder can be used for a range of analyses including mode decomposition

(Fukami, Nakamura & Fukagata 2020; Mo, Traverso & Magri 2024), dynamical modelling (Fukami *et al.* 2021b; Constante-Amores & Graham 2024; Solera-Rico *et al.* 2024), shape optimisation (Tran *et al.* 2024) and flow control (Linot, Zeng & Graham 2023; Liu, Beckers & Eldredge 2025).

Although a nonlinear autoencoder can be employed as a powerful data compressor of unsteady flows, it is important to note that careful use of an autoencoder by incorporating prior knowledge of physics is essential to promote understanding of flows in a low-order latent space (Fukami & Taira 2023). It is challenging to use compressed variables obtained through a naive application of a standard autoencoder for characterising and controlling unsteady flows (Fukami, Nakao & Taira 2024; Smith *et al.* 2024). In response, we incorporate aerodynamic coefficients into the nonlinear autoencoder formulation in identifying a low-order subspace. Equipped with this observable-augmented autoencoder, this study reveals the existence of a three-dimensional representation of transonic airfoil buffet flows, which describes the complex phenomena over the buffet cycle dynamics in a compact manner. Furthermore, the current model trained at a wind-tunnel-scale Reynolds number based on a chord length $Re \sim 10^6$ can be used for sparse-sensor reconstruction of aerodynamic responses at the level of a real-aircraft-operation high Reynolds number $Re \sim 10^7$. The present approach may facilitate data-driven analysis of transonic buffet flows across a range of Reynolds numbers.

This paper is organised as follows. The simulation set-up used for data generation and flow physics are expressed in § 2. The present autoencoder technique is described in § 3. Results and discussion are presented in § 4. Conclusions are offered in § 5.

2. Transonic airfoil buffet flows at high Reynolds numbers

This study seeks a low-dimensional representation of two-dimensional transonic airfoil buffet flows, capturing time-varying characteristics over the buffet cycle using nonlinear machine learning. We consider datasets of Fukushima & Kawai (2018) generated by wall-modelled LES of the transonic buffet over the OAT15A supercritical airfoil at a high Reynolds number of $Re = u_\infty c / \nu_\infty = 3 \times 10^6$ for nonlinear machine-learning compression. Here, u_∞ , c and ν_∞ describe the free-stream velocity, the chord length and the kinematic viscosity, respectively. Following the observation in our previous study (Fukushima & Kawai 2018), we consider two different Mach numbers of $M_\infty = u_\infty / a_\infty = (0.715, 0.730)$, where a_∞ is the free-stream sonic speed. While a steady shock wave is observed at $M_\infty = 0.715$, the unsteady shock oscillating buffet phenomena emerge on increasing the Mach number to 0.730. Involving both non-buffet and buffet conditions in the present datasets for the nonlinear machine-learning analysis enables extracting the difference between them in a low-order manner. All the physical variables throughout the paper are normalised using combinations of c , a_∞ and the density ρ_∞ . We further consider a higher-Reynolds-number case of $Re = 3 \times 10^7$ with $M_\infty = 0.730$, exhibiting unsteady buffet phenomena, to evaluate the applicability of the current technique trained at a wind-tunnel-scale Reynolds number $Re \sim 10^6$ to a scenario at a real aircraft-scale Reynolds number $Re \sim 10^7$.

The computational mesh used in the present study is shown in figure 1. The spatially filtered compressible Navier–Stokes equations are numerically solved, where the LES with modelled wall shear stresses and wall heat fluxes resolves the outer-layer turbulence (Fukushima & Kawai 2018). We follow our previous studies (Kawai & Larsson 2012, 2013; Fukushima & Kawai 2018) for the numerical schemes as well as the treatment of the boundary conditions.




Figure 1. The computational grid used in the present wall-modelled LES of two-dimensional transonic airfoil buffet at a high Reynolds number (Fukushima & Kawai 2018). An instantaneous streamwise velocity field u near the wall and the density gradient magnitude $|\nabla \rho|$ are superposed. The grey grid lines are displayed every fifth point in the g_1 and g_2 (wall-normal) directions. The inset is focused on the region of the shock wave–turbulent boundary layer interactions with the grey grid lines plotted every fifteenth point in the g_1 direction and every fifth point in the g_2 direction.

The spatial derivatives at interior grid points are evaluated using the sixth-order compact differencing scheme (Lele 1992). Time integration is performed with the third-order total variation diminishing Runge–Kutta scheme (Gottlieb & Shu 1998). To accurately resolve the shock wave, the localised artificial diffusivity method is employed with the sixth-order compact scheme (Kawai, Shankar & Lele 2010). While we compute the subgrid-scale turbulent eddy viscosity with a selective mixed-scale model (Lenormand, Sagaut & Ta Phuoc 2000), the equilibrium wall model (Kawai & Larsson 2012) is considered.

The computational mesh for the present wall-modelled LES is designed based on the grid resolution requirements (Kawai & Larsson 2012; Larsson *et al.* 2016). Although we use the same mesh at both Reynolds numbers, $Re = 3 \times 10^6$ and 3×10^7 , the employed mesh satisfies the resolution requirements across the streamwise domain of the attached fully turbulent boundary layer upstream of the shock wave ($0.2 \lesssim x/c \lesssim 0.35$), providing more than 23–25 grid points in each direction per boundary-layer thickness. Specifically, the mesh resolves the boundary layer with at least 29, 34 and 38 points in the wall-normal direction at $x/c \approx 0.2$, 0.25 and 0.3, respectively. In the wall-parallel directions, the resolution corresponds to at least 23, 28 and 33 grid points per local boundary-layer thickness at the same stream locations. These values meet the standards for wall-modelled LES resolution (Kawai & Larsson 2012).

Furthermore, previous studies have reported that wall-modelled LES with the equilibrium wall model can reasonably produce the flow states associated with the interaction between the shock waves and turbulent boundary layer even with the simplification of the equilibrium wall model (Bermejo-Moreno *et al.* 2014; Fukushima & Kawai 2018; De Vanna *et al.* 2022; Sashida *et al.* 2024). Therefore, the present wall-modelled LES provides a high-fidelity dataset for the present nonlinear machine-learning analysis. Further details on the simulation set-up are provided in Fukushima & Kawai (2018).




Figure 2. Lift coefficient and pressure fields at $M_\infty = 0.715$ (a–d) and 0.730 (e–h). A note concerning the shock location is provided underneath each contour of $M_\infty = 0.730$. The arrow in each subcontour represents the direction of shock movement.

The temporal evolution of lift coefficient C_L and a sectional pressure field p extracted from the wing centre in the spanwise direction at $Re = 3 \times 10^6$ obtained through the present simulation is presented in figure 2. The case for $M_\infty = 0.715$ shows statistically steady states, producing small fluctuations of lift over time. The shock mostly appears at $x/c \approx 0.55$ while slightly oscillating in the streamwise direction on the wing.

In contrast, the case for $M_\infty = 0.730$ clearly presents its time-varying feature associated with self-sustained large-scale shock oscillation. The shock wave periodically moves in large amplitude over the wing while the separation near the trailing edge is triggered depending on the shock location, which coincides with observations in wind-tunnel experiments (Jacquin *et al.* 2009). Correspondingly, the lift response also exhibits a periodic signal over the buffet cycle. Hence, the phase of shock location over the buffet cycle is almost identical to that of lift. The separation height is particularly increased when the shock wave moves upstream, which is shown later. The interaction between the wake and separation at this stage causes the upstream-travelling wave (Lee 2001; D’Aguanno *et al.* 2021; Iwatani *et al.* 2023). The lift response is greatly affected by the time-varying area size of supersonic flow along with the aforementioned processes. Note that these buffet dynamics are further discussed and quantified later with observation in the machine-learning-based low-dimensional subspace.

3. Nonlinear machine-learning-based compression of transonic airfoil buffet flows

To seek a low-dimensional representation of transonic airfoil buffet flows from data, we consider a nonlinear autoencoder-based data compression (Hinton & Salakhutdinov 2006). An autoencoder $\mathcal{F}_{AE}$ aims to reconstruct (or output) the same data as the input data $\mathbf{q} \in \mathbb{R}^n$. The autoencoder is designed to possess a bottleneck, referred to as a latent space $\xi \in \mathbb{R}^m$, as illustrated in figure 3. The latent dimension m is generally set to be much smaller than the original data dimension n such that $m \ll n$. Hence, the latent vector ξ can be considered as a compressed representation of the given data $\mathbf{q}$ if the autoencoder $\mathcal{F}_{AE}$ accurately reconstructs the data. The aforementioned process is described as

$$\mathbf{q} \approx \mathcal{F}_{AE}(\mathbf{q}) = \mathcal{F}_d(\mathcal{F}_e(\mathbf{q})), \quad \xi = \mathcal{F}_e(\mathbf{q}), \quad \mathbf{q} \approx \hat{\mathbf{q}} = \mathcal{F}_d(\xi), \quad (3.1)$$




Figure 3. Lift-augmented nonlinear autoencoder (Fukami & Taira 2023).

where $\widehat{(\cdot)}$ denotes a reconstructed variable and $\mathcal{F}_e$ and $\mathcal{F}_d$ correspond to an encoder and a decoder, respectively. A range of neural-network models with nonlinear activation functions can be considered for the construction of autoencoder $\mathcal{F}_{AE}$. The use of nonlinear activation functions promotes network capabilities, providing better compression than linear techniques, which is mathematically proven through the relationship between a linear activation autoencoder and other linear compression approaches (Oja 1982; Bourlard & Kamp 1988; Fukami *et al.* 2021a).

We consider a sectional pressure field sampled from the wing centre in the spanwise direction as the input and output $\mathbf{q}$ of a nonlinear autoencoder to extract the underlying characteristics of transonic airfoil buffet flows. While a standard autoencoder achieves significant data compression of fluid flows, it is often challenging to interpret the identified subspace in a physically understandable manner. To facilitate the present latent identification from the viewpoint of aerodynamics, this study uses a lift-augmented nonlinear autoencoder (Fukami & Taira 2023) producing a lift response from the latent vector through a branch network, as illustrated in figure 3. The optimisation for the parameters (or weights) $\mathbf{w}$ inside the lift-augmented autoencoder is performed with

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w}} \left[\|\mathbf{q} - \hat{\mathbf{q}}\|_2^2 + \beta \|C_L - \hat{C}_L\|_2^2 \right], \quad (3.2)$$

where β balances the pressure field and lift reconstruction loss terms. This weighting parameter β is set to 0.03 and 0.05 based on the L-curve analysis (Hansen & O’Leary 1993) for the observable-augmented autoencoder, while a regular autoencoder, i.e. $\beta = 0$, is also considered for comparison. To minimise the above cost function, the model needs to accurately estimate $C_L(t)$ while performing data compression of the pressure field data $\mathbf{q}(t)$. In other words, the current formulation enables $\mathbf{w}$ to be tuned to capture structures appearing over the buffet cycle that are associated with the lift response. As the periodic shock movement over an OAT15A airfoil, clearly observed in the pressure field, is highly correlated with the lift coefficient $C_L(t)$, the resulting low-dimensional representation is expected to emphasise aerodynamically important events during the buffet cycle.

The current dataset for the nonlinear autoencoder analysis is composed of 6800 snapshots with $M_\infty = 0.715$ (non-buffet condition) over 30.8 non-dimensional time, $t/(c/u_\infty)$, and 17 300 snapshots with $M_\infty = 0.730$ (buffet condition) over 26.4 non-dimensional time. We consider a subdomain of $(x, y)/c \in [-0.6, 1.5] \times [-0.5, 1.3]$ with spatially uniform grid points $(N_x, N_y) = (480, 200)$ extracted from the entire computational domain for the data-driven analysis, where the leading edge of the wing is

Encoder		Decoder		Lift subnetwork	
Layer	Data size	Layer	Data size	Layer	Data size
Input $\mathbf{q}$	(480, 200)	Latent vector ξ	(3)	Latent vector ξ	(3)
Conv. (3, 3, 16)	(480, 200, 16)	MLP	(16)	MLP	(32)
Conv. (3, 3, 16)	(480, 200, 16)	MLP	(32)	MLP	(64)
Maxpooling (2, 2)	(240, 100, 16)	MLP	(256)	MLP	(32)
Conv. (3, 3, 16)	(240, 100, 16)	MLP	(480)	Output $\hat{C}_L$	(1)
Conv. (3, 3, 16)	(240, 100, 16)	(Reshape)	(12, 5, 8)		
Maxpooling (2, 2)	(120, 50, 16)	Conv. (3, 3, 8)	(12, 5, 8)		
Conv. (3, 3, 8)	(120, 50, 8)	Conv. (3, 3, 8)	(12, 5, 8)		
Conv. (3, 3, 8)	(120, 50, 8)	Upsampling (5, 5)	(60, 25, 8)		
Maxpooling (2, 2)	(60, 25, 8)	Conv. (3, 3, 8)	(60, 25, 8)		
Conv. (3, 3, 8)	(60, 25, 8)	Conv. (3, 3, 8)	(60, 25, 8)		
Conv. (3, 3, 8)	(60, 25, 8)	Upsampling (2, 2)	(120, 50, 8)		
Maxpooling (5, 5)	(12, 5, 8)	Conv. (3, 3, 16)	(120, 50, 16)		
Conv. (3, 3, 8)	(12, 5, 8)	Conv. (3, 3, 16)	(120, 50, 16)		
Conv. (3, 3, 8)	(12, 5, 8)	Upsampling (2, 2)	(240, 100, 16)		
(Reshape)	(480)	Conv. (3, 3, 16)	(240, 100, 16)		
MLP	(256)	Conv. (3, 3, 16)	(240, 100, 16)		
MLP	(64)	Upsampling (2, 2)	(480, 200, 16)		
MLP	(32)	Conv. (3, 3, 16)	(480, 200, 16)		
MLP	(16)	Conv. (3, 3, 16)	(480, 200, 16)		
Latent vector ξ	(3)	Output $\hat{\mathbf{q}}$	(480, 200)		

Table 1. The architecture of observable-augmented nonlinear autoencoder. The convolutional layers are denoted as ‘Conv.’. The size of the convolutional filter F and the number of the filter K are shown for each convolutional layer as (F, F, K) . The maxpooling/upsampling ratio R is shown for each layer as (R, R) .

positioned at the origin. The interior of the wing is set to be zero. As a fixed angle of attack is considered for all the data in this study, the model is not affected by this operation. The present autoencoder is composed of convolutional neural networks (LeCun *et al.* 1998) and multi-layer perceptrons (Rumelhart, Hinton & Williams 1986) following the original study of the lift-augmented autoencoder, as summarised in table 1. While the convolutional network learns large-scale structures in a flow field through filter-based operations, the multi-layer perceptrons are used for the bottleneck part of the autoencoder, where the data dimension is very low and the spatial coherence is less important than the complex relationship among the latent variables (Fukagata & Fukami 2025). This combination enables data-driven compression of fluid flow data with reasonable computational costs compared with a model based solely on a multi-layer perceptron that often encounters the curse of dimensionality (Fukami *et al.* 2021a; Morimoto *et al.* 2021). Further details on machine-learning set-ups with the present L-curve analysis for the decision of β are given in Appendix A and a sample code at <https://github.com/kfukami/Observable-AE>.

4. Results and discussion

4.1. Latent space identification of transonic airfoil buffet flows

This section discusses data-driven compression and the resulting subspace identification of transonic airfoil buffet flows. Let us first examine the latent dimension that accurately reproduces the original flow state. The relationship between the latent dimension n_ξ and the L_2 reconstruction error norm ε_q is shown in figure 4. Here, the L_2 reconstruction error norm between a variable $\mathbf{f}$ and its reconstruction $\hat{\mathbf{f}}$ is defined as $\varepsilon_f = \|\mathbf{f} - \hat{\mathbf{f}}\|_2^2 / \|\mathbf{f}'\|_2^2$,




Figure 4. Comparison of compression performance for transonic airfoil buffet flow data between linear POD and a standard nonlinear autoencoder (AE, $\beta = 0$). (a) The relationship between the latent dimension n_ξ and the L_2 reconstruction error ε . (b) Representative reconstructed pressure snapshots with $n_\xi = (1, 3, 5)$ for $M_\infty = 0.730$ with (c) the reference field. (d) The absolute error field $e_{L1} = |\mathbf{q} - \hat{\mathbf{q}}|$ corresponding to panels in (b).

where $\mathbf{f}'$ represents the fluctuation of $\mathbf{f}$ from the time-averaged value. While a standard nonlinear autoencoder without lift incorporation, i.e. $\beta = 0$, is considered for this analysis, linear POD is also used for comparison.

The nonlinear autoencoder is superior to POD across the latent dimension, suggesting that the use of nonlinear activation functions inside the model facilitates compression performance. Compared with the POD-based reconstruction exhibiting high error near the shock, the autoencoder accurately reproduces a flow state, as presented in figure 4. We also find that the error curve of the autoencoder plateaus once the latent dimension reaches three. This reveals that the primary large-scale feature of the pressure fields for the present transonic airfoil buffet flows at $Re = 3 \times 10^6$ can be represented with solely three-dimensional latent variables with nonlinear machine learning. To achieve a similar reconstruction level of $\varepsilon_q \approx 0.1$ to a nonlinear autoencoder with $n_\xi = 3$, 85 linear POD modes are needed.

The plateau behaviour for the autoencoder is in part due to the present network architecture shown in table 1, which compresses data with 480 dimensions given by the portion of the convolutional network to be $\mathcal{O}(10^0)$ using multi-layer perceptrons. A similar observation of producing plateau behaviour in capturing dominant large-scale features has recently been found (Fukami, Smith & Taira 2025) for extremely strong vortex–airfoil interactions with turbulent vortical structures. It is anticipated that the error would be further reduced once fine-scale structures begin to be captured in the latent space with much larger latent dimensions. Since large-scale motions have already been extracted with $n_\xi = 3$, the resulting curve for the autoencoder likely exhibits a step-type behaviour in




Figure 5. Latent subspace identified by a standard autoencoder ($\beta = 0$) and the lift-augmented autoencoder ($\beta = 0.03$ and 0.05) coloured by the cases of different Mach numbers $M_\infty = (0.715, 0.730)$ (a) and the time-varying lift coefficient $C_L(t)$ (b). The pressure fields over time corresponding to the points (i)–(v) in the latent space are also shown. The arrow in each subcontour represents the direction of shock movement. The zoomed-in view of wake and the downstream region visualised with a different colour scheme are also depicted to emphasise the interaction between the wake, shock and turbulent boundary layer.

which the plateaued error reduces again once the latent dimension becomes sufficiently large. Hereafter, we choose a latent dimension of 3 for the discussions.

Next, we examine the behaviour of low-dimensionalised transonic airfoil buffet flows in the latent space. The three-dimensional subspace identified by a standard autoencoder ($\beta = 0$) and the lift-augmented autoencoder ($\beta = 0.03$ and 0.05) is exhibited in figure 5. For all cases, the trajectory for the non-buffet and buffet cases appears in different regions of the latent space. The non-buffet case for $M_\infty = 0.715$ across the autoencoders is described in a similar way, that is, a small-sized circle-like orbit. This representation likely corresponds to the statistically steady dynamics with small oscillations of aerodynamic responses for the present non-buffet flows, which is evident from the reconstruction of lift response and pressure fields for the non-buffet case presented in Appendix B.

While all the present subspaces capture the relationship between the non-buffet and buffet cases and the characteristics of the non-buffet flow in a low-order manner, the latent expression for the buffet case of $M_\infty = 0.730$ shows a clear difference by introducing




Figure 6. Time trace of latent vectors ξ obtained from nonlinear autoencoders, shock location x_s , lift coefficient C_L and separation height h for the buffet case.

the lift augmentation. This can be observed with the difference in the relative location of the low-dimensionalised flow states (i), (iv) and (v). Here, the shock in the flow field (i) moves downstream while that in (iv) and (v) moves upstream. The standard model encodes them into nearby regions in the latent space. In contrast, their locations begin to differ due to the lift augmentation. Consequently, the low-order trajectory with $\beta = 0.05$ presents a geometric structure possessing two wings, while that with $\beta = 0$ and 0.03 rather shows a regular cyclic orbit.

To discuss what physics are captured in the present low-order representation, the temporal behaviour of latent vectors $\xi(t)$ is compared with the shock location $x_s(t)$, the lift coefficient $C_L(t)$ and the separation height $h(t)$, as shown in figure 6. Here, the shock location x_s is defined as a streamwise position at which the density gradient magnitude $|\nabla\rho|$ takes the maximum value. The separation height h is set to be a distance from the wall in which the streamwise momentum ρu becomes 0 at $x/c = 0.6$ in measuring across the wall-normal direction.

The latent expression from the standard autoencoder emphasises the cyclic behaviour of shock location as the notable peak of latent vectors at $t \approx 13$. With the lift incorporation of $\beta = 0.05$, the latent vectors possess an additional dominant peak around $t = 20$, corresponding to the emergence of the wing-type geometric structure in the low-order subspace. While this moment is under-evaluated with $\beta = 0$ and 0.03, we find that the peak appearing at $t \approx 20$ coincides with the timing when the separation height h is increased, as shown in figure 6. This increase in the separation height h is attributed to the upstream-moving shock wave, not only producing a strong shock due to the increase of relative shock Mach number but also inducing a large separation due to a strong shock adverse pressure gradient. In this manner, the separation height varies depending on the direction of shock movement across the streamwise direction, i.e. relative shock Mach number, in

addition to the shock location. Hence, it can be argued that the current lift augmentation well captures the relationship between the shock motion and the aerodynamic responses in its latent representation. Although the flow field data themselves given as the input may also include phase information of the buffet cycle as the phase of shock location matches that of lift response as presented in [figure 6](#), the present observation suggests that providing an aerodynamic variable as an observable output through the subnetwork is essential to identify the physically interpretable subspace. The dependence of the latent representation geometry on the number of training samples and the initial random seed assigned to the weights in the observable-augmented autoencoder is examined in [Appendices C and D](#), respectively.

Note that all the latent spaces across β represent the cyclic transonic buffet dynamics while achieving the same level of reconstruction through the decoder. The latent expression hence becomes stretched by highlighting the events associated with a given observable. In other words, all the latent subspaces are regarded as the compact representation of transonic airfoil buffet flows, although their ways of presentation are different from each other. The present lift augmentation can highlight aerodynamically important events as a manifold geometry while a regular model does not capture them in an interpretable manner, e.g. points (i), (v) and (iv) in [figure 5](#).

4.2. Sparse-sensor reconstruction of transonic airfoil buffet flows via low-order subspace

The current findings through autoencoder compression imply that the right set of variables may capture the essence of transonic airfoil buffet flows. This also makes us anticipate that sparse sensors could also be such a set of low-order variables, thereby achieving sparse-sensor-based reconstruction. Furthermore, of interest here is whether it is possible to gain situational awareness from sparse sensors towards guiding flight operations based on insights into the physically interpretable latent subspace. Based on this viewpoint, we further consider leveraging the discovered low-order subspace for the data-driven global flow field reconstruction.

Since the decoder $\mathcal{F}_d$ provides the pressure field from the latent vector, we aim to estimate the latent vector $\boldsymbol{\xi}(t)$ from sparse sensors $\boldsymbol{s}(t)$ by preparing an independent machine-learning model $\mathcal{F}_s$. By feeding the estimated latent vector $\hat{\boldsymbol{\xi}} = \mathcal{F}_s(\boldsymbol{s}(t))$ into the pretrained decoder $\mathcal{F}_d$, a pressure field $\boldsymbol{q}(t)$ is reconstructed, as illustrated in [figure 7](#). The above-mentioned procedure is expressed as

$$\boldsymbol{q}(t) \approx \hat{\boldsymbol{q}}(t) = \mathcal{F}_d(\hat{\boldsymbol{\xi}}(t)) = \mathcal{F}_d(\mathcal{F}_s(\boldsymbol{s}(t))), \quad (4.1)$$

with an optimisation for the weights $\boldsymbol{w}_s$ of the latent vector estimator $\mathcal{F}_s$:

$$\boldsymbol{w}_s^* = \operatorname{argmin}_{\boldsymbol{w}_s} \|\boldsymbol{\xi} - \mathcal{F}_s(\boldsymbol{s}; \boldsymbol{w}_s)\|_2^2. \quad (4.2)$$

We use multi-layer perceptrons (Rumelhart *et al.* 1986) with the units of 14–32–64–128–32–3 across the layers for constructing the latent vector estimator $\mathcal{F}_s$ that maps sensor measurements $\boldsymbol{s} \in \mathbb{R}^{n_s}$ to $\boldsymbol{\xi} \in \mathbb{R}^3$, where n_s represents the number of sensors. This low-order mapping between sparse sensors and the latent vector enables avoiding a naive learning for the relationship between the sensor inputs and the global field output (Fukami, Fukagata & Taira 2023; Eldredge & Mousavi 2025). While such a field reconstruction problem often becomes computationally expensive due to a significant difference in data dimension between the input and output, this approach can save costs by leveraging the pretrained decoder.




Figure 7. Sparse-sensor-based reconstruction via the low-order subspace. (a) Pressure sensor placements on the wall and responses in time. (b) The present full state reconstruction combined with a latent vector estimator $\mathcal{F}_s$ and the pretrained decoder $\mathcal{F}_d$. An example of the reconstructed field with the L_2 error norm ε_q and reproduced lift coefficient from 14 sensors is shown.

An example of the reconstructed pressure field and estimated lift coefficient from 14 sensors is shown in figure 7. Here, these sensors are placed along the airfoil surface in an equispaced manner, enabling a comprehensive analysis of data-driven sensor reduction performed later. We use the latent vector ξ extracted from the lift-augmented autoencoder with $\beta = 0.05$. In addition to the flow state including wake shedding and shock location, the lift response is accurately reproduced from the sensor readings. As implied through the discovery of a low-dimensional subspace, sparse-sensor-based reconstruction is indeed possible for the present transonic airfoil buffet flow.

Furthermore, the minimal number and appropriate placements of sensors can be quantified with the latent vector estimator trained with 16 sensors above and the lift subnetwork prepared for subspace identification. This is achieved by performing a sensitivity analysis between a machine-learning estimate and a given input (Morimoto *et al.* 2022; Chen *et al.* 2024). Considering the gradient between the sensor input s and the output of machine-learning model $\hat{z}$, $\gamma(t) = \partial \hat{z}(t) / \partial s(t)$, the importance of each sensor for estimation, i.e. sensitivity $S(t)$, is quantified as a weighted input:

$$S_j(t) = \gamma_j(t) s_j(t), \quad (4.3)$$

where j is an index of pressure sensor s_j . As an output variable $\hat{z}$, the estimated lift coefficient $\hat{C}_L$ and latent vector $\hat{\xi}$ are considered.

The sensitivity S with respect to the lift and latent vectors is shown in figure 8. In addition to the time trace, the absolute time-averaged values are also presented to further gain insights into the general trend of sensitivities over the buffet cycle. As the present autoencoder is trained such that the latent vector extracts the flow features associated with




Figure 8. Gradient-based sensor sensitivities with respect to the lift and latent vectors. Both the time trace and the time-averaged sensitivities over 14 sensors are shown. The sensor index here corresponds to that shown in figures 7, 9 and 10.

the lift coefficient, both sensitivity maps present a consistent trend in the direction of time and sensor index. Note that high-frequency fluctuations of the sensitivity are caused because the present sensitivity is calculated using the estimate by the machine-learned model, which includes the estimation error varying in time. We have confirmed that the rank of sensor importance is not affected by such high-frequency fluctuations through a preliminary analysis by taking moving averages.

Focusing on the lift estimation, the sign of sensor sensitivity seems to be opposite between the suction (index 2–7) and pressure (index 9–14) sides due to their different role in contributing to lift. The responsible sensors are clearly shown where $|\overline{S}| > 0.02$: sensor 1 at the leading edge, sensor 8 at the trailing edge, sensors 5, 6 and 7 placed on the suction side and sensors 10 and 13 placed on the pressure side. In turn, less sensitive sensors are also identified. Sensors 2 and 3, placed in the supersonic region, particularly show very small $|\overline{S}|$, likely because their sensor signals are less affected by the shock movement compared with others according to figure 7(a).

The present sensitivity information is further leveraged to reduce the number of sensors for subspace estimation. Let us consider removing the sensors following the rank of absolute time-averaged sensitivity $|\overline{S}|$ so that sensors with small contribution to estimation are eliminated while keeping the highly contributing sensors. The relationship between the number of sensors n_s and the estimation errors is shown in figure 9. The error for the latent vector and lift response is depicted on a single plot. The error curves are flat between $n_s = 7$ and 14, exhibiting that accurate estimation of lift and flow fields is achieved up to $n_s = 7$. This is also evident from the reconstructed flow field shown in figure 9, and the estimated latent subspace and lift response presented in figure 10.

All seven sensor readings here report the absolute time-averaged value of $|\overline{S}| > 0.02$, exhibiting a relatively larger value compared with other less-contributing sensors observed




Figure 9. Sensitivity-based sensor reduction. The relationship between the number of sensors n_s and the estimation errors of latent vectors ε_ξ and lift ε_{C_L} is shown. The reconstructed fields are presented with the L_2 error norm ε_q underneath each contour.

in figure 8. Once the sensors are further removed, the error of the latent vector starts to increase. However, the error curve for the lift coefficient presents a slower slope at $n_s \leq 6$ compared with that for the latent space. In fact, the lift response at $n_s = 3$ still exhibits reasonable agreement with the reference data. This is likely because a global quantity of lift coefficient aggregating the flow information over the entire body is easier to estimate than the latent subspace, a representation of the whole flow field itself.

To examine the dependence of reconstruction performance on the choice of sensor-selection technique and compression approach, we further consider the QR pivot-based sensor placement optimisation (Manohar *et al.* 2018) with $n_s = 7$. Their approach finds the optimal sensor locations through QR factorisation with column pivoting applied to the POD bases. Further details on this linear technique are provided in Manohar *et al.* (2018). The original placements of sensors before performing the QR pivot-based reduction are constrained on the wing surface and the same as those used in the autoencoder-based analysis shown in figure 7(a). Here, three approaches are considered:




Figure 10. The estimated latent subspace and estimated lift coefficient across n_s reduced via the sensitivity analysis.

- (i) Estimate the three-dimensional latent vectors ξ based on the sensors reduced via the gradient sensitivity and decode a flow using the nonlinear decoder $\mathcal{F}_d$ (the original formulation).
- (ii) Estimate the dominant three POD coefficients $\mathbf{a}$ based on the sensors reduced via the QR pivot and decode a flow with POD modes.
- (iii) Estimate the three-dimensional latent vectors ξ based on the sensors reduced via the QR pivot and decode a flow using the nonlinear decoder $\mathcal{F}_d$.

For fair comparison, we use the same multi-layer perceptron architecture for all three cases in estimating the latent vectors and the three dominant POD coefficients. The flow fields are then decoded using the nonlinear decoder $\mathcal{F}_d$ or POD modes Φ . While (4.1) is applied for cases (i) and (iii), case (ii) with the POD multi-layer perceptron model with QR pivot-based sensor reduction is expressed as

$$\mathbf{q}(t) \approx \Phi \hat{\mathbf{a}}(t) = \Phi(\mathcal{F}_s(s(t))). \quad (4.4)$$

Let us compare the reduced sensor placements in figure 11. Four sensors (index 5, 7, 8 and 13), reporting high $|\bar{S}|$ with the gradient-based approach, are commonly kept with both sensor-reduction methods through the reduction process. However, the remaining three sensors are placed in a different way. While sensors chosen by the QR pivot are grouped with neighbours (s_3 – s_5 , s_7 – s_8 and s_{13} – s_{14}), the gradient-based method seems to attempt to cover the entire wing surface. This result suggests that the dominant features captured by both POD and the autoencoder make the reduction approach keep the common four sensors, while the subdominant characteristics that are better compressed with the nonlinear autoencoder cause the difference in the location of the remaining three sensors.

The reconstruction fields with cases (i)–(iii) are also shown in figure 11. When using the nonlinear decoder, the reconstruction with the gradient-based approach is slightly better




Figure 11. Dependence of sparse-sensor reconstruction performance on the choice of sensor-reduction technique and compression approach with $n_s = 7$.

than that with the QR pivot. These accurate reconstructions suggest that the high error for case (ii) is primarily due to the use of linear POD modes as a decoder rather than the sensor placements determined by the QR pivot. We note that the error for case (iii) of the QR pivot and the autoencoder latent variables starts to increase with $n_s \leq 6$, similarly to case (i) using the gradient-based method, although not shown. While both the gradient-based method and the QR pivot currently provide a similar level of sensor reduction performance, they could be further improved by accounting for redundancy between sensor readings, which can be quantified with inter-correlations and mutual information.

The present analysis is focused on transonic airfoil buffet flow at $Re = 3 \times 10^6$. While the current Reynolds number may be higher than those often considered for numerical and data-driven analyses in the community, this still resides in the range of wind-tunnel-scale conditions. Of particular interest here is whether the current model trained at a wind-tunnel-scale Reynolds number can be applied to a scenario under a real aircraft operation level of Reynolds number. In response, this study lastly evaluates the applicability of the present method to a transonic airfoil buffet flow at $Re = 3 \times 10^7$ with $M_\infty = 0.730$.

The wall-modelled LES is performed for the case with $(Re, M_\infty) = (3 \times 10^7, 0.730)$ at $\alpha = 3.5^\circ$, as presented in figure 12(a). There is a self-sustained shock buffet cycle that produces almost the same frequency and oscillation amplitude of aerodynamic coefficients as those for $Re = 3 \times 10^6$, as seen in figure 12(b). The difference in the flow between the two Reynolds numbers is examined with the instantaneous streamwise velocity u sampled at the same phase $t/T = 0.70$, where T denotes the time window across the buffet cycle, as depicted in figure 12(c). The shock location moves downward and the separation height becomes greater on increasing the Reynolds number, strengthening the shock wave accompanied by a large adverse pressure gradient and triggering a larger separation, which is also evident from the time-averaged flow fields shown in figure 12(d). Due to the trade-off relationship between the suppression effect of separation due to the increment of Reynolds number and the separation induced by the strong shock wave, the resulting shock-wave oscillation is sustained.

Let us finally apply the present sensor-based reconstruction model trained at $Re = 3 \times 10^6$ to the level of real aircraft operation at $Re = 3 \times 10^7$, as shown in figure 13. Here, we use the latent vector estimator $\mathcal{F}_p$ trained with seven sensors following the observation in figures 9 and 10. The reconstructed fields exhibit a smaller height of shock compared




Figure 12. (a) An instantaneous snapshot of transonic airfoil buffet flows at $Re = 3 \times 10^7$ visualised by the isocontours of the Q -criterion. Comparison of (b) lift coefficient and (c) instantaneous streamwise velocity fields sampled at $t/T = 0.70$ with $Re = 3 \times 10^6$ and $Re = 3 \times 10^7$. (d) Time- and spanwise-averaged streamwise velocity fields at $Re = 3 \times 10^6$ and $Re = 3 \times 10^7$.

with the reference snapshots as such a shock with a greater height does not appear in the training data at $Re = 3 \times 10^6$. However, it is worth noting that the shock locations of the machine-learning reconstruction are constantly evaluated forward compared with that of the reference at $Re = 3 \times 10^7$. Since the shock moves downward on increasing the Reynolds number while keeping its phase as presented in figure 12, this constant shift indicates that the present model may correctly capture the phase information across the buffet cycle even at the current real-aircraft-level Reynolds number. This is further evident from the reproduced lift response. While the magnitude of lift is underestimated due to the difference in Reynolds number between the training and testing data, the temporal trend of the lift signal accurately matches the reference. This observation suggests that nonlinear machine learning can be transferred to scenarios where the characteristics of variables of interest remain relatively consistent across different Reynolds numbers.

5. Concluding remarks

This study examined a low-dimensional representation of transonic airfoil buffet flows at a high Reynolds number with nonlinear machine learning. Wall-modelled LES of flow over the OAT15A supercritical airfoil at Mach numbers $M_\infty = 0.715$ and 0.730 , corresponding to non-buffet and buffet conditions, were performed at a chord-based Reynolds number




Figure 13. Application of the sparse-sensor reconstruction model trained at $Re = 3 \times 10^6$ to a flow at the level of a real aircraft operation of $Re = 3 \times 10^7$. The reconstructed pressure field and lift response are shown.

of $Re = 3 \times 10^6$ to generate the datasets used in the present data-driven analysis. To derive a low-order expression from the data, we considered nonlinear lift-augmented autoencoder-based compression. We found that there exists a compact three-dimensional latent subspace reflecting the characteristics of transonic airfoil buffet flow. The discovered representation captures key flow features, including shock movement and shock-induced separation, in a reduced-order manner.

Based on these physical implications, sparse-sensor-based reconstruction via the learned representation was further performed. Equipped with the sensitivity analysis, the sensor configuration required for accurately reproducing aerodynamic responses can be determined. Finally, the model trained at a wind-tunnel-scale Reynolds number of $Re = 3 \times 10^6$ was assessed at a real aircraft operational level of $Re = 3 \times 10^7$, revealing its ability to reasonably predict phase dynamics of aerodynamic loads from sparse sensors.

While we considered two configurations of buffet/non-buffet conditions at a fixed angle of attack, additional cases with a range of different parameters, including angle of attack, Reynolds number and Mach number, would be needed to fully characterise the whole picture of buffet onset. Although it is anticipated that a low-order subspace capturing the difference in such parameters and the occurrence of transonic buffet could be identified, a major challenge arises from a collection of datasets through large-scale simulations. From this aspect, one can consider data fusion between LES, unsteady Reynolds-averaged Navier–Stokes and experimental measurements to supplement the pros and cons across different datasets with each other in extracting a low-order submanifold with observable-augmented learning (Fukami & Taira 2025).

The present analysis reveals that three latent variables are needed to represent transonic airfoil buffet flows. Although buffet dynamics is often modelled as a self-sustained oscillator subjected to stochastic forcing (Feldhusen-Hoffmann *et al.* 2021; Sansica *et al.* 2022; Crouch, Ahrabi & Kamenetskiy 2024), our findings suggest the necessity of a third dimension. This additional latent dimension likely corresponds to aerodynamic phenomena related to the separation height, according to the observation in figure 6. To characterise this dynamics more precisely, it is essential to investigate the nonlinear modal structures associated with each latent variable. This can be achieved by integrating mode-decomposing autoencoders (Fukami *et al.* 2020; Murata *et al.* 2020), which we plan to pursue in future work.

With the present formulation of observable-augmented learning, users have to choose an appropriate observable from the candidates, and it currently takes some level of computational effort to find a physically relevant subspace. Note, however, that the former point of the non-automatic process enables us to have the opportunity to incorporate

physical or mathematical knowledge based on what we would like to associate with, while the computational cost for the latter point is still manageable as the degree of freedom of observables is much less than that of the original simulations. A series of recent studies on observable-augmented manifold learning have revealed that an appropriate choice of observable assists in compactly extracting physics for a range of unsteady flow scenarios including vortex–airfoil interactions (Fukami & Taira 2023; Fukami *et al.* 2024; Liu *et al.* 2025; Mousavi & Eldredge 2025), vehicle aerodynamics (Tran *et al.* 2024), turbulent boundary layers (Fukami & Taira 2025) and roughness turbulence (Nair *et al.* 2025), enabling the enjoyment to learn physics from data for fluid mechanicians. More broadly, an ‘observable’ here does not need to be a variable. Some applied mathematical techniques, such as persistent homology (Smith *et al.* 2024) and information theory (Fukami & Araki 2025), can also be considered as observables depending on the physics of interest. Hence, adding an observable may be regarded as one approach to support data-driven analysis for unsteady flows.

Based on the current findings considering flows around a wing, the applicability of the present data-driven subspace identification to transonic buffet conditions around a full-aircraft configuration would also be of interest (Asada *et al.* 2023; Tamaki & Kawai 2024). For such cases, a combination of linear, scalable compression techniques such as POD and the present observable augmentation would be helpful to reduce the computational burden (Linot & Graham 2023; Tran *et al.* 2024; Asada & Kawai 2025). The current study may offer a new perspective on the analysis and determination of flight envelopes towards next-generation air vehicle operations.

Acknowledgements. We thank H. Itsui and T. Hattori for fruitful discussions.

Funding. We acknowledge support from MEXT as a programme for promoting research on the Supercomputer Fugaku (research toward DX in aircraft development led by digital flight, JPMXP1020230320). K.F. acknowledges support from the JSPS KAKENHI (grant no. JP25K23418), the JST PRESTO (grant no. JPMJPR25KA), and the MEXT Coordination Funds for Promoting Aerospace Utilization (grant no. JPJ000959). Y.I. acknowledges support from JSPS Grant-in-Aid for JSPS Fellows (JP23KJ0167). S.M. acknowledges support from the Japan Science and Technology Agency (JST) SPRING (JPMJSP2114).

Declaration of interests. The authors report no conflict of interest.

Appendix A. Training procedures and L-curve analysis

Here, we provide details on training procedures and the choice of weighting parameter β in (3.2) for the present observable-augmented nonlinear autoencoder. The Adam optimiser (Kingma & Ba 2014) with the default parameter sets in Keras is used to update the weights through machine-learning training. The maximum number of training iterations is set to be 50 000, while early stopping (Prechelt 1998) with the criterion of a series of 100 continuous epochs is employed to avoid overfitting. We use 70 % of the datasets for training and the remaining 30 % are prepared for validation. The number of grid points $(N_x, N_y) = (480, 200)$ for the current data-driven analysis is determined such that the shock can be represented without exhibiting any discontinuous artefacts, which is evident from a comparison with other resolutions $(N_x, N_y) = (240, 100)$ and $(960, 400)$ shown in figure 14. In using the entire data set of 24 100 snapshots, the training process takes approximately two hours in an NVIDIA A100 GPU environment, and the inference time for each snapshot is 0.003 seconds.

The weighting parameter β in (3.2) is determined based on the L-curve analysis (Hansen & O’Leary 1993) that finds an appropriate regularisation parameter of the cost function, as shown in figure 15. We consider nine different values of β (0.005, 0.01, 0.03, 0.05, 0.1,




Figure 14. Pressure field interpolated onto a spatially uniform grid with a resolution of $(n_x, n_y) = (240, 100)$, $(480, 200)$ and $(960, 400)$.




Figure 15. L-curve analysis for the present observable-augmented autoencoder.

0.5, 1, 5 and 10). The cases with $\beta = 0.03$ and 0.05 , providing low reconstruction errors for the lift response and the pressure field in a balanced manner, are chosen for the present analysis.

Appendix B. Reconstructed variables for the non-buffet case

We exhibit in [figure 16](#) the decoded lift coefficient and pressure fields obtained from the present lift-augmented autoencoder with $\beta = 0.05$. While achieving accurate estimation of the lift coefficient, the reconstructed pressure fields are in agreement with the reference data, reporting less than 8 % L_2 norm error over time. Along with the observation of a small-sized cyclic orbit in [figure 5](#) and small oscillations of the pressure coefficient C_p in [figure 16](#), it is argued that the present model well represents statistically steady dynamics of the non-buffet case in the identified low-order subspace.

Appendix C. Effect of the number of training samples

We examine the dependence of reconstruction performance and latent space geometry on the number of training snapshots by subsampling them to be 25 % and 50 % of the original amount, as presented in [figure 17](#). We use the same autoencoder network with the same weighting parameter β of 0.05 as that used in the original case. The case with 50 % presents a similar result to the original model. However, the latent geometry with the 25 % case starts to deform from the original shape, although there still exists a two-wing-shaped submanifold. Since the dimensionality in the subspace is determined based on whether the given data cover the entire space of the attractor or not, rather than the number of




Figure 16. Decoded lift coefficient and pressure fields via a lift-augmented autoencoder with $\beta = 0.05$ for the non-buffet case with $M_\infty = 0.715$. The flow fields (a–d) correspond to those shown in figure 2. The whole (bottom left) and zoom-in (bottom right) views of the pressure coefficient C_p on the wing surface for the snapshots (a–d) are also presented.

snapshots, the latent dimension is not affected for this analysis, in which we subsample the snapshots while keeping the entire time window.

The deformation of latent space geometry is caused by several factors. There may exist an optimal weighting parameter β for the case with 25 % data. Furthermore, the primary reason is likely less temporal density of data compared to the original case, which may cause miscapturing of some events over the buffet cycle. A sufficient temporal resolution is needed to obtain an interpretable low-order subspace in a data-driven manner.

Appendix D. Uniqueness of latent representation

To consider the uniqueness of latent representation, we examine the dependence of the latent geometry on the initial random seed assigned to the weights in the observable-augmented autoencoder, as shown in figure 18. A weighting parameter β of 0.05 is used for




Figure 17. Dependence of field reconstruction performance and latent space geometry on the number of training snapshots.




Figure 18. Dependence of the latent geometry on the initial random seed assigned to the weights in the observable-augmented autoencoder.

this analysis. The model exhibits reasonable robustness across the three runs, presenting a two-wing-shaped submanifold while distinguishing the non-buffet and buffet cases in a low-order manner. Although this paper only considers a single-network configuration of observable-augmented autoencoder, the results above indicate that a model may provide a similar wing-shaped geometry over a range of the network capacities by choosing the optimal value of β through the L-curve analysis.

REFERENCES

- ASADA, H. & KAWAI, S. 2025 Exact parallelized dynamic mode decomposition with Hankel matrix for large-scale flow data. *Theor. Comput. Fluid Dyn.* **39**, 8.
- ASADA, H., TAMAKI, Y., TAKAKI, R., YUMITORI, T., TAMURA, S., HATANAKA, K., IMAI, K., MAEYAMA, H. & KAWAI, S. 2023 FFVHC-ACE: fully automated Cartesian-grid-based solver for compressible large-eddy simulation. *AIAA J.* **61** (8), 3466–3484.
- BERMEJO-MORENO, I., CAMPO, L., LARSSON, J., BODART, J., HELMER, D. & EATON, J.K. 2014 Confinement effects in shock wave/turbulent boundary layer interactions through wall-modelled large-eddy simulations. *J. Fluid Mech.* **758**, 5–62.

- BOURLARD, H. & KAMP, Y. 1988 Auto-association by multilayer perceptrons and singular value decomposition. *Biol. Cybern.* **59** (4), 291–294.
- CHEN, D., KAISER, F., HU, J.-C., RIVAL, D.E., FUKAMI, K. & TAIRA, K. 2024 Sparse pressure-based machine learning approach for aerodynamic loads estimation during gust encounters. *AIAA J.* **62** (1), 275–290.
- CONSTANTE-AMORES, C.R. & GRAHAM, M.D. 2024 Data-driven state-space and Koopman operator models of coherent state dynamics on invariant manifolds. *J. Fluid Mech.* **984**, R9.
- CROUCH, J.D., AHRABI, B.R. & KAMENETSKIY, D.S. 2024 Weakly nonlinear behaviour of transonic buffet on airfoils. *J. Fluid Mech.* **999**, A8.
- CROUCH, J.D., GARBARUK, A. & MAGIDOV, D. 2007 Predicting the onset of flow unsteadiness based on global instability. *J. Comput. Phys.* **224** (2), 924–940.
- CROUCH, J.D., GARBARUK, A., MAGIDOV, D. & TRAVIN, A. 2009 Origin of transonic buffet on aerofoils. *J. Fluid Mech.* **628**, 357–369.
- CROUCH, J.D., GARBARUK, A. & STRELETS, M. 2018 Global instability analysis of unswept-and swept-wing transonic buffet onset. AIAA paper, 2018-3229.
- CROUCH, J.D., GARBARUK, A. & STRELETS, M. 2019 Global instability in the onset of transonic-wing buffet. *J. Fluid Mech.* **881**, 3–22.
- CUONG NGUYEN, N., TERRANA, S. & PERAIRE, J. 2022 Large-eddy simulation of transonic buffet using matrix-free discontinuous Galerkin method. *AIAA J.* **60** (5), 3060–3077.
- DANDOIS, J. 2016 Experimental study of transonic buffet phenomenon on a 3D swept wing. *Phys. Fluids* **28**, 016101.
- DE VANNA, F., BERNARDINI, M., PICANO, F. & BENINI, E. 2022 Wall-modeled LES of shock-wave/boundary layer interaction. *Intl J. Heat Fluid Flow* **98**, 109071.
- DECK, S. 2005 Numerical simulation of transonic buffet over a supercritical airfoil. *AIAA J.* **43** (7), 1556–1566.
- D’AGUANO, A., SCHRIJER, F.F.J. & VAN OUDHEUSDEN, B.W. 2021 Experimental investigation of the transonic buffet cycle on a supercritical airfoil. *Exp. Fluids* **62**, 1–23.
- ELDRIDGE, J.D. & MOUSAVI, H. 2025 A review of Bayesian sensor-based estimation and uncertainty quantification of aerodynamic flows. arXiv: 2502.20280.
- FELDHUSEN-HOFFMANN, A., LAGEMANN, C., LOOSEN, S., MEYSONNAT, P., KLAAS, M. & SCHRÖDER, W. 2021 Analysis of transonic buffet using dynamic mode decomposition. *Exp. Fluids* **62**, 1–17.
- FUKAGATA, K. & FUKAMI, K. 2025 Compressing fluid flows with nonlinear machine learning: mode decomposition, latent modeling, and flow control. *Fluid Dyn. Res.* **57**, 041401.
- FUKAMI, K. & ARAKI, R. 2025 Information-theoretic machine learning for time-varying mode decomposition of separated aerodynamic flows. *AIAA J.* <https://doi.org/10.2514/1.J065914>
- FUKAMI, K., FUKAGATA, K. & TAIRA, K. 2023 Super-resolution analysis via machine learning: a survey for fluid flows. *Theor. Comput. Fluid Dyn.* **37**, 421–444.
- FUKAMI, K., HASEGAWA, K., NAKAMURA, T., MORIMOTO, M. & FUKAGATA, K. 2021a Model order reduction with neural networks: application to laminar and turbulent flows. *SN Comput. Sci.* **2** (467), 1–16.
- FUKAMI, K., MURATA, T., ZHANG, K. & FUKAGATA, K. 2021b Sparse identification of nonlinear dynamics with low-dimensionalized flow representations. *J. Fluid Mech.* **926**, A10.
- FUKAMI, K., NABAE, Y., KAWAI, K. & FUKAGATA, K. 2019 Synthetic turbulent inflow generator using machine learning. *Phys. Rev. Fluids* **4**, 064603.
- FUKAMI, K., NAKAMURA, T. & FUKAGATA, K. 2020 Convolutional neural network based hierarchical autoencoder for nonlinear mode decomposition of fluid field data. *Phys. Fluids* **32**, 095110.
- FUKAMI, K., NAKAO, H. & TAIRA, K. 2024 Data-driven transient lift attenuation for extreme vortex gust–airfoil interactions. *J. Fluid Mech.* **992**, A17.
- FUKAMI, K., SMITH, L. & TAIRA, K. 2025 Extreme vortex-gust airfoil interactions at Reynolds number 5000. *Phys. Rev. Fluids* **10**, 090701.
- FUKAMI, K. & TAIRA, K. 2023 Grasping extreme aerodynamics on a low-dimensional manifold. *Nat. Commun.* **14**, 6480.
- FUKAMI, K. & TAIRA, K. 2025 Observable-augmented manifold learning for multi-source turbulent flow data. *J. Fluid Mech.* **1010**, R4.
- FUKUSHIMA, Y. & KAWAI, S. 2018 Wall-modeled large-eddy simulation of transonic airfoil buffet at high Reynolds number. *AIAA J.* **56** (6), 2372–2388.
- GIANNELIS, N.F., VIO, G.A. & LEVINSKI, O. 2017 A review of recent developments in the understanding of transonic shock buffet. *Prog. Aerosp. Sci.* **92**, 39–84.
- GOC, K.A., AGRAWAL, R., BOSE, S.T. & MOIN, P. 2025 Studies of transonic aircraft flows and prediction of initial buffet using large-eddy simulation. *J. Aircraft* **62** (3), 530–550.

- GOTTLIEB, S. & SHU, C.-W. 1998 Total variation diminishing Runge–Kutta schemes. *Maths Comput.* **67** (221), 73–85.
- GRAHAM, M.D. & FLORYAN, D. 2021 Exact coherent states and the nonlinear dynamics of wall-bounded turbulent flows. *Annu. Rev. Fluid Mech.* **53** (1), 227–253.
- HANSEN, P.C. & O’LEARY, D.P. 1993 The use of the L-curve in the regularization of discrete ill-posed problems. *SIAM J. Sci. Comput.* **14** (6), 1487–1503.
- HINTON, G.E. & SALAKHUTDINOV, R.R. 2006 Reducing the dimensionality of data with neural networks. *Science* **313** (5786), 504–507.
- HOUTMAN, J., TIMME, S. & SHARMA, A. 2023 Resolvent analysis of a finite wing in transonic flow. *Flow* **3**, E14.
- IOVNOVICH, M. & RAVEH, D.E. 2012 Reynolds-averaged Navier–Stokes study of the shock-buffet instability mechanism. *AIAA J.* **50** (4), 880–890.
- IOVNOVICH, M. & RAVEH, D.E. 2015 Numerical study of shock buffet on three-dimensional wings. *AIAA J.* **53** (2), 449–463.
- IWATANI, Y., ASADA, H. & KAWAI, S. 2022 POD, DMD, and resolvent analysis of transonic airfoil buffet. *AIAA paper*, 2022-0461.
- IWATANI, Y., ASADA, H., YEH, C.-A., TAIRA, K. & KAWAI, S. 2023 Identifying the self-sustaining mechanisms of transonic airfoil buffet with resolvent analysis. *AIAA J.* **61** (6), 2400–2411.
- JACQUIN, L., MOLTON, P., DECK, S., MAURY, B. & SOULEVANT, D. 2009 Experimental study of shock oscillation over a transonic supercritical profile. *AIAA J.* **47** (9), 1985–1994.
- KAWAI, S. & LARSSON, J. 2012 Wall-modeling in large eddy simulation: length scales, grid resolution, and accuracy. *Phys. Fluids* **24**, 015105.
- KAWAI, S. & LARSSON, J. 2013 Dynamic non-equilibrium wall-modeling for large eddy simulation at high Reynolds numbers. *Phys. Fluids* **25**, 015105.
- KAWAI, S., SHANKAR, S.K. & LELE, S.K. 2010 Assessment of localized artificial diffusivity scheme for large-eddy simulation of compressible turbulent flows. *J. Comput. Phys.* **229** (5), 1739–1762.
- KINGMA, D.P. & BA, J. 2014 Adam: a method for stochastic optimization. arXiv: 1412.6980.
- KOIKE, S., UENO, M., NAKAKITA, K. & HASHIMOTO, A. 2016 Unsteady pressure measurement of transonic buffet on NASA common research model. *AIAA paper*, 2016-4044.
- LAGEMANN, E., BRUNTON, S.L., SCHRÖDER, W. & LAGEMANN, C. 2024 Towards extending the aircraft flight envelope by mitigating transonic airfoil buffet. *Nat. Commun.* **15** (1), 5020.
- LARSSON, J., KAWAI, S., BODART, J. & BERMEJO-MORENO, I. 2016 Large eddy simulation with modeled wall-stress: recent progress and future directions. *Mech. Engng Rev.* **3** (1), 15–00418.
- LECUN, Y., BOTTOU, L., BENGIO, Y. & HAFNER, P. 1998 Gradient-based learning applied to document recognition. *Proc. IEEE* **86** (11), 2278–2324.
- LEE, B.H.K. 1990 Oscillatory shock motion caused by transonic shock boundary-layer interaction. *AIAA J.* **28** (5), 942–944.
- LEE, B.H.K. 2001 Self-sustained shock oscillations on airfoils at transonic speeds. *Prog. Aerosp. Sci.* **37** (2), 147–196.
- LELE, S.K. 1992 Compact finite difference schemes with spectral-like resolution. *J. Comput. Phys.* **103** (1), 16–42.
- LENORMAND, E., SAGAUT, P. & TA PHUOC, L. 2000 Large eddy simulation of subsonic and supersonic channel flow at moderate Reynolds number. *Intl J. Numer. Meth. Fluids* **32** (4), 369–406.
- LINOT, A.J. & GRAHAM, M.D. 2023 Dynamics of a data-driven low-dimensional model of turbulent minimal Couette flow. *J. Fluid Mech.* **973**, A42.
- LINOT, A.J., ZENG, K. & GRAHAM, M.D. 2023 Turbulence control in plane Couette flow using low-dimensional neural ODE-based models and deep reinforcement learning. *Intl J. Heat Fluid Flow* **101**, 109139.
- LIU, Z., BECKERS, D. & ELDREDGE, J.D. 2025 Model-based reinforcement learning for control of strongly-disturbed unsteady aerodynamic flows. *AIAA J.* **63** (10), 4105–4125.
- LUMLEY, J.L. 1967 The structure of inhomogeneous turbulent flows. In *Atmospheric Turbulence and Radio Wave Propagation*, (ed. A.M. Yaglom & V.I. Tatarski). Nauka.
- LUSHER, D.J., SANSICA, A. & HASHIMOTO, A. 2024 Effect of tripping and domain width on transonic buffet on periodic NASA-CRM airfoils. *AIAA J.* **62** (11), 4411–4430.
- LUSHER, D.J., SANSICA, A. & HASHIMOTO, A. 2025 Implicit large eddy simulations of three-dimensional turbulent transonic buffet on wide-span infinite wings. *J. Fluid Mech.* **1007**, A26.
- MANOHAR, K., BRUNTON, B.W., KUTZ, J.N. & BRUNTON, S.L. 2018 Data-driven sparse sensor placement for reconstruction: demonstrating the benefits of exploiting known patterns. *IEEE Control Syst. Mag.* **38** (3), 63–86.

- MASINI, L., TIMME, S. & PEACE, A.J. 2020 Analysis of a civil aircraft wing transonic shock buffet experiment. *J. Fluid Mech.* **884**, A1.
- MENEVEAU, C. & KATZ, J. 2000 Scale-invariance and turbulence models for large-eddy simulation. *Annu. Rev. Fluid Mech.* **32** (1), 1–32.
- MO, Y., TRAVERSO, T. & MAGRI, L. 2024 Decoder decomposition for the analysis of the latent space of nonlinear autoencoders with wind-tunnel experimental data. *Data-Centric Engng* **5**, e38.
- MORIMOTO, M., FUKAMI, K., ZHANG, K. & FUKAGATA, K. 2022 Generalization techniques of neural networks for fluid flow estimation. *Neural Comput. Applics.* **34** (5), 3647–3669.
- MORIMOTO, M., FUKAMI, K., ZHANG, K., NAIR, A.G. & FUKAGATA, K. 2021 Convolutional neural networks for fluid flow analysis: toward effective metamodeling and low dimensionalization. *Theor. Comput. Fluid Dyn.* **35** (5), 633–658.
- MOUSAVI, H. & ELDREDGE, J.D. 2025 Low-order flow reconstruction and uncertainty quantification in disturbed aerodynamics using sparse pressure measurements. *J. Fluid Mech.* **1013**, A41.
- MURATA, T., FUKAMI, K. & FUKAGATA, K. 2020 Nonlinear mode decomposition with convolutional neural networks for fluid dynamics. *J. Fluid Mech.* **882**, A13.
- NAIR, S.S., KUNZ, R.F., ZHANG, W. & YANG, X.I.A. 2025 Rough-wall modeling on a low-dimensional manifold. *Phys. Rev. Fluids* **10** (6), 064606.
- OHMICHI, Y., ISHIDA, T. & HASHIMOTO, A. 2018 Modal decomposition analysis of three-dimensional transonic buffet phenomenon on a swept wing. *AIAA J.* **56** (10), 3938–3950.
- OJA, E. 1982 Simplified neuron model as a principal component analyzer. *J. Math. Biol.* **15**, 267–273.
- OMATA, N. & SHIRAYAMA, S. 2019 A novel method of low-dimensional representation for temporal behavior of flow fields using deep autoencoder. *AIP Adv.* **9** (1), 015006.
- PAGE, J., HOLEY, J., BRENNER, M.P. & KERSWELL, R.R. 2024 Exact coherent structures in two-dimensional turbulence identified with convolutional autoencoders. *J. Fluid Mech.* **991**, A10.
- PALADINI, E., BENEDDINE, S., DANDOIS, J., SIPP, D. & ROBINET, J.-C. 2019 Transonic buffet instability: from two-dimensional airfoils to three-dimensional swept wings. *Phys. Rev. Fluids* **4** (10), 103906.
- PLANTE, F., DANDOIS, J. & LAURENDEAU, É. 2020 Similarities between cellular patterns occurring in transonic buffet and subsonic stall. *AIAA J.* **58** (1), 71–84.
- POPLINGER, L., RAVEH, D.E. & DOWELL, E.H. 2019 Modal analysis of transonic shock buffet on 2D airfoil. *AIAA J.* **57** (7), 2851–2866.
- PRECHELT, L. 1998 Automatic early stopping using cross validation: quantifying the criteria. *Neural Networks* **11** (4), 761–767.
- RUMELHART, D.E., HINTON, G.E. & WILLIAMS, R.J. 1986 Learning representations by back-propagation errors. *Nature* **322**, 533–536.
- SANSICA, A. & HASHIMOTO, A. 2023 Global stability analysis of full-aircraft transonic buffet at flight Reynolds numbers. *AIAA J.* **61** (10), 4437–4455.
- SANSICA, A., LOISEAU, J.-C., KANAMORI, M., HASHIMOTO, A. & ROBINET, J.-C. 2022 System identification of two-dimensional transonic buffet. *AIAA J.* **60** (5), 3090–3106.
- SARTOR, F., METTOT, C. & SIPP, D. 2015 Stability, receptivity, and sensitivity analyses of buffeting transonic flow over a profile. *AIAA J.* **53** (7), 1980–1993.
- SASHIDA, H., AOYAMA, T., KAWAI, S. & KAWAI, S. 2024 Wall-modeled LES of pitching airfoil at high Reynolds number transonic flows. *AIAA paper*, 2024-0299.
- SCHAUERTE, C.J. & SCHREYER, A.-M. 2023 Experimental analysis of transonic buffet conditions on a two-dimensional supercritical airfoil. *AIAA J.* **61** (8), 3432–3448.
- SMITH, L., FUKAMI, K., SEDKY, G., JONES, A. & TAIRA, K. 2024 A cyclic perspective on transient gust encounters through the lens of persistent homology. *J. Fluid Mech.* **980**, A18.
- SOLERA-RICO, A., SANMIGUEL VILA, C., GÓMEZ-LÓPEZ, M., WANG, Y., ALMASHJARY, A., DAWSON, S.T.M. & VINUESA, R. 2024 β -variational autoencoders and transformers for reduced-order modelling of fluid flows. *Nat. Commun.* **15** (1), 1361.
- SUGIOKA, Y., KOIKE, S., NAKAKITA, K., NUMATA, D., NONOMURA, T. & ASAI, K. 2018 Experimental analysis of transonic buffet on a 3D swept wing using fast-response pressure-sensitive paint. *Exp. Fluids* **59**, 108.
- SUGIOKA, Y., KOUCHI, T. & KOIKE, S. 2022 Experimental comparison of shock buffet on unswept and 10-deg swept wings. *Exp. Fluids* **63** (8), 132.
- SUGIOKA, Y., NAKAKITA, K., KOIKE, S., NAKAJIMA, T., NONOMURA, T. & ASAI, K. 2021 Characteristic unsteady pressure field on a civil aircraft wing related to the onset of transonic buffet. *Exp. Fluids* **62**, 1–18.
- TAMAKI, Y. & KAWAI, S. 2024 Wall-modeled large-eddy simulation of transonic buffet over NASA-CRM using FFWHC-ACE. *AIAA J.* **62** (7), 2489–2504.

- TIJDEMAN, H. & SEEBASS, R. 1980 Transonic flow past oscillating airfoils. *Annu. Rev. Fluid Mech.* **12** (1), 181–222.
- TIMME, S. 2020 Global instability of wing shock-buffet onset. *J. Fluid Mech.* **885**, A37.
- TRAN, J., FUKAMI, K., INADA, K., UMEHARA, D., ONO, Y., OGAWA, K. & TAIRA, K. 2024 Aerodynamics-guided machine learning for design optimization of electric vehicles. *Commun. Engng* **3** (1), 174.
- YOUSIF, M.Z., YU, L. & LIM, H. 2022 Physics-guided deep learning for generating turbulent inflow conditions. *J. Fluid Mech.* **936**, A21.