

## MIT Open Access Articles

### *Prediction of Intermittent Fluctuations from Surface Pressure Measurements on a Turbulent Airfoil*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

**Citation:** Rudy, Samuel H. and Sapsis, Themistoklis P. 2022. "Prediction of Intermittent Fluctuations from Surface Pressure Measurements on a Turbulent Airfoil." AIAA Journal, 60 (7).

**As Published:** 10.2514/1.j061163

**Publisher:** American Institute of Aeronautics and Astronautic

**Persistent URL:** <https://hdl.handle.net/1721.1/154220>

**Version:** Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

**Terms of use:** Creative Commons Attribution-Noncommercial-ShareAlike



# Prediction of Intermittent Fluctuations from Surface Pressure Measurements on a Turbulent Airfoil

Samuel H. Rudy\* and Themistoklis Sapsis

Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139

## Abstract

This work studies the effectiveness of several machine learning techniques for predicting extreme events occurring in the flow around an airfoil at low Reynolds. For certain Reynolds numbers the aerodynamic forces exhibit intermittent fluctuations caused by changes in the behavior of vortices in the airfoil wake. Such events are prototypical of the unsteady behavior observed in airfoils at low Reynolds and their prediction is extremely challenging due to their intermittency and the chaotic nature of the flow. We seek to forecast these fluctuations in advance of their occurrence by a specified length of time. We assume knowledge only of the pressure at a discrete set of points on the surface of the airfoil, as well as offline knowledge of the state of the flow. Methods include direct prediction from historical pressure measurements, flow reconstruction followed by forward integration using a full order solver, and data-driven dynamic models in various low dimensional quantities. Methods are compared using several criteria tailored for extreme event prediction. We show that methods using data-driven models of low order dynamic variables outperform those without dynamic models and that unlike previous works, low dimensional initializations do not accurately predict observables with extreme events such as drag.

*Keywords*– Extreme events, unsteady aerodynamics, machine learning

## 1 Nomenclature

$\mathbf{u}, p$	= Fluid velocity and pressure
$P$	= Airfoil surface pressure time series
$C_d, C_l$	= Drag and lift coefficients
$q$	= Smoothed drag coefficient
$\omega$	= Extreme event rate
$\tau$	= Load time for prediction of $q$
$m$	= Mass matrix from spectral element grid
$w$	= Weights for proper orthogonal decomposition and flow reconstruction
$\mathcal{C}$	= chord-length
$\mathcal{L}_\bullet$	= Loss function used to train neural network •
$dt$	= Time step
$Re$	= Reynolds number
$u_\infty$	= inlet velocity
$\Phi\Sigma\Psi^T$	= components of proper orthogonal decomposition of flow field
$\xi$	= Latent space representation of flow field in full-field neural network
$r$	= Rank of low dimensional representations of flow
$G$	= Network for estimating proper orthogonal decomposition time series
$E$	= Encoder network for flow reconstruction

\* Corresponding author (shrudy@mit.edu).

**Software:** [https://github.com/snagcliffs/Airfoil\\_EE](https://github.com/snagcliffs/Airfoil_EE).

$D$	=	Decoder network for flow reconstruction
$F_{p/\psi/\xi}$	=	Non-dynamic models for mapping $P/\psi/\xi$ to $q$
$H_{p/\psi/\xi}$	=	Dynamic models for variables $P$ , $\psi$ , and $\xi$
$S, R, F_1, \alpha$	=	Precision, recall, $F_1$ score, and area under S-R curve
$F_{1,opt}$	=	Optimal test set $F_1$ score
$\alpha^*$	=	Maximum adjusted area under precision-recall curve

## 2 Introduction

Extreme events are common features in engineering and scientific disciplines including climate, ocean engineering, and fluid structure interaction that are characterized by observables of a dynamical system exhibiting heavy tails [40]. The outlier events populating these tails are of particular interest due to their effects on aerodynamics and fatigue, or other potentially adverse consequences. However, the rarity and intermittency of such events also makes their prediction challenging. There has been significant recent interest in sampling strategies [33, 39, 2], optimization schemes [3], and tailored loss functions [18, 36] for the prediction of extreme events. A common goal of many of the past works on extreme events and of the present work is the prediction of extreme events in advance of their occurrence.

The focus of this work is on the two-dimensional incompressible flow around an airfoil at low ( $\mathcal{O}(10^4)$ ) Reynolds number. Dynamics of the flow around the airfoil at this Reynolds regime are highly nontrivial [26] and have been shown to be characteristically different from those at higher Reynolds [22]. Previous works using both experimental and computational tools have found that slow moving airfoils exhibit a large range of wake behaviors, with qualitative changes in the nature of the flow occurring with small changes in angle of attack and Reynolds number [30, 17, 47]. Similar unstable behavior has been observed in the flow around a cylinder, in the so called transitional regime between ordered and disordered behavior [49] as well as in vortex-induced-vibrations of flexible cylinders [32]. Due to these instabilities, as well as an apparent lack of fidelity between computational and experimental results, it has been suggested that construction of rigid winged slow flying vehicles may be challenging if not impossible [43].

Despite apparent challenges, there has long been considerable interest in the study of low Reynolds airfoils [26]. In particular, recent works have explored the use of machine learning to estimate flow field and aerodynamic data from sensors on the surface of the airfoil. These include methods for flow reconstruction from limited sensors using neural networks [29, 15], filtering based flow estimation [8, 6], as well as prediction of aerodynamic coefficients [7]. Deep learning has also been used for estimating properties of the flow used in low order vortex models. In [21, 24], authors use neural network based methods to estimate the leading edge suction parameter(LESP). The model studied in this work has a Reynolds number of 17,500, substantially higher than other computational works which focus on Reynolds number in the range of  $\mathcal{O}(10^2) - \mathcal{O}(10^3)$  [24, 8, 6, 29]. This is closer to the lower end of the range considered by experimental work [7, 15]. Unlike some other works [7, 15, 21, 24], this work does not study the effects of pitching motions or disturbances in the incident velocity. We instead focus on prediction of intermittent fluctuations in the aerodynamic coefficients of a stationary airfoil, which have not been the focus of previous works.

This work applies several machine learning techniques to predict fluctuations in the drag coefficient of an airfoil in the transitional regime where we observe chaotic and intermittent behavior. It follows a broader trend of the application of tools from machine learning to problems in fluid dynamics. For a broader view, the interested reader may re-

fer to a number of recent articles outlining and discussing the role of machine learning in fluid dynamics; [5] provides an excellent overview of many of the applications machine learning has seen in fluids, [13] provides an assessment of several common supervised learning methods applied to flow reconstruction, super-resolution, and coefficient estimation, and [4] provides a discussion of the role of machine learning in fluids, as well as some pitfalls and concerns.

The paper is organized as follows: In section 3 we formulate the problem of predicting aerodynamic fluctuations. Section 4 describes the methodology used in this work, including flow reconstruction methods in section 4.1, forecasting methods for aerodynamic fluctuations in section 4.2 and a discussion of performance metrics in section 4.3. Results are shown in section 5 with comparisons between all methods. In section 6 we offer closing thoughts, pitfalls, and potential future directions based on this paper.

### 3 Problem description

We consider a NACA 4412 airfoil at chord-length based Reynolds number of  $Re = 17500$  and  $5^\circ$  angle of attack. Flow around the airfoil is simulated using the spectral element code Nek5000 [35] according to the incompressible Navier-Stokes equations given by,

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u} \\ \nabla \cdot \mathbf{u} &= 0. \end{aligned} \quad (1)$$

The computational grid uses 4368 elements with spectral order 7. Statistics of observables relevant to this work were found to be in agreement with those from a shorter simulation using a more resolved grid having 14144 elements. A convective boundary condition is used for the outflow [9]. The spectral element grid without Gauss-Lobatto Legendre interpolation points and a snapshot of the vorticity are shown in the top row of Figure 1. Further details of the numerical simulation are provided in Appendix A and software for reproducing data used in this work is available online.

Pressure recordings along the surface of the airfoil  $P(t)$  are taken at a discrete set of 50 points around the airfoil at intervals of  $dt = 0.01$  throughout the simulation. These locations are shown in the bottom panel of Fig. 1. Aerodynamic force is computed using the pressure and skin friction. This decomposes into the streamwise ( $x$ ) and cross-flow ( $y$ ) directions, defined by:

$$\vec{F}(t) = \oint \boldsymbol{\tau}(t) - p(t)\mathbf{n} ds = D(t)\mathbf{e}_x + L(t)\mathbf{e}_y \quad (2)$$

where  $\boldsymbol{\tau}$ ,  $p$ , and  $\mathbf{n}$  are the skin shear stress, pressure, and wall normal vector, and the integral is taken over the airfoil surface. Forces are then used to compute the non-dimensional drag-coefficient  $C_d$  and lift coefficient  $C_l$ , defined as:

$$C_d(t) = \frac{2D(t)}{\rho u_\infty^2 C}, \quad C_l(t) = \frac{2L(t)}{\rho u_\infty^2 C}, \quad (3)$$

where choord length  $C = 1$ , density  $\rho = 1$ , and free-stream velocity  $u_\infty = 1$ .

The two dimensional simulation yields a quasi-stable behavior in which intermittent fluctuations are observed in the aerodynamic coefficients, shown in the first two rows of Figure 2 alongside the density functions of their absolute deviations. We note that the density functions clearly exhibit the expected "heavy tails" associated with observables of




Figure 1: Left: Domain of computational problems showing outlines of spectral elements without internal interpolation points. Right: Snapshot of vorticity. Bottom: Sensor placement showing indexing from 0-25 along top and 25-49 along bottom of airfoil. Axes on bottom figure not drawn to scale.

dynamical systems with extreme events [41]. This regime of intermittent fluctuations is persistent for the entirety of the simulation used in this work, but with altered conditions may exhibit mode switching to a state with more regular oscillations. Further details of this case are given in Appendix C. In this work we focus solely on the regime where intermittent fluctuations are observed.

The goal of the present work is to predict these intermittent fluctuations in the aerodynamic coefficients using information regarding the surface pressure of the airfoil in advance of their occurrence by some lead time  $\tau$ . To focus on non-periodic behavior, we consider predictions on a smoothed time series derived from the drag coefficient. Specifically,

$$q(t) = \left( C_d(s) * \mathcal{N} \left( s \left| 0, \frac{1}{(2f_{peak})^2} \right. \right) \right) (t), \quad (4)$$

where  $f_{peak} \approx 1.44$  is the peak frequency of the drag coefficient. In practice, the convolution in Eq. (4) is computed using a compactly supported kernel having width  $3/f_{peak}$ . The time series  $q(t)$  is also normalized to be mean zero and unit variance. The quantity  $q(t)$  captures the non-periodic behavior of the drag coefficient, in particular maintaining extreme events and the heavy tailed deviation. The goal of this work is thus concisely stated as learning data-driven models for the prediction problem,

$$P(s \leq t) \rightarrow q(t + \tau), \quad (5)$$

at various  $\tau \geq 0$ .

## 4 Methods

In this section we formulate several data-driven models for Eq. (5). We assume knowledge of measurements of pressure at a discrete set of points along the surface of the airfoil up




Figure 2: Left: Aerodynamic coefficients  $C_d$  and  $C_l$ , and smoothed drag coefficient  $q$  as defined in Eq. (4). Red dashed lines indicate  $\pm 2$  standard deviations away from the mean. Right: Histograms of the deviation of each quantity showing typical heavy tails of systems with extreme events.

to the current time  $t$  and seek to predict the value of the extreme event indicator at time  $t + \tau$ . We denote the time series for pressure measurements up to the current time  $t$  as  $P_t = P(s \leq t)$ . The minimal problem is thus learning a function directly from historical pressure measurements to the  $q(t + \tau)$ , which may be done using a variety of standard machine learning tools.




Figure 3: Schematic of methods for EE prediction showing various flows of information from assumed knowledge of historical pressure measurements  $P_t$  to future value of quantity of interest  $q(t + \tau)$ . Dynamic model for pressure measurements has been omitted.

A baseline predictor for Eq. (5) is found using a standard recurrent neural network, such as a long-short-term-memory networks (LSTM) [19], to interpolate a function directly from  $P_t$  to  $q(t + \tau)$ . Alternatively, we may try to improve forecasts of  $q(t)$  through the offline use of flow field data. Previous works [3] have used modal representations of a flow field combined with adjoint equations to learn precursor states to extreme events. In this work we discuss two methods for compressing flow field data, sensing expansion coefficients in the compressed basis, and exploiting this knowledge for potentially improved

prediction of  $q(t+\tau)$ : the proper orthogonal decomposition and neural network based flow reconstruction. In each case we study the predictive capability of initializing a flow solver with the reduced order initial condition, predicting directly from historical representations in the reduced space, and neural network based reduced order models. We also consider data-driven dynamic models for the pressure measurements.

The methods considered in this work are summarized in Figure 3. Methods are separated into an offline compression stage, a sensing stage where we infer the reduced order state from point pressure measurements, reconstruction of the flow field from the reduced order state, and finally forecast of the quantity of interest,  $q(t)$ . Learned functions  $G$ ,  $E$ , and  $D$  are predictors of the POD mode time series, latent space encoder, and flow reconstruction network, respectively. The letter  $F$  has been used to denote LSTM predictions that do not use a dynamic model, with subscript indicating input.  $H$  is used to denote neural network based reduced order models. In each case, we use time series for pressure at discrete points on the surface of the airfoil as the starting point of the online prediction.

Details on each method are provided in the following sections. We assume familiarity with common implementations of neural networks and stochastic optimization. In particular, the work in this manuscript makes frequent use of deep LSTMs [19] and the Adam method for optimization [23]. The unfamiliar reader may find an excellent reference in the free online textbook [16]. Further details on the structure and implementation of neural networks is given in Appendix B.

#### 4.1 Flow field compression and reconstruction

We begin with a discussion of the two offline methods for flow reconstruction: the proper orthogonal decomposition (POD) and an LSTM based encoder-decoder pair. In this subsection we provide an overview of the methodology used to form the reduced rank representation for each of these two cases as well as methods for approximating time series associated with each POD mode.

##### 4.1.1 Proper orthogonal decomposition and sensing

The proper orthogonal decomposition is a standard tool for decomposing a flow field into spatial modes that are orthogonal with respect to a given inner-product and whose time evolution are also orthogonal [20, 42]. We apply the POD to the velocity field around the airfoil using a weighted inner product. Specifically, the POD finds matrices  $\Phi$ ,  $\Sigma$ , and  $\Psi$  which are discretized solutions to,

$$\mathbf{u}(\mathbf{x}, t) - \bar{\mathbf{u}}(\mathbf{x}) = \Phi(\mathbf{x})\Sigma\Psi(t)^T, \quad (6)$$

where,

$$\begin{aligned} \Sigma &= \text{diag}(\sigma_1, \dots, \sigma_r), \quad \sigma_i \geq \sigma_{i+1} \geq 0 \\ \langle \phi_i, \phi_j \rangle_w &= \int_{\Omega} \phi_i(\mathbf{x})^T \phi_j(\mathbf{x}) w(\mathbf{x}) dx = \delta_{i,j} \\ \langle \psi_i, \psi_j \rangle_w &= \int_0^T \psi_i(t) \psi_j(t) dt = \delta_{i,j}, \end{aligned} \quad (7)$$

and  $w(\mathbf{x}) > 0$  is a weight used to focus the inner product, and thus variation explained by POD, on regions of interest near the surface of the airfoil. A similar weighted approach was used in [3] where a wall focused POD was used as a basis for predicting extreme

dissipation events in channel flow. We use the family of weights given by,

$$w(\mathbf{x}) = \frac{1 - \epsilon}{1 + e^{(d(\mathbf{x})-l)/\delta}} + \epsilon \quad (8)$$

where  $d(\mathbf{x})$  is the distance from  $\mathbf{x}$  to the surface of the airfoil. Equation (8) describes a smooth sigmoidal curve that decays from 1 (assuming  $l \gg \delta$ ) at the airfoil surface to  $\epsilon$  in the far field. For small  $\delta$ , this transition is localized around  $d = l$  and the weights are approximately 1 for  $d < l - \delta$  and  $\epsilon$  for  $d > l + \delta$ . We use parameter values  $l = 1$ ,  $\delta = 0.1$ , and  $\epsilon = 0.1$ . The POD is therefore principally focused on describing variation in the velocity field within one chord length of the airfoil surface, with approximately one tenth the weighting for variation outside this region.

In the case where  $w(\mathbf{x}) = 1$  the POD is equivalent to the singular value decomposition of the mean subtracted data, also known as principal component analysis. For non-identity weights,  $\Sigma$  and  $\Psi(t_j)$ ,  $j = 1, \dots, m$  are given by the eigenvalue decomposition of  $\mathbf{U}^T \mathbf{W}_m \mathbf{U}$  where  $\mathbf{U} \in \mathbb{R}^{n \times m}$  is the mean-subtracted velocity data and  $\mathbf{W}_m$  is a diagonal matrix with  $w(\mathbf{x}_i)m(\mathbf{x}_i)$  along the diagonal where  $m(\mathbf{x}_i)$  is the mass associated with that grid point for the spectral element grid [20]. Modes  $\Phi$  are subsequently computed using their definition in Eq. (6). Alternative methods may compute  $\Phi$  before  $\Sigma$  and  $\Psi$ , but these suffer from numerical issues for  $\epsilon \ll 1$ .




Figure 4: The mean velocity field and three POD modes. Axes not drawn to scale.

Applying POD to the airfoil data yields modes  $\phi_i$ , shown in Fig. 4, singular values  $\Sigma$ , and time series  $\psi_i(t)$  corresponding to each mode  $\phi_i$ . In the online phase of any prediction method, we will only have access to  $P_t$ , not  $\psi_i(t)$ . The latter may be estimated from sparse or gappy measurements [46]. We therefore train a deep LSTM model to estimate the current POD representation of the flow from pressure measurements. Letting  $\psi_r(t)$  be the rank  $r$  truncation of the time series of the POD expansion we have,

$$\hat{\psi}_r(t) = G(P_t). \quad (9)$$

The exact form of and training procedure for  $G$  is described in greater detail in Appendix B. True values of POD coefficients as well as their estimates from pressure via Eq. (9) are shown in Fig. 5. Note that reconstructed time series are filled in on a much denser grid than true values since they are computed from the finely sampled pressure time series. Temporal resolution on  $\psi(t)$  is limited by the number of output files saved during numerical simulation as well as memory limitations in the computation of the POD.

It is worth noting that the time series  $\psi(t)$  are normalized to unit variance and  $\Sigma$  is not considered in loss function. Hence, error in higher modes is treated the same as error in lower modes. Higher modes did still have higher error, perhaps because they tended to exhibit more chaotic behavior. The authors did not explore the loss function exhaustively since doing so would be a significant research endeavor on its own.



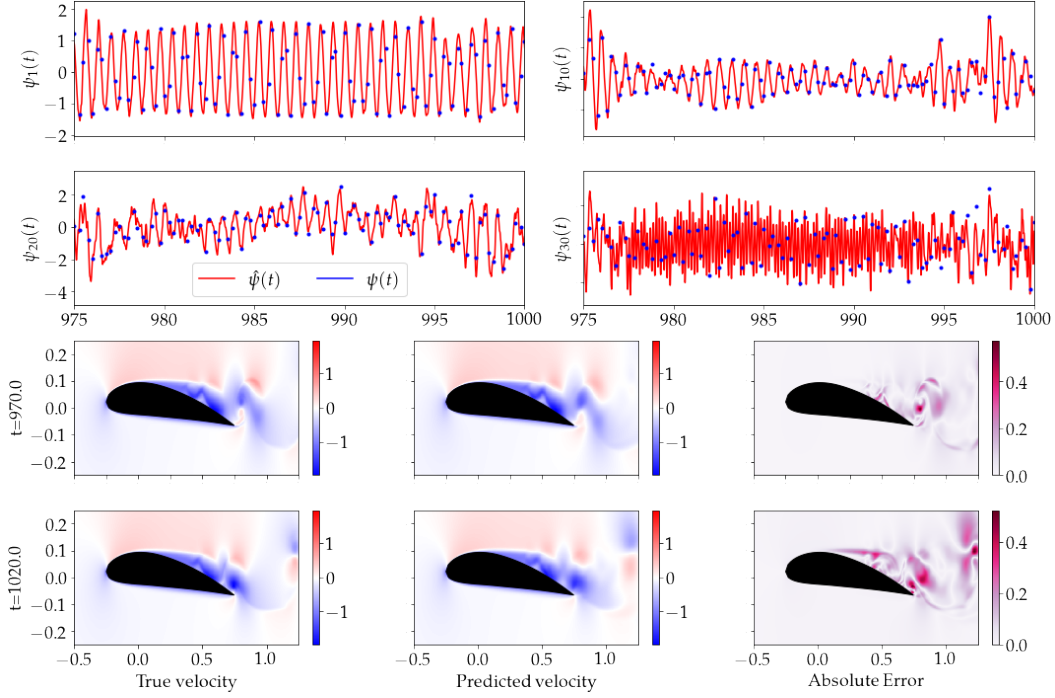


Figure 5: Flow field reconstruction using LSTM to predict time series for first 32 POD modes using pressure measurements,  $P_t$ . Top [Operator  $G(P_t)$ ]: LSTM prediction of POD time series for 1<sup>st</sup>, 10<sup>th</sup>, 20<sup>th</sup>, and 30<sup>th</sup> modes (red) and sparser true values (blue). Bottom [Operator  $\Phi\Sigma G$ ]: POD reconstruction of streamwise velocity compared to two true fields within test set data.

After prediction of  $\hat{\psi}(t)$  using  $G$  one may reconstruct an approximation of the full flow field using Eq. (6). The lower two rows on Fig. 5 show the results of this in the streamwise direction as well as the true velocity and absolute error for two snapshots taken from the testing dataset, i.e. snapshots not seen by the optimization algorithm used to learn  $G$ . These reconstructions suffer from multiple sources of error. Expanding the difference between true and reconstructed fields, we get,

$$\begin{aligned} \mathbf{u}(t) - \hat{\mathbf{u}}_{POD}(t) &= \Phi\Sigma\psi(t)^T - \Phi_r\Sigma_r G(P_t)^T \\ &= \underbrace{\Phi_{-r}\Sigma_{-r}\psi_{-r}(t)^T}_{\text{Unresolved}} - \Phi_r\Sigma_r \underbrace{(\psi_r(t) - G(P_t))^T}_{\text{NN Error}}, \end{aligned} \quad (10)$$

where the subscript  $-r$  denotes those components not in the first  $r$ . Error in the neural network increases for modes with high frequency and chaotic behavior, such as mode 30 shown in Fig. 5. However, significant error is also incurred from unresolved modes due to the slow decay of singular values  $\Sigma$ . While spatially periodic translational behavior may be represented with pairs of modes (see Eq. (34) in [42]) the system studied in the present work tends to shed isolated pairs of vortices. Moreover, it also exhibits intermittency, which may be difficult to capture in a POD basis. The authors are not aware of efficient linear methods for representing translation of sparse structures such as the wake vortices observed in this data.

#### 4.1.2 Neural network flow reconstruction

In light of some of the deficiencies of the POD and motivated by the successful applications of neural networks to problems in fluid dynamics [31, 27, 37, 5, 13] we also consider neural network based approaches to flow reconstruction. The approach in this work is to use an LSTM-based encoder coupled with a fully connected network predicting the velocity at each grid point. The fully connected decoder resembles that used in [10] and [29] though do to the large computational grid we have not used the probabilistic formulation considered in the later. We note however, that neither of the aforementioned papers included the use of history terms in their prediction of the fluid state, as we will show is done by the LSTM encoder for the networks considered in this work. For brevity, we will call networks of this type full field neural networks (FFNN), indicating that the output of the network is the values of fluid velocity at each grid point used by a solver.

As in the POD case, we represent the fluid velocity around an airfoil at time  $t$  using a low dimensional representation,  $\boldsymbol{\xi}(t) \in \mathbb{R}^r$ . The time history of the pressure sensors is encoded to the state using an LSTM given by,

$$\boldsymbol{\xi}(t) = E(P_t). \quad (11)$$

Since the full state of the flow field is encoded in  $\boldsymbol{\xi}(t)$ , one may tune the dimension of  $\boldsymbol{\xi}$  to acquire a desired rank for the reduced order representation of  $\mathbf{u}$ . We found improvements in reconstruction accuracy up to approximately rank  $r = 32$ , with minimal improvement at high values. We therefore use  $r = 32$  for the remainder of this work unless noted otherwise. The reconstructed velocity field is then given by,

$$\hat{\mathbf{u}}(\mathbf{x}_i, t_m) = D(\boldsymbol{\xi})_i, \quad (12)$$

where  $D$  is a standard fully connected neural network with the final layer being linear. Taken together, Eq. (11) and (12) form a recurrent network from pressure sensor time histories to the fluid velocity given by,

$$\hat{\mathbf{u}}(x_i, t_j) = D \circ E(P(t_m))_i, \quad (13)$$

which are trained together using numerical simulation data. Further details on the network structure may be found in Appendix B.

We train the networks using a weighted square-error loss function designed to favor accurate reconstruction near the wall boundary;

$$\mathcal{L}_{E,D} = \sum_j \int_{\Omega} \|u(\mathbf{x}, t_j) - \hat{u}(\mathbf{x}, \boldsymbol{\xi})\|^2 w(\mathbf{x}) dx, \quad (14)$$

where  $\Omega$  is the computational domain and  $w(\mathbf{x})$  is as defined in Sec. 4.1.1 using  $\epsilon = \delta = 0.1$  and  $l = 1$ . It is also possible to weigh the loss function to more heavily penalize errors in times immediately preceding an extreme event, but doing so was observed to have little effect on prediction performance. In practice, the integral in Eq. (14) is approximated using the mass matrix  $m$  obtained from the spectral element grid. The expression simplifies to a simple weighted sum of squares error given by,

$$\mathcal{L}_{E,D} \approx \sum_j \sum_i \|\mathbf{u}(\mathbf{x}_i, t_j) - \hat{\mathbf{u}}(\mathbf{x}_i, \boldsymbol{\xi})\|^2 w(\mathbf{x}_i) m(\mathbf{x}_i). \quad (15)$$

Snapshots of the true streamwise velocity, the FFNN predicted streamwise velocity, and absolute error are shown in Fig. 6, which includes the same fields shown in Fig. 5 to illustrate POD reconstruction. Note that the maximum absolute error using neural networks




Figure 6: [Operator  $D \circ E(P_t)$ ] Reconstruction of streamwise velocity ( $u_\infty = 1$  subtracted) using FFNN for flow compression with 32 dimensional latent space and LSTM to map the pressure measurements,  $P_t$ , to the latent variables,  $\xi$ . Note maximum error is approximately half that of POD, and regions of high error are significantly more localized.

is considerably lower than that of POD and that regions of high error are more localized. This clearly indicates the superior performance of FFNN to POD for flow reconstruction in this particular case.

The choice to use a very large output layer of the neural network makes predictions specific to the particular grid used in training, though interpolation schemes could be used in other cases. The FFNN decoder's size also makes it highly memory intensive, which limits batch size in training. It is plausible that similar networks for three dimensional flows would require coordinate descent like algorithms where fractions of the output weights are updated on any given batch. This is in contrast to operator type networks, where spatial coordinate  $x$  is given as input [27]. Methods based on the latter were implemented without physical constraints for flow reconstruction from pressure measurements, but were found to underperform the full-field neural networks discussed in this work. We note that this could be in part due to using a neural encoder of pressure measurements rather than sparse function evaluations, as used in [27]. It is also possible that the use of physics-informed methods could improve prediction accuracy of the operator type networks and this is noted as a potentially interested research question. However, we consider such an approach to be outside of the scope of this work.

## 4.2 Forecasting aerodynamic fluctuations

We now consider online methods for Eq. (5). We separate these methods into two broad categories; those that do not use dynamic models, and those that use dynamic models such as the Navier Stokes equations or data-driven dynamic models.

### 4.2.1 Non-dynamic methods

We first consider methods that do not employ any sort of dynamic model. These are simply interpolations a function from the time history of an input quantity to the future value of  $q(t)$ . The general form is given by,

$$q(t + \tau) = F_\bullet(\bullet_t), \quad (16)$$

where  $\bullet \in [P, \psi, \xi]$  and  $F$  is a deep LSTM mapping some time series of historical measurements to the future quantity of interest. In the case where the reduced representation of the flow state is used we have:

$$\begin{aligned} q(t + \tau) &= F_\psi \circ G(P_t) \\ q(t + \tau) &= F_\xi \circ E(P_t), \end{aligned} \tag{17}$$

so the POD sensing network  $G$  and neural network encoder  $E$  may be considered as feature maps for the forecasting networks  $F_{\psi/\xi}$ . Each of  $F_{\psi/\xi}$  are trained using the mean square error loss function. Other loss functions more specific to extreme events were considered but found to make negligible difference to the resulting trained network.

Non-dynamic methods of the form given by Eq. (16) may be favorable for several reasons. They are simple, easy to train, and based on the ubiquitous and highly effective LSTM network structure. They may be of particular interest due to the computational savings offered by avoiding dynamic models. Approximation of  $q(t + \tau)$  may be rapidly computed from the state of the LSTM, which is updated online from streaming data. However there are also downsides to the non-dynamic approach. As  $\tau$  becomes large, we are approximating larger steps of a chaotic time series. The problem thus becomes very sensitive to inputs, and balancing sufficient model complexity with overfitting becomes challenging. Moreover, prediction of  $q(t + \tau)$  for any given lead time  $\tau$  requires its own trained network.

#### 4.2.2 Full order dynamical model based prediction methods

Previous work predicting extreme events for turbulent flows has used a low dimensional representation of the flow as initial condition for a fluid simulation and employed the adjoint to evaluate the gradient of the future quantity of interest with respect to coefficients in the low dimensional expansion [3]. This is possible in cases where the low order initial condition is sufficiently close to the full order state to track its behavior. Here we evaluate if this is the case for low dimensional initialization from both the POD basis and FFNN. Numerical simulations are initialized using either the FFNN based reconstructed velocity or the low dimensional reconstruction using the known POD modes and singular values, along with the estimated temporal coefficients  $\hat{\psi}(t) = G(P_t)$ . Future values of  $q(t)$  are computed using the same convolution as in Eq. (4);

$$\begin{aligned} \hat{q}(t + \tau) &= \left( \hat{C}_d(s; \hat{\mathbf{u}}) * \mathcal{N} \left( s \mid 0, \frac{1}{(2f_{peak})^2} \right) \right) (t + \tau) \\ \text{where, } \hat{\mathbf{u}}(t) &= D(\xi) \text{ or } \hat{\mathbf{u}}(t) = \Phi \Sigma G(\xi), \end{aligned} \tag{18}$$

where  $\hat{\mathbf{u}}(s)$ ,  $s \in [t, t + \tau]$  is found via numerical integration of the Navier-Stokes equations. Since the time series for  $\hat{C}_d$  is much shorter than the full simulation used for training neural networks, the common practice of zero padding distorts the values of  $\hat{q}$  near the endpoints. We instead truncate and re-normalize the smoothing kernel. For valid comparison, the same is done for the true drag coefficient truncated to the interval  $[t, t + \tau]$ . Following Eq. (18), time series for  $q$  and  $\hat{q}$  are normalized using the mean and variance of the full length time series defined in Eq. (4).

#### 4.2.3 Data-driven dynamical model based prediction methods

We also consider forecasts of either i) the pressure signal,  $P(t)$ , or ii) the reduced order state of the fluid flow, using data-driven dynamic models. In the later case we formulate and evaluate reduced order dynamical models for both the reduced POD state and the latent

representation learned by the FFNN. For the POD ROM, we train the dynamic model on the estimated POD coefficients ( $\hat{\psi} = G(P_t)$ ), instead of the true coefficients  $\psi$ , since the former may be found exactly from airfoil surface pressure. This eliminates error due to imprecise initial conditions. Learning a data-driven model of the pressure signal along the surface of the airfoil has, to the best of the authors knowledge, not been used in previous works.

Forecasting models in each case are LSTM networks mapping historical measurements of a given quantity ( $P/\hat{\psi}/\xi$ ) to its value  $\kappa$  timesteps in the future. They are represented as,

$$P(t_{j+\kappa}) = H_P(P(t_j), P(t_{j-\kappa}), \dots), \quad (19)$$

with similar networks for  $\hat{\psi}$  and  $\xi$ . Here  $\kappa$  is taken to be three in order to alleviate some of the numerical difficulties with training data-driven dynamic models in the small timestep limit [25]. Thus, the LSTM maps historical measurements of the dynamic quantity  $P$  to the value of that quantity three steps in the future:  $P(t_{j+3})$ . Since  $P$  is measured every  $dt = 0.01$  time units, the LSTM is effectively a dynamic model for  $P$  with timestep equal to  $\kappa dt = 0.03$ . Networks are trained using the mean square error loss over a prediction window of 20 steps;

$$\mathcal{L}_{H_P} = \sum_i \sum_{j=1}^{20} \|H_P^j(P_i) - P(t_{i+\kappa j})\|^2 \quad (20)$$

where we define composition of  $H_P$  with itself by,

$$\begin{aligned} H_P^1(P_i) &= H_P(P(t_i), P(t_{i-\kappa}), P(t_{i-2\kappa}), \dots), \\ H_P^2(P_i) &= H_P(H_P^1(P_i), P(t_i), P(t_{i-\kappa}), \dots), \\ H_P^3(P_i) &= H_P(H_P^2(P_i), H_P^1(P_i), P(t_i), \dots) \dots \end{aligned} \quad (21)$$

The sum over index  $i$  in Eq. (20) is taken over all initial times in the training dataset. The same network structure and loss function is used for the full set of 50 pressure measurements as well as the 32 dimensional reduced order representations of the flow field. Layer sizes are scaled to account for the difference in dimension between  $P$  and  $\psi/\xi$ .

For any  $\tau \geq 0$ , we obtain the estimated forecast,

$$\hat{P}(t + \tau) = H_P^{m_\tau}(P_t), \quad (22)$$

where  $m_\tau = \tau/(\kappa dt)$ . Following the LSTM based forecast of  $P/\hat{\psi}/\xi$  we may use  $F_{P/\hat{\psi}/\xi}$  trained for zero lead time to evaluate  $\hat{q}(t + \tau) = F_P \circ H_P^{m_\tau}(P_t)$  and likewise for  $\hat{\psi}/\xi$ . Since prediction of  $q$  with zero lead time is a much simpler problem, we use standard feed-forward neural networks in place of the LSTMs used for non-dynamic predictions with non-zero lead time.

### 4.3 Error metrics

Each of the neural networks in this work is trained using the mean square error (MSE). The MSE is differentiable and may be evaluated on subsets of the training data, allowing for the use of stochastic optimization schemes run on a graphics processing unit. However, the MSE may not be a good indicator of success in predicting extreme events. We have therefore adapted several extreme-event-tailored error metrics to compare the various methods considered in this work. Specifically, results are compared using the batch relative entropy

loss, the maximum adjusted area under the precision-recall curve, and the extreme event rate dependent area under the precision recall curve, and optimal  $F_1$  score.

The batch relative entropy loss (BRE) is inspired by the work in [36], where authors use a relative entropy loss function to train convolutional neural networks that are capable of making accurate prediction of a system governed by the truncated Korteweg-de Vries (tKdV) equation in regimes with extreme events. Their work uses empirical partition functions similar to the soft-max activation commonly used in neural networks to transform high-dimensional predictions into probability distribution functions highlighting outlier values. Loss is subsequently measured using the KL-divergence. This approach was shown to significantly improve prediction accuracy over the MSE on the tKdV problem. A similar approach was adapted for the present work using partition functions over mini-batches rather than output dimensions. Specifically, we define the batch relative entropy loss as,

$$BRE = \sum_i z_i \log \left( \frac{z_i}{\hat{z}_i} \right) \quad (23)$$

where  $z_i, \hat{z}_i$  are given by the empirical partition functions,

$$z_i = \frac{e^{q(t_i)}}{\sum_j e^{q(t_j)}}, \quad \hat{z}_i = \frac{e^{\hat{q}(t_i)}}{\sum_j e^{\hat{q}(t_j)}}, \quad (24)$$

and where the sum is taken over a mini-batch. We note that this work considers partition functions that weigh positive outliers, as opposed to the symmetric variant used in [36], since all events of interest to this work skew positive. The BRE loss was tested as a means of training the neural networks described in previous sections, but taken over mini-batches was found to perform comparably with the mean square error. For evaluation, we use the batch relative entropy loss taken over the full testing dataset.

The performance of a predictor of extreme events may also be measured by the ability of that predictor to classify events based on a threshold value of the quantity of interest. The maximum adjusted area under the precision-recall curve ( $\alpha^*$ ) was proposed in [18] as a loss function for predicting extreme events and shown to perform well for predicting extreme dissipation events in Kolmogorov flow and rouge waves in the Majda-McLaughlin-Tabak model. To define  $\alpha^*$  we first introduce the quantity  $\omega$  for the extreme event rate and the corresponding threshold  $\hat{a}(\omega)$ , such that  $p(q > \hat{a}(\omega)) = \omega$ . Introducing a second threshold  $\hat{b}$  for the prediction  $\hat{q}$  yields a classifier for which we can compute the precision ( $S = \text{true positives} / \text{predicted positives}$ ) and recall ( $R = \text{true positives} / \text{total positives}$ ). Noting that precision is uniquely determined by the extreme event rate and recall, [18] computes the area under the precision recall curve,

$$\alpha(\omega) = \int_0^1 S(R, \omega) dR = \int_{\mathbb{R}} S(\hat{b}, \omega) \left| \frac{\partial R(\hat{b}, \omega)}{\partial \hat{b}} \right| d\hat{b}, \quad (25)$$

and the maximum adjusted area under the precision-recall curve as

$$\alpha^* = \max_{\omega} \alpha(\omega) - \omega. \quad (26)$$

We compute the integral in Eq. (25) using a finite grid of  $\hat{b}$  values spread evenly between  $\min(\hat{q})$  and  $\max(\hat{q})$  with finite difference scheme to evaluate the derivative term. It's value is bounded between 0 and 1 but occasionally falls slightly outside this range due to numerical issues and is subsequently clipped. The maximization in Eq. (26) is taken over a discrete set of samples evenly spaced  $\omega$  between 0.01 and 0.25.

Finally, we consider the extreme event rate dependent optimal  $F_1$  score, defined as the  $F_1$  score on the testing dataset using the threshold that optimizes the  $F_1$  score on the training and validation datasets. That is,

$$F_{1,opt}(q, \hat{q}, \omega) = F_1 \left( (q_{test} > \hat{a}(\omega)), (\hat{q}_{test} > \hat{b}_{opt}) \right) \quad (27)$$

where,

$$\hat{b}_{opt} = \underset{\hat{b}}{argmax} F_1 \left( (q_{train/val} > \hat{a}(\omega)), (\hat{q}_{train/val} > \hat{b}) \right). \quad (28)$$

Taken together this yields five metrics of predictor performance. Mean square error, batch relative entropy, and  $\alpha^*$  are independent of extreme event rate and yield simple scalar metrics of performance, though do not indicate performance at a particular extreme event rate. The extreme event rate dependent  $\alpha(\omega)$  and  $F_{1,opt}(\omega)$  each seek to measure a balance between precision and recall at variable extreme event rates. The dependency on extreme event rate allows for a more descriptive quantification of error, since we can compare methods for a variety of extreme event rates.

## 5 Results

In this section we present results for each of the methods for Eq. (5). Results for prediction using the full order model with reduced order initial condition are kept distinct from those using neural networks. This is due to the weak performance by the former, as well as its considerable computational expense, which limits the number of trials we use to evaluate it. For predictions of  $q(t)$  using dynamic methods, we apply a smoother to remove higher frequency oscillations. This smoother weights previous predictions with exponentially decreasing weights and does not use any future prediction beyond  $\tau$ . Examples of full test set predictions for  $\hat{q}(t)$  using each of the six neural network based methods are shown in Figure 11 for lead time  $\tau = 7.0$  and in Appendix D for other lead times.

### 5.1 Simulations with reduced order initial conditions

We first present results showing the failure of approaches taken by the authors to forecast aerodynamic fluctuations using the Navier-Stokes solver and reduced order initial condition. For each of a variety of ranks, reduced order initial conditions were formed using the estimated POD time series  $\hat{\psi}(t)$  and FFNN flow reconstruction  $D \circ E(P_t)$  at 50 evenly spaced times throughout the portion on data reserved for testing. Examples of the resulting smoothed drag coefficient are shown in Figure 7, which compares predicted time series for  $\hat{q}(t)$  using several ranks of POD reconstructions (left column) as well as FFNN based reconstructions (right column) to the true value  $q(t)$  (black curves). Each row represents a different initial condition from within the testing dataset. We note that all the initializations using different reduction methods/orders exhibit poor agreement with the true time series in at least 1 of the three cases shown.

The mean absolute error in the drag coefficient and smoothed drag coefficient  $q$  for each rank are shown in Figure 8. There are several noteworthy features; the approximated values of  $C_d$  for both methods are initially fairly accurate, falling within 10% of the true values, but very quickly diverge. By the end of the 10.5 time unit simulations, the error in each quantity is roughly as large as the standard deviation of the true data. Initial error in  $q(t)$  is higher, as is expected since this quantity includes some information from future estimates of  $C_d$ . We also note the lack of noticeable correlation between the rank of the latent




Figure 7: Examples of forecasts of  $q(t + \tau)$  using Eq. (18) with various ranks of POD (left column) and FFNN (right column). Examples show prediction from initial conditions are spaced evenly throughout the testing dataset. True value is shown as solid black line. Each row represents a different realization of the flow.




Figure 8: Mean absolute error, normalized by standard deviations of true time series for the drag coefficient and  $q(t)$  using various ranks of reduced order initial conditions from the estimated POD time series and FFNN.

space used for flow reconstruction and prediction accuracy. In the POD case, it appears that the lowest rank ( $r = 8$ ) performs poorly compared to higher ranks, but  $r = 64$  is not uniformly better than  $r = 16$  or  $r = 32$ . A plausible explanation for this is that as the rank is increased, so too is the difficulty of the sensing problem outlined at the end of Sec. 4.1.1. Indeed, using the true (non-reconstructed)  $\psi(t)$  we see improved agreement with full order results if  $r$  is let to become large. The connection between rank and accuracy for the FFNN examples is more opaque. While higher rank initial conditions are slightly more accurate at predicting the initial drag, the difference is small and decays quickly as the simulation progresses.

While the results here are not sufficiently exhaustive to preclude the use of a full order model for predicting extreme events in the flow around an airfoil they at least provide strong evidence of its difficulty. We include them to show the sensitivity of the problem to small changes in the initial condition (note that the FFNN was able to accurately reconstruct the fluid velocity) and to contrast to other examples in fluids where a representation in low rank basis was found to be an effective precursor to extreme events [3].



## 5.2 Comparison of reduced order dynamic methods

We now consider the data-driven dynamic models for the pressure signals and reduced order representations of the flow field. Figure 9 shows the time evolution of the mean square error in the dynamic quantity and absolute error in prediction of  $q$  for each of the three methods  $H_{P/\psi/\xi}$ . The dynamic quantities in each case have been normalized to unit variance so direct comparison of error magnitudes across methods and across indices are meaningful. Error statistics have been binned for each timestep and normalized such that vertical slices of any subplot are density functions across all examples from the testing dataset. Apparent recurring features in the error plots for  $q$  are likely due to the methods missing the same feature from multiple closely sampled initial conditions.

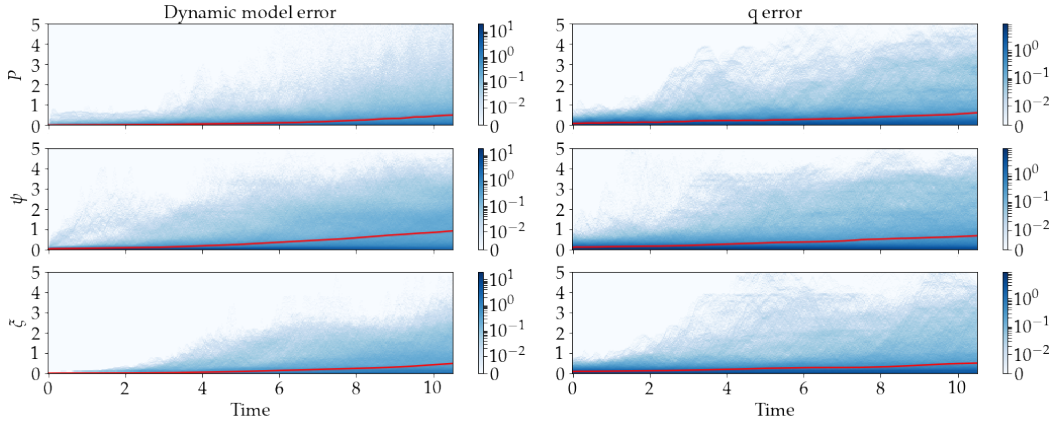


Figure 9: Test set empirical probability density functions of the mean square error of the LSTM forecast (left) and absolute error in predicting  $q(t+\tau)$  for 50 dimensional data-driven forecast of pressure measurements (top), and 32 dimensional reduced order models using POD modes (middle) and FFNN latent space (bottom). Mean error quantities are shown by red curves.

In each of the three models, the bulk of the error remains low throughout the forecast window, as is shown by the curves for mean error. Error statistics for the dynamic model for  $P$  remain low for approximately three time units before some trajectories diverge. Outlier errors for the dynamic model for  $\psi$  grow more rapidly than others initially and by the end of the 10.5 convective unit interval are noticeably larger, on average than the other two. The dynamic model for  $\xi$  clearly has both the lowest error in the dynamic variable and has error for  $q$  of a similar magnitude to the dynamic model in  $P$ .

For the models in  $P$  and  $\psi$  where indices in the dynamic variable carry meaning is also instructive to see what variables accumulate error at what rates. Figure 10 shows the mean square error for each index in each dynamic variable taken across test set examples. The sharp discontinuity in the error for pressure sensors is a consequence of their position on the airfoil, shown in the lower panel of Fig. 1. The first 26 pressure sensors are equally placed on the top of the airfoil from front to tail including endpoints and the subsequent 24 on evenly spaced from front to tail on the lower side. We see that error towards the tail of the airfoil on both the suction and pressure sides is initially lower than towards the front. However, as time progresses, there is considerable error in pressure sensors towards the rear of the suction edge of the airfoil.

As expected error in the low order POD modes is lower than in the high order, less




Figure 10: Time evolution of mean (across test set samples) square error for each of 50 pressure sensors (left), 32 POD modes (middle) and dimensions of FFNN latent space (right) using LSTM based dynamic models.

energetic, modes. This may be explained by the fact that these low order modes tend to track large scale features of the flow and have slower and less chaotic trajectories. It is worth noting that, while the left hand column of Fig. 9 makes clear that error in  $\psi$  grows much more rapidly in the initial forecast than  $P$  or  $\xi$ , the same is not as obvious for low  $\tau$  prediction of  $q(t + \tau)$  using the dynamic model for  $\hat{\psi}$ . This suggests that the value of  $q$  may largely be a function of the low order energetic POD modes, which Fig. 10 shows are accurately tracked for longer lead times. Indices for the FFNN latent space variable  $\xi$  do not have meaning and there is no correlation between index and the rate at which error increases.

### 5.3 Forecasting aerodynamic fluctuations

In this section we compare the six neural network based prediction strategies using the metrics outlined in Sec. 4.3. Examples of predictions for lead time  $\tau = 7.0$  using all methods are shown in Fig. 11. Blue curves indicate true values  $q(t)$  and red dashed curves show predictions with lead time  $\tau = 7.0$  using the three non-dynamic and three dynamics methods. Plots for other lead times are shown in Appendix D.




Figure 11: True (blue) and predicted (red) time series for  $q(t)$  with lead time  $\tau = 7$ .

Scalar values (MSE,  $\alpha^*$ , BRE) of prediction accuracy for each method and lead times  $\tau$  ranging from 0 to 10.5 convective time units in intervals of  $0.7 \approx 1/f_{peak}$  are shown in Fig. 12. Metrics for prediction via non-dynamic methods are shown as solid lines and those for dynamic models are shown as dashed lines.

Figure 12 clearly indicates some qualitative differences between dynamic and non-dynamic methods. In particular, error statistics for the dynamic models at a particular




Figure 12: Mean square error, maximum adjusted area under the precision-recall curve ( $\alpha^*$ ), and test set relative entropy for each of the six neural network based prediction methods. Evaluation is performed on partition of data reserved for method comparison (test set).

value of  $\tau$  are highly correlated with those for similar lead times. The curves indicating their performance as a function of lead time are therefore smooth. The same is not true for non-dynamic models, where predictions for each lead time are performed via their own trained neural network. It is plausible that more care could be taken to yield consistency across lead times for non-dynamic models but methods for doing so are not immediately apparent.

The scalar error metrics also clearly show that the dynamic models in  $P$  and  $\xi$ ,  $H_P$  and  $H_\xi$  are the most accurate of the six models. Differences between the other four are less pronounced, though  $H_\psi$  does, on average, slightly outperform non-dynamic methods. This difference is most notable in the MSE, though still apparent for the extreme event specific metrics. This suggests that a non-negligible portion of the improvement  $H_\psi$  has over non-dynamic methods manifests in the non-extreme values of  $q$ , though the scalar valued metrics do not resolve this feature.

Values for the remaining extreme event rate dependent evaluation metrics are shown in Fig. 13 for values of  $\omega$  between 0.01 and 0.25. The top row shows values for  $F_{1,opt}$  and bottom row shows the area under the precision-recall curve. Within each subplot, values towards the bottom of the image indicate the prediction accuracy of the model for classifying very rare events (1%) while those at the very top show more common events, with lead time increasing across the horizontal axis. While differences in values of prediction accuracy using Fig. 13 may not be immediately clear, the same qualitative features seen in Fig. 12 are again apparent. In particular, errors for the dynamic methods are much smoother in time than those of the non-dynamic methods.

Differences between methods become much more apparent when plotted directly. Figure 14 shows the differences in values between all methods for both  $F_{1,opt}$  (left) and area under the precision-recall curve,  $\alpha$  (right). Values within each subfigure indicate the metric evaluated with the method on the corresponding row minus the metric for the method on the corresponding column. Hence, blue (negative) indicates the the method corresponding to that column performed better while red (positive) indicates the method corresponding to the row is better.

Figure 14 shows  $H_P$  and  $H_\xi$  outperforming other methods across all lead times  $\tau$  and extreme event rates  $\omega$  except for a small number of datapoints clustered around  $\tau = 4$  and low  $\omega$ . Comparisons between other methods are less easily summarized. There does not appear to be discernible trends in the comparison between non-dynamic methods using  $P$ ,  $\psi$  or  $\xi$  as input. From this it seems reasonable to surmise that the accuracy of non-




Figure 13: Extreme event rate dependent area under the precision-recall curve ( $\alpha(\omega)$  and  $F_{1,opt}$ ) for each method evaluated on test set.

dynamic methods is not improved in a meaningful manner by exploiting flow field data. The dynamic model for  $\psi$  outperforms non-dynamic methods for mid-range lead times and lower extreme event rates, but under perform in prediction of less rare events (higher  $\omega$ ) at longer lead times. Differences between methods are largely consistent between  $F_{1,opt}$  and  $\alpha$ .

## 6 Discussion

We have investigated the feasibility of a variety of methods for forecasting aerodynamic fluctuations occurring in a simulation of two dimensional incompressible flow around a NACA 4412 airfoil using only pressure recordings along the surface of the airfoil. Some of these methods also exploited, in an offline manner, knowledge of the flow field to construct low dimensional representations that were either used as inputs to machine learned predictors or variables in dynamic models. Comparison between the methods considered in this work suggests that the complexity of the flow, though low Reynolds and only two dimensional, precludes the use of a full order computational fluid dynamic model with low dimensional input. Other methods were largely comparable, except for the learned dynamic models of the pressure measurements and of the latent variable of the FFNN based flow reconstruction, which performed notably better than others.

The authors highlight that neural networks allow for tremendous freedom with regards to structure, hyper-parameters, regularization, and other factors that may affect performance. We have tested several architectures and selected the best result for this paper, but acknowledge that our work falls considerably short of testing across all sizes and training procedures. Doing so would require an extensive computational resources and results would have no guarantee of generality beyond the problem specifically discussed in this work. Rather, we sought to investigate the use of various classes of models (non-dynamic, ROM, full order simulation) and representations of data (POD modes, FFNN) for the purpose of forecasting a quantity exhibiting extreme events. We also note that the work contained in this manuscript has been done without the inclusion of artificial noise, as is common in works applying machine learning to synthetic datasets. The authors suspect that including noise would not change the results in a meaningful way, since flow reconstruction methods have been shown to be robust to noise [29, 10] and quantities with




Figure 14: Comparison of  $\alpha$  (left) and  $F_{1,opt}$  (right) across various methods, lead times  $\tau$ , and extreme event rates  $\omega$ . Each plot shows  $\alpha/F_{1,opt}$  for method corresponding to row minus  $\alpha/F_{1,opt}$  for method corresponding to column. Blue (negative) indicates column-method outperforms row-method. Hence lower left plot indicates  $H_\zeta$  has higher  $\alpha$  (favorable performance) than  $H_\psi$  in most cases.

dynamic models may be estimated with filtering. Nonetheless, further work studying the effect of noise would be necessary to confirm this.

In Section 5 we suggested that low order initializations of the Navier Stokes solver failed to accurately predict fluctuations, this motivated the use of data-driven dynamic models. This is not surprising, given the complexity and non-periodic nature of the flow. It is possible that a mixed strategy of projecting the governing Navier-Stokes equations onto POD modes coupled with an LSTM closure model would outperform both approaches, as suggested in [25]. This approach, however would be highly non-trivial when using the latent space of the FFNN, since it is unclear how one might project the known governing equations.

We also stress the nature of this work as a computational study, rather than one that may be directly applied to engineering. Numerical simulation for the current work was implemented with a blockage ratio of 3.18%. This falls within the range that might be expected for experiments on bluff bodies [50, 48] but is lower than some works studying the aerodynamic properties of airfoils. Several works have noted dependence of aerodynamic properties including the drag coefficient, Strouhal frequency certain critical Reynolds numbers on the blockage ratio [45]. This effect was observed in the case used for the present work, but not explored in depth. We also note that the simulation is of two dimensional flow over a smooth airfoil, which exhibits qualitative differences from three dimensional simulations and experiments [43]. Thus, the present work should be considered as a study of a computational model of prototypical flow phenomena, rather than experimental or application conditions.

The present work considered the case where pressure measurements are taken at 50 positions around the perimeter of the airfoil. This is, of course, not practical and future work could consider the use of various sensor placement techniques for determining opti-

mal placement on the airfoil with respect to forecasting fluctuations [24, 8].

We believe this work establishes compelling comparisons and baselines for the prediction of extreme events using measurements on the surface of an airfoil. Numerous techniques, including several adapted from other works on extreme events were tested and compared. The results suggest that the use of data-driven dynamic models for quantities subsequently used to predict extreme events outperform those that ignore dynamics. This work also provides numerous opportunities for future research. In particular, sensor placement and uncertainty quantification are critical elements of practical engineering that may be explored in the context developed in this work.

## Acknowledgments

This material is based upon work supported by the Army Research Office (Grant No. W911NF-17-1-0306) and used the Extreme Science and Engineering Discovery Environment (XSEDE) [44], which is supported by National Science Foundation under Grant No. ACI-1548562. TPS has also been supported by the Air Force Office of Scientific Research (Grant No. MURI FA9550-21-1-0058). Simulations were performed on Expanse at the San Diego Supercomputing Center through allocation TG-MTH210003.

## References

- [1] Martín Abadi and et. al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] Antoine Blanchard and Themistoklis Sapsis. Bayesian optimization with output-weighted optimal sampling. *Journal of Computational Physics*, 425:109901, 2021.
- [3] Patrick J Blonigan, Mohammad Farazmand, and Themistoklis P Sapsis. Are extreme dissipation events predictable in turbulent fluid flows? *Physical Review Fluids*, 4(4):044606, 2019.
- [4] MP Brenner, JD Eldredge, and JB Freund. Perspective on machine learning for advancing fluid mechanics. *Physical Review Fluids*, 4(10):100501, 2019.
- [5] Steven L Brunton, Bernd R Noack, and Petros Koumoutsakos. Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics*, 52:477–508, 2020.
- [6] Andre FC da Silva and Tim Colonius. Flow state estimation in the presence of discretization errors. *Journal of Fluid Mechanics*, 890, 2020.
- [7] Scott T Dawson, Nicole K Schiavone, Clarence W Rowley, and David R Williams. A data-driven modeling framework for predicting forces and pressures on a rapidly pitching airfoil. In *45th AIAA fluid dynamics conference*, page 2767, 2015.
- [8] Andre Fernando de Castro da Silva and Tim Colonius. An enkf-based flow state estimator for aerodynamic flows. In *8th AIAA Theoretical Fluid Mechanics Conference*, page 3483, 2017.
- [9] Suchuan Dong. A convective-like energy-stable open boundary condition for simulations of incompressible flows. *Journal of Computational Physics*, 302:300–328, 2015.
- [10] N Benjamin Erichson, Lionel Mathelin, Zhewei Yao, Steven L Brunton, Michael W Mahoney, and J Nathan Kutz. Shallow neural networks for fluid flow reconstruction with limited sensors. *Proceedings of the Royal Society A*, 476(2238):20200097, 2020.
- [11] Paul Fischer and Julia Mullen. Filter-based stabilization of spectral element methods. *Comptes Rendus de l’Académie des Sciences-Series I-Mathematics*, 332(3):265–270, 2001.
- [12] PF Fischer. Implementation considerations for the oifs/characteristics approach to convection problems. *Argonne National Laboratory*, 2003.
- [13] Kai Fukami, Koji Fukagata, and Kunihiko Taira. Assessment of supervised machine learning methods for fluid flows. *Theoretical and Computational Fluid Dynamics*, 34(4):497–519, 2020.

- [14] Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre-and post-processing facilities. *International journal for numerical methods in engineering*, 79(11):1309–1331, 2009.
- [15] Daniel F Gomez, Francis Lagor, Phillip B Kirk, Andrew Lind, Anya R Jones, and Derek A Paley. Unsteady dmd-based flow field estimation from embedded pressure sensors in an actuated airfoil. In *AIAA Scitech 2019 Forum*, page 0346, 2019.
- [16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [17] Muralikrishnan Gopalakrishnan Meena, Kunihiko Taira, and Keisuke Asai. Airfoil-wake modification with gurney flap at low reynolds number. *AIAA journal*, 56(4):1348–1359, 2017.
- [18] Stephen Guth and Themistoklis P Sapsis. Machine learning predictors of extreme events occurring in complex dynamical systems. *Entropy*, 21(10):925, 2019.
- [19] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [20] Philip Holmes, John L Lumley, Gahl Berkooz, and Clarence W Rowley. *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge university press, 2012.
- [21] Wei Hou, Darwin Darakananda, and Jeff D Eldredge. Machine-learning-based detection of aerodynamic disturbances using surface pressure measurements. *AIAA Journal*, 57(12):5079–5093, 2019.
- [22] Dong-Ha Kim, Jo-Won Chang, and Joon Chung. Low-reynolds-number effect on aerodynamic characteristics of a naca 0012 airfoil. *Journal of aircraft*, 48(4):1212–1215, 2011.
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [24] Mathieu Le Provost, Wei Hou, and Jeff Eldredge. Deep learning and data assimilation approaches to sensor reduction in estimation of disturbed separated flows. In *AIAA Scitech 2020 Forum*, page 0799, 2020.
- [25] Matthew E Levine and Andrew M Stuart. A framework for machine learning of model error in dynamical systems. *arXiv preprint arXiv:2107.06658*, 2021.
- [26] PBS Lissaman. Low-reynolds-number airfoils. *Annual review of fluid mechanics*, 15(1):223–239, 1983.
- [27] Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- [28] Yvon Maday, Anthony T Patera, and Einar M Rønquist. An operator-integration-factor splitting method for time-dependent problems: application to incompressible fluid flow. *Journal of Scientific Computing*, 5(4):263–292, 1990.
- [29] Romit Maulik, Kai Fukami, Nesar Ramachandra, Koji Fukagata, and Kunihiko Taira. Probabilistic neural networks for fluid flow surrogate modeling and data recovery. *Physical Review Fluids*, 5(10):104401, 2020.
- [30] Karthik Menon and Rajat Mittal. Aerodynamic characteristics of canonical airfoils at low reynolds numbers. *AIAA Journal*, 58(2):977–980, 2020.
- [31] Michele Milano and Petros Koumoutsakos. Neural network modeling for near wall turbulent flow. *Journal of Computational Physics*, 182(1):1–26, 2002.
- [32] Y Modarres-Sadeghi, F Chasparis, MS Triantafyllou, M Tognarelli, and P Beynet. Chaotic response is a generic feature of vortex-induced vibrations of flexible risers. *Journal of Sound and Vibration*, 330(11):2565–2579, 2011.
- [33] Mustafa A Mohamad and Themistoklis P Sapsis. Sequential sampling strategy for extreme event statistics in nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 115(44):11138–11143, 2018.
- [34] Anthony T Patera. A spectral element method for fluid dynamics: laminar flow in a channel expansion. *Journal of computational Physics*, 54(3):468–488, 1984.
- [35] James W. Lottes Paul F. Fischer and Stefan G. Kerkemeier. nek5000 Web page, 2008. <http://nek5000.mcs.anl.gov>.
- [36] Di Qi and Andrew J Majda. Using machine learning to predict extreme events in complex systems. *Proceedings of the National Academy of Sciences*, 117(1):52–59, 2020.

- [37] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [38] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. *arXiv preprint arXiv:1710.05941*, 2017.
- [39] Themistoklis P Sapsis. Output-weighted optimal sampling for bayesian regression and rare event statistics using few samples. *Proceedings of the Royal Society A*, 476(2234):20190834, 2020.
- [40] Themistoklis P Sapsis. Statistics of extreme events in fluid flows and waves. *Annual Review of Fluid Mechanics*, 53, 2020.
- [41] Themistoklis P Sapsis. Statistics of extreme events in fluid flows and waves. *Annual Review of Fluid Mechanics*, 53:85–111, 2021.
- [42] Kunihiro Taira, Steven L Brunton, Scott TM Dawson, Clarence W Rowley, Tim Colonius, Beverley J McKeon, Oliver T Schmidt, Stanislav Gordeyev, Vassilios Theofilis, and Lawrence S Ukeiley. Modal analysis of fluid flows: An overview. *Aiaa Journal*, 55(12):4013–4041, 2017.
- [43] J Tank, Lelanie Smith, and GR Spedding. On the possibility (or lack thereof) of agreement between experiment and computation of flows over wings at moderate reynolds number. *Interface focus*, 7(1):20160076, 2017.
- [44] John Towns, Timothy Cockerill, Maytal Dahan, Ian Foster, Kelly Gaither, Andrew Grimshaw, Victor Hazlewood, Scott Lathrop, Dave Lifka, Gregory D Peterson, et al. Xsede: accelerating scientific discovery. *Computing in science & engineering*, 16(5):62–74, 2014.
- [45] S Turki, H Abbassi, and S Ben Nasrallah. Effect of the blockage ratio on the flow in a channel with a built-in square cylinder. *Computational Mechanics*, 33(1):22–29, 2003.
- [46] Daniele Venturi and George Em Karniadakis. Gappy data and reconstruction procedures for flow past a cylinder. *Journal of Fluid Mechanics*, 519:315–336, 2004.
- [47] S Wang, Y Zhou, Md Mahbub Alam, and Hongxing Yang. Turbulent intensity and reynolds number effects on an airfoil at low reynolds numbers. *Physics of Fluids*, 26(11):115107, 2014.
- [48] GS West and CJ Apelt. The effects of tunnel blockage and aspect ratio on the mean flow past a circular cylinder with reynolds numbers between 104 and 105. *Journal of Fluid mechanics*, 114:361–377, 1982.
- [49] Charles HK Williamson. Vortex dynamics in the cylinder wake. *Annual review of fluid mechanics*, 28(1):477–539, 1996.
- [50] Serhiy Yarusevych and Michael SH Boutilier. Vortex shedding of an airfoil at low reynolds numbers. *AIAA journal*, 49(10):2221–2227, 2011.



## Appendix A: Numerical simulation details

Numerical simulations of the airfoil used in this work were performed using the open-source spectral element solver Nek5000 [35] and run on Expanse at the San Diego Supercomputing Center using a grant through the National science Foundation funded Extreme Science and Engineering Discovery Environment (XSEDE) [44]. The spectral element method, proposed in [34], partitions the computational domain into non-overlapping elements and using polynomial interpolation within each element to represent the solution. Mesh was generated using gmsh [14].

Time integration was performed using a second order semi-implicit scheme described in [12]. The operator-integrating-factor splitting method proposed in [28] was used to allow for stable time integration with a fixed timestep of 0.001. The filtering method proposed in [11] was also used to stabilize the simulation.

The simulation was initialized with a velocity of  $\mathbf{u} = (1, 0)$ . This led to a short transient which was avoided when training neural networks and taking the proper orthogonal decomposition. The inflow boundary condition was held at a fixed value of  $\mathbf{u} = (1, 0)$  throughout the simulation. Boundary conditions on the wall of the airfoil were no slip (Dirichlet), along the top and bottom on the domain were symmetric. The outflow boundary condition was of the convective type proposed in [9]. This allowed for the passage of strong vortices out of the domain while avoiding numerical issues.

Simulation for 1020 convective units took approximately 3.5 hours running on 128 cores or approximately 5 days running on 16 cores on the author’s local computer. Aerodynamic quantities are saved every timestep and pressure at discrete points on the surface of the airfoil is saved every 10 steps, or 0.01 time units. Velocity and pressure data on the full domain are saved every 0.25 time units. Code for recreating datasets, as well as files for time series of pressure and aerodynamic coefficients are available online at [https://github.com/snagcliffs/Airfoil\\_EE](https://github.com/snagcliffs/Airfoil_EE).

## Appendix B: Neural network structures and implementation

All neural networks used in this work were implemented in Python using the TensorFlow library [1]. The swish activation function [38] was used between fully connected non-recurrent layers and the sigmoid function was used for in LSTM layers. In section 4.2.3 we note that prediction of  $q$  following LSTM dynamic models are performed using non-recurrent networks. This step, which does not forecast any quantity, is a fairly trivial problem and details of networks have been omitted here but are available in the online code repository.

*Structure:*

Structures for neural networks used in this work are given in Eq. (29). We use the notation (input shape)  $\rightarrow$  (layer type, layer sizes) with multiple integer layer sizes indicating repeated layers. Hence, the first line tells us  $F_p$  uses 140 history points of dimension 50 as input, then a 32 dimensional fully connected layer, two LSTM layers of size 32, and finally a series of fully connected layers having sizes 32, 16, 8, 4, and 1. The final layer of size 1 is the output layer. For networks including a latent representation of the dynamics, the quantity  $r$  has been used as a stand in for rank. This work includes examples having  $r \in \{8, 16, 32, 64\}$ . The output size of  $E$  is the number of interpolation points in the spectral element grid  $n = 279, 552$ , multiplied by the dimension  $d = 2$ , for a total size of  $n \cdot d = 559, 104$ . In

all cases using history terms, a stride of 3 was used between inputs. This was found to yield a slight performance advantage over using a stride of 1. Connections in the decoder network, excluding those to the final large layer we modeled after residual type networks, having a linear connection added to those feeding into nonlinear activation functions. A larger number of history points were used for  $F_p$  to account for the history terms used to compute  $\hat{\psi}$  and  $\xi$ , though this was not found to significantly affect results.

$$\begin{aligned}
F_p &: (50 \times 140) \rightarrow (FC, 32) \rightarrow (LSTM, 32, 32) \rightarrow (FC, 32, 16, 8, 4, 1) \\
F_\psi &: (r \times 70) \rightarrow (FC, 32) \rightarrow (LSTM, 32, 32) \rightarrow (FC, 32, 16, 8, 4, 1) \\
F_\xi &: (r \times 70) \rightarrow (FC, 32) \rightarrow (LSTM, 32, 32) \rightarrow (FC, 32, 16, 8, 4, 1) \\
G &: (50 \times 70) \rightarrow (FC, 64) \rightarrow (LSTM, 128) \rightarrow (FC, 64, r) \\
E &: (50 \times 70) \rightarrow (FC, 64) \rightarrow (LSTM, 64) \rightarrow (FC, 64, r) \\
D &: (r \times 1) \rightarrow (FC, 64, 128, 256, n \cdot d) \\
H_p &: (50 \times 70) \rightarrow (FC, 50) \rightarrow (LSTM, 100, 100) \rightarrow (FC, 100, 50, 50) \\
H_{\psi/\xi} &: (r \times 70) \rightarrow (FC, 32) \rightarrow (LSTM, 64, 64) \rightarrow (FC, 64, 32, 32)
\end{aligned} \tag{29}$$

As noted in section 6, designing neural networks allows tremendous freedom on seemingly arbitrary choices such as layer sizes and activation functions. When testing different networks for this work we found significant differences between network architectures (LSTM, fully connected, branch-trunk, etc.). Small changes such as slightly altering layer sizes or activation functions did not, in general, significantly affect results.

#### *Data partitions:*

In many machine learning tasks and particularly those with a high dimensional parameter space such as neural networks, a dataset is split into distinct sets for training, validation, and testing. The training dataset is used to update model parameters according to the gradient of the cost function. The validation set is used to prevent overfitting through the use of early stopping when performance metrics on validation data have stalled. Finally, the testing set is reserved for a comparison between models. We use a (70/15/15%) split, meaning 70 percent of the available data in each case is used for training, 15 for validation, and 15 for testing. Training and validation datasets are mixed randomly from within an interval of time spanning the initial 85% of the training data and testing data is the remaining contiguous 15%. The length of data available for training models exhibits slight variability due to input sequence lengths and lead times. Comparison between models in section 5 is therefore performed on the final 15000 datapoints, or equivalently 150 convective units of the simulation.

#### *Training procedure:*

Parameters for the training procedures used for each neural network are given in Table 1. We used the Adam optimizer [23] with batch size as given in the last column. Restarts, set to 3 for all networks except  $E/D$ , indicates the number of optimizations from random initial weights that were performed. Each was run until validation error failed to decrease for a specified number of epochs, called the patience. The encoder-decoder pair was only trained once due to computational expense. The quantities  $\ell^{1/2}$ -reg describe the  $\ell^1$  and  $\ell^2$  regularization weights applied to connections in non-recurrent layers. LSTM layers as well as connections to the output layer in each network were left unregularized. We note that considerably inflated values used for  $F_\xi$  are due to significant overfitting observed at lower

regularization. A pair (0.98,2) for decay indicates that learning rate was multiplied by 0.98 every two epochs. The initial learning rate in each case was set to 0.001.

Network	Restarts	$\ell^1$ reg.	$\ell^1$ reg.	Decay	Patience	Batch
$F_P$	3	0	1e-3	0.98, 2	10	1000
$F_\psi$	3	0	1e-3	0.98, 2	3	1000
$F_\xi$	3	1e-2	1e-2	0.98, 2	3	1000
$H_P$	3	0	1e-3	0.95, 2	5	250
$H_\psi$	3	0	1e-3	0.95, 2	5	250
$H_\xi$	3	0	1e-3	0.95, 2	5	250
$G$	3	0	1e-5	none	20	100
$E/D$	1	0	1e-3	0.95,1	5	10

Parameters for training procedures were largely set based on empirical evidence and considerations for computational resources. While we claim to have put forward due diligence in tuning all networks for the sake of a valid comparison, we make no claim that these values represent the optimal set for the problem at hand. The authors are not aware of any convincing methods for optimizing such hyper-parameters.

### Appendix C: Wake instability

In section 3 we noted that the intermittent behavior examined in this work is quasi-stable and that in some cases a change in the wake pattern occurs, resulting in less chaotic behavior. To illustrate this, we include here a trajectory similar to the one considered in this work where this mode switching does occur. Figure 15 shows  $C_d(t)$  for a simulation exhibiting a shift in wake behavior around  $t = 400$ . Plots of the vorticity that clearly illustrate the transition are shown in 16. Wake patterns before and after transition resemble the  $P$  and  $2S$  behaviors discussed in [17] and also shown in [30] at lower Reynolds numbers and higher angle of attack.




Figure 15: Time series for the drag coefficient showing transition for intermittent to more regular behavior. Lower row shows close up of time series from two shaded regions.

### Appendix D: Time series of predicted quantity of interest using all methods

Time series prediction of  $\hat{q}$  compared to the true value for test set examples at various values of lead time  $\tau$  are shown in Fig. 17. Networks used to generate these time series are




Figure 16: Vorticity snapshots of the simulation exhibiting mode switching at time  $t = 100$  before switching,  $t = 400$  soon before switching, and  $t = 700$  after switching.

described in further detail in Appendix B. The reduced rank representation of  $\mathbf{u}$  using both POD and FFNN use  $r = 32$ . Figures showing time series result for  $\tau$  sampled between 0 and 10.5 in increments of 0.7 may be found online in the code repository for this work.




Figure 17: Time series for  $q(t)$  and  $\hat{q}(t)$  using each of the neural network based methods on the test set for lead times  $\tau = 1.4, 3.5, 7, 0,$  and  $10.5$ .