

REDUCED-ORDER MODELING OF FLUIDS SYSTEMS,
WITH APPLICATIONS IN UNSTEADY AERODYNAMICS

SCOTT T. M. DAWSON

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
MECHANICAL AND AEROSPACE ENGINEERING
ADVISER: CLARENCE W. ROWLEY

JANUARY 2017

© Copyright by SCOTT T. M. DAWSON, 2017.

ALL RIGHTS RESERVED

Abstract

This thesis focuses on two major themes: modeling and understanding the dynamics of rapidly pitching airfoils, and developing methods that can be used to extract models and pertinent features from datasets obtained in the study of these and other systems in fluid mechanics and aerodynamics. Much of the work utilizes in some capacity dynamic mode decomposition (DMD), a recently developed method to extract dynamical features and models from data.

The investigation of pitching airfoils includes both wind tunnel experiments and direct numerical simulations. Experiments are performed on a NACA 0012 airfoil undergoing rapid pitching motion, with the focus on developing a switched linear modeling framework that can accurately predict unsteady aerodynamic forces and pressure distributions throughout arbitrary pitching motions.

Numerical simulations are used to study the behavior of sinusoidally pitching airfoils. By systematically varying the amplitude, frequency, mean angle and axis of pitching, a comprehensive database of results is acquired, from which interesting regions in parameter space are identified and studied. Attention is given to pitching at “preferred” frequencies, where vortex shedding in the wake is excited or amplified, leading to larger lift forces.

More generally, the ability to extract nonlinear models that describe the behavior of complex fluids systems can assist in not only understanding the dominant features of such systems, but also to achieve accurate prediction and control. One potential avenue to achieve this objective is through numerical approximation of the Koopman operator, an infinite-dimensional linear operator capable of describing finite-dimensional nonlinear systems, such as those that might describe the dominant dynamics of fluids systems. This idea is explored by showing that algorithms designed to approximate the Koopman operator can indeed be utilized to accurately model nonlinear fluids systems, even when the data available is limited or noisy.

Data-driven algorithms can be adversely affected by noisy data. Focusing on DMD, it is shown analytically that the algorithm is biased to sensor noise, which explains a previously observed sensitivity to noisy data. Using this finding, a number of modifications to DMD are proposed, which all give better approximations of the true dynamics using noise-corrupted data.

Acknowledgements

This dissertation would not have come together without the dependable support of many people in the Princeton community and beyond. First I would like to thank my adviser, Clancy Rowley, whose advice, guidance and encouragement have been invaluable throughout my graduate studies, and who always managed to appropriately balance encouraging me to explore new ideas, and guiding me in the right direction.

To everyone who has been in and around H-125 during my time here—Steve Brunton, Lauren Padilla, Jonathan Tu, Brandt Belson, Kevin Chen, Imène Goumiri, Anthony DeGennaro, Hao Zhang, Daniel Floryan, Leonid Pogorelyuk, Vivian Steyert, Mark Luchtenburg, Matt Williams, Maziar Hemati, Utku Senturk and Michael Fairchild—I am thankful for the stimulating environment you all created, and for the countless discussions of ideas both within and outside of the realms of research. I must thank Steve Brunton in particular, who gave me the most ideal of starts to my graduate research, and who continues to provide enthusiastic support.

I would also like to thank my PhD committee and thesis examiners, Lex Smits and Rob Stengel, who provided much insight and advice throughout my graduate studies, and my thesis readers, Luigi Martinelli and Marcus Hultmark, who provided valuable feedback that improved the quality of this dissertation.

I must thank David Williams, for being a welcoming and generous host throughout each of my five trips to the Illinois Institute of Technology, and for giving me free reign over his wind tunnel while I was there. I would also like to thank Bruno Monnier, Xuanhong An, Lou Grimaud, Simion Iliev, Jeremy Weirich, Florian Reißner, David Finol and all of the other members of Professor Williams’ lab, for providing many stimulating discussions, patiently teaching me how to use the wind tunnel, and helping to fix things whenever I broke them. My experience at IIT has given me much more than just Chapter 3 of my thesis. I would additionally like to thank Nicole Schiavone, who was instrumental in helping to collect and analyze the data while she was an undergraduate student working in Professor Rowley’s lab.

There are many other people I need to thank, for being great classmates, roommates, generals study mates, softball and soccer teammates, exercise buddies, and those who do not fall into any of these categories but are fantastic people who made my time here so enjoyable and rewarding. I will not be able to name everybody who falls in at least one of these categories, but such a list would include Brendan Andrade, Steven Atkinson, Anthony DeGennaro, Amanda DeGiorgi, Matthew Feldman, Jie Feng, Katie Fitch, Michael Fairchild, Max Freidfeld, Katie Hartl, Michael Hepler, Kaity Kratsios, Jon Kuyper, Lucy Nicholls, Matt Plasek, AJ Riggs, Will Scott, Xiaoyu Tang, Hannah Waight, and Sheng Yang. I would additionally like to thank Jill Ray and the other MAE administration staff, for keeping everything running smoothly.

Lastly, I would like to thank my family, who let me move so far away, and who continue to be an incredible source of support from the other side of the world.

The work presented in this dissertation was supported by the Air Force Office of Scientific Research (grants FA9550-12-1-0075 and FA9550-14-1-0289). I am additionally grateful for the support of a Francis Robbins Upton Graduate Fellowship throughout my first four years at Princeton, and a Britt & Eli Harari Fellowship for my fifth year of study.

This dissertation carries the number T-3333 in the records of the Department of Mechanical and Aerospace Engineering.

Contents

Abstract	iii
Acknowledgements	iv
List of Tables	ix
List of Figures	x
1 Introduction	1
1.1 Unsteady aerodynamics	2
1.2 Data-driven modeling	3
1.3 Organization and contributions	5
2 Background and motivating results	9
2.1 Data-driven modeling of fluids systems	9
2.1.1 Proper orthogonal decomposition	9
2.1.2 Galerkin projection	10
2.1.3 Dynamic mode decomposition	11
2.1.4 Eigensystem realization algorithm	13
2.2 Motivating results	14
3 A data-driven modeling framework for predicting forces and pressures on a rapidly pitching airfoil	19
3.1 Introduction	19
3.2 System identification method	20
3.3 Experimental method	21
3.4 Results	22
3.4.1 System identification results	23
3.4.2 Multiple pitch-up and pitch-down maneuvers	27
3.4.3 Sinusoidal pitching	28
3.4.4 Pseudo-random pitching	32
3.5 Discussion and conclusions	32
3.6 Appendices	35
3.6.1 Further static airfoil results and analysis	35
3.6.2 Kalman filter design	39
4 Lift enhancement at high angles of attack with periodic pitching	40
4.1 Introduction	40
4.2 Numerical method and scope of investigation	41

4.3	Results	43
4.3.1	Static data	43
4.3.2	Force analysis	44
4.3.3	Wake analysis with dynamic mode decomposition	53
4.3.4	Experimental Results	62
4.4	Discussion and conclusions	63
5	Nonlinear reduced-order models of fluids systems using extended dynamic mode decomposition	65
5.1	Introduction	65
5.2	Reduced-order models using extended dynamic mode decomposition	66
5.2.1	Extended dynamic mode decomposition and nonlinear models	66
5.2.2	A modification to DMD/EDMD	68
5.3	Example: flow past a circular cylinder	69
5.3.1	Data arrangement and selection	70
5.3.2	Comparison between EDMD and GP models, regularization, and model order dependence	72
5.3.3	Model prediction for untrained conditions	74
5.3.4	Noisy data	74
5.3.5	Data from a restricted spatial domain	76
5.3.6	Data from limited temporal sampling	77
5.4	Discussion and conclusions	79
6	Characterizing and correcting for the effect of sensor noise in the dynamic mode decomposition	81
6.1	Introduction	81
6.2	Characterizing noise in dynamic mode decomposition	82
6.2.1	Sensor noise in DMD	82
6.2.2	Direct correction of sensor noise bias in DMD	85
6.2.3	Forward-backward DMD	86
6.2.4	Total least-squares DMD	88
6.3	Results with synthetic data	90
6.3.1	Example: A periodic linear system	90
6.3.2	A periodic linear system with a high-dimensional state of observables	93
6.3.3	Comparison to other modified DMD algorithms	94
6.3.4	Identifying hidden dynamics	95
6.3.5	Differentiating between process and sensor noise	98
6.4	Results with numerical and experimental data	99
6.4.1	Cylinder wake: simulation data	99
6.4.2	Cylinder wake: experimental data	101
6.5	Discussion and conclusions	104
6.6	Appendix: Quantifying the size of the bias in DMD	108
7	Conclusions and future work	110

List of Tables

3.1	Qualitative features observed during pitch-up and pitch-down maneuvers for each pressure and force measurement	27
-----	--	----

List of Figures

2.1	(a) Step responses of various magnitude for the closed-loop system in DNS, and (b) normalized step responses	16
2.2	Step responses of various magnitude for the closed-loop system in wind tunnel experiments.	17
2.3	Controller performance in tracking sinusoidally varying reference lift coefficient.	17
3.1	Schematic of the experimental setup.	22
3.2	Force and pressure coefficient data for a static airfoil at angles of attack between 0° and 22° , as well as for an airfoil sinusoidally pitching between 5° and 15° at a rate $k = \frac{\pi f c}{U} = 0.051$	24
3.3	Force and pressure data for 5° pitch-up maneuvers for starting angles of 0° , 5° , 10° and 15° , which were used for system identification	25
3.4	Force and pressure data for 5° pitch-down maneuvers for starting angles of 5° , 10° , 15° and 20° , which were used for system identification	26
3.5	Example of inaccurate prediction of a pitch-up model ($5^\circ - 10^\circ$ model, red curves) on pitch down maneuver ($10^\circ - 5^\circ$ maneuver, blue curves).	28
3.6	Switched model performance in predicting pressure and force coefficients for double pitch up/down maneuvers with different pitch-rates (increasing from (a) to (c)), between 10° and 20°	29
3.7	Performance of a single ($5^\circ - 10^\circ$ pitch-up) model in predicting pressure and force coefficients for double pitch up/down maneuver.	30
3.8	Performance of switched model in predicting pressure and force coefficients for quadruple pitch up/down maneuver	31
3.9	Predicted and actual pressure and force coefficients for high-amplitude sinusoidal pitching at dimensionless frequencies $k = \frac{\pi f c}{U} = 0.051$ and 0.103	33
3.10	Performance of the switched model, equipped with a Kalman filter and measurements of the first and last pressure coefficient, in predicting pressure and force coefficients for high-amplitude pseudo-random pitching maneuver.	34
3.11	Location of static separation point (determined by pressure maxima) as function of α	36
3.12	Power spectral densities of the pressures, forces, and angle of attack measurements made with the tunnel turned off.	37
3.13	Power spectral densities of the pressures, forces, and angle of attack measurements for angles of attack ranging between 0° (black) and 21° (red), with pertinent frequencies identified by vertical lines.	37

3.14	Power spectral density of pressure and force signals at selected frequencies as a function of α	38
4.1	Computational domain used for this study, with the size and location of the airfoil shown	42
4.2	Resolution study to determine appropriate size of computational domain . .	43
4.3	Lift curve for stationary airfoil, showing regions where the equilibrium is stable ($\alpha < 27^\circ$) and unstable ($\alpha > 27^\circ$), above which periodic vortex shedding occurs	44
4.4	Mean lift coefficient for a range of pitching amplitudes α_A , frequencies, f^* , and mean angles of attack, α_M	45
4.5	Mean drag coefficient for a range of pitching amplitudes α_A , frequencies, f^* , and mean angles of attack, α_M	46
4.6	Mean of lift-to-drag ratio for a range of pitching amplitudes α_A , frequencies, f^* , and mean angles of attack, α_M	47
4.7	Mean lift coefficient increment over the fixed airfoil value, normalized by the amplitude of pitching α_A , pitching about midchord	49
4.8	Mean lift coefficient increment over the fixed airfoil value, normalized by the amplitude of pitching α_A , pitching about leading edge	50
4.9	A comparison between the maximum, minimum, and mean lift coefficient attained for pitching with amplitude $\alpha_A = 5^\circ$, for a range of mean angles of attack, α_M , and frequencies, f^*	51
4.10	Power spectral densities of the lift force for pitching motion about the midchord with amplitude $\alpha_A = 1^\circ$	52
4.11	Instantaneous vorticity fields for pitching about the leading edge (left) and midchord (right) at a variety of frequencies, with $\alpha_M = 20^\circ$ and $\alpha_A = 1^\circ$. . .	54
4.12	Instantaneous vorticity fields for pitching about the leading edge (left) and midchord (right) at a variety of frequencies, with $\alpha_M = 20^\circ$ and $\alpha_A = 5^\circ$. . .	55
4.13	Instantaneous vorticity fields for pitching about the midchord (right) at a variety of frequencies, with $\alpha_A = 5^\circ$ and $\alpha_M = 25^\circ$ (left) and $\alpha_M = 30^\circ$ (right). .	56
4.14	DMD eigenvalue amplitudes and frequencies (left) and eigenmodes corresponding to the four largest amplitude eigenmodes, for pitching with $\alpha_M = 25^\circ$, $\alpha_A = 5^\circ$, and $f^* = 0.05, 0.1, 0.2$ and 0.25	58
4.15	DMD eigenvalue amplitudes and frequencies (left) and eigenmodes corresponding to the four largest amplitude eigenmodes, for pitching with $\alpha_M = 25^\circ$, $\alpha_A = 5^\circ$, and $f^* = 0.3, 0.35, 0.4$ and 0.5	59
4.16	DMD eigenvalue amplitudes and frequencies (left) and eigenmodes corresponding to the four largest amplitude eigenmodes, for pitching with $\alpha_M = 30^\circ$, $\alpha_A = 5^\circ$, and $f^* = 0.05, 0.1, 0.2$ and 0.25	60
4.17	DMD eigenvalue amplitudes and frequencies (left) and eigenmodes corresponding to the four largest amplitude eigenmodes, for pitching with $\alpha_M = 30^\circ$, $\alpha_A = 5^\circ$, and $f^* = 0.3, 0.35, 0.4$ and 0.5	61
4.18	Comparison between experimental data for a NACA 0009 airfoil at a Reynolds number of 50,000, with numerical results previously shown in Figure 4.7 for $\alpha_M = 20^\circ$	63

5.1	Computational domain used for numerical simulation of flow past a circular cylinder. Dashed lines represent the borders of each nested grid. The gray shading denotes the region from which data was collected for modeling purposes	69
5.2	Contours of streamwise velocity for the (a) mean mode and (b)-(j) first nine POD modes, along with the corresponding time-varying coefficients of the POD modes identified from data collected as the flow transitions from near the unstable equilibrium to the limit cycle	71
5.3	(a) Relative and (b) cumulative energy content of the first 30 POD modes, for the same data used in Figure 5.2	72
5.4	Performance of 3rd order GP and EDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, showing (a) time evolution and (b) phase portrait plots	72
5.5	Performance of 9th order GP and EDMD nonlinear models in predicting the time-evolution of POD coefficients for transitional flow past a cylinder, showing (a) time evolution and (b) phase portrait plots	73
5.6	Prediction of the limit cycle amplitude (a) and frequency (b) of the first POD mode from both EDMD and GP models	74
5.7	Performance of 3rd order GP and EDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, with the initial condition taken to be the mean of the limit cycle	75
5.8	Performance of 3rd order GP and EDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, with models identified from data that is corrupted by noise	75
5.9	Performance of 3rd order GP and EDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, using data from a restricted spatial domain	76
5.10	Performance of 3rd order GP and EDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, using data obtained at three spatial locations	77
5.11	Performance of 3rd order GP and eDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, using data from restricted temporal samplings	78
6.1	Illustrative diagram showing how the error in estimation of a given quantity can be decomposed into bias error and random error	85
6.2	Eigenvalues (in continuous-time) identified by using variants of DMD from 100 snapshots of data from Equation (6.3.1)	91
6.3	Error in the estimated propagation matrix identified from performing DMD and ncDMD on noise-corrupted data generated from Equation (6.3.1), for different values of m and σ_N^2	92
6.4	Error in the estimated propagation matrix identified from performing DMD, ncDMD, fbDMD, and tlsDMD on noise-corrupted data generated from Equation (6.3.1), for different values of m and σ_N^2	92

6.5	Mean and 95% confidence ellipses of continuous-time eigenvalues identified by applying unbiased variants of DMD to 1000 trials of noisy data generated by Equation (6.3.1), and observed as in Equation (6.3.2)	94
6.6	Mean and 95% confidence ellipses of continuous-time eigenvalues identified by applying DMD, fbDMD, OMD and spDMD to 1000 trials of noisy data generated by Equation (6.3.1), and observed as in Equation (6.3.2)	96
6.7	Visualization of data generated by Equation (6.3.3), with $k_1 = 1$, $\omega_1 = 1$, $\gamma_1 = 1$, $k_2 = 0.4$, $\omega_2 = 3.7$, $\gamma_2 = -0.2$, and $\sigma = 0.5$	97
6.8	Mean and 95% confidence ellipses of continuous-time eigenvalues identified by applying unbiased variants of DMD to 1000 trials of noisy data generated by Equation (6.3.3)	98
6.9	Eigenvalues identified using (a) DMD, (b) ncDMD, and (c) tlsDMD for the Stuart-Landau equation (Equation (6.3.4)), with 100,000 snapshots of data from Equation (6.3.5), with and without sensor and process noise	100
6.10	(a) Eigenvalues and (b) POD coefficients identified from applying DMD, ncDMD fbDMD, and tlsDMD to noise-corrupted DNS vorticity data from a cylinder wake at $Re = 100$	102
6.11	(a) A subset of the DMD modes (real components) computed from applying various variants of DMD to DNS data of flow around a cylinder. (b) Inner product between (normalized) clean modes, and modes obtained from noisy data (with $\sigma_N = 0.2 \frac{U_\infty}{D}$)	103
6.12	(a) Eigenvalues and (b) POD coefficients identified from applying DMD and tlsDMD to experimental vorticity data	105

Chapter 1

Introduction

The study of aerodynamics and fluid mechanics predate the birth of the scientific method itself. From the initial musings of Aristotle and Archimedes, to the sketches of Leonardo Da Vinci, to George Caley’s conception of a modern airplane configuration in 1799, the field has often been at the forefront of scientific and engineering enquiry [5]. Yet, in spite of this long history, and the fact that the underlying governing equations of viscous fluid flow being known for hundreds of years, there is still much to learn. For example, in many respects we are still catching up to the mastery exhibited by biological swimmers and fliers, in terms of maneuverability and efficiency. This existence of ongoing work is not by any means due to a lack of imperative. Indeed, reductions in aerodynamic drag on cars, trucks, airplanes and ships can lead to billions of dollars in fuel savings [55], and significant reductions of CO₂ emissions, to speak nothing of the enhanced safety that would come with, for example, a more comprehensive understanding of the dynamics of aircraft in deep stall [e.g., 93].

On top of developments in materials, manufacturing, and the miniaturization of processors allowing for new research questions to become of practical interest, advances in both experimental techniques and equipment, and computational power and storage capacity, researchers in fluid dynamics can now generate more high-fidelity data than ever before. This is crucial to advancing the field as a whole, since all but the most idealized, simple systems cannot be fully understood by analytical deductions alone. new questions to become of practical interest. This is not to suggest that the field of fluid dynamics and aerodynamics are now confined to the realms of data science. Indeed, progress is made by combining the insight attained from study of the fundamental equations with the additional dimension of large datasets. Indeed, obtaining the right data relies on understanding of the physical system. Beyond this, however, data collection should also be informed by a proper understanding of the capabilities and limitations of the algorithmic tools that are required to assist in data analysis.

The presence of increasingly large data sets necessitates the use of such post-processing techniques that are able to extract tractable and physically relevant information from the data. This thesis will consider methods to extract pertinent information and models from data obtained from fluids simulations and experiments. While the focus will be on unsteady aerodynamic systems, specifically pitching airfoils, the techniques used and developed are applicable to a wider range of applications, both within and outside the realms of fluid mechanics.

1.1 Unsteady aerodynamics

The flight of small, highly maneuverable aircraft, whether biological or manmade, is greatly impacted by unsteady aerodynamic effects, which can be either beneficial or detrimental to flight. Accurate understanding of such effects can allow for the design of aircraft that are more efficient, responsive, and robust.

The need to account for unsteady effects has been recognized since soon after the breakthrough of powered manmade flight, in the classical works of Wagner [167], Theodorsen [154], and Garrick [56]. Indeed, many failed attempts at flight can probably be attributed to a severe lack of understanding of how to utilize such effects. These classical models give significant insight into the fundamental flow physics associated with unsteady flight, such as relative contributions to lift of the added mass, quasi-steady bound circulation, and wake vortices. For example, the Theodorsen model gives the relationship between the airfoil kinematics (α , $\dot{\alpha}$, and $\ddot{\alpha}$, where α is the angle of attack) and the lift coefficient by:

$$C_L = \frac{\pi}{2} \left(\dot{\alpha} - \frac{1}{2} a \ddot{\alpha} \right) + 2\pi \left(\alpha + \frac{1}{2} \dot{\alpha} \left(\frac{1}{2} - a \right) \right) C(k), \quad (1.1.1)$$

where a is a parameter that defines the pitch axes, with $a = -1$ and $+1$ corresponding to pitching about the leading and trailing edge of the airfoil, respectively. $k = \frac{\pi f c}{U}$ is the reduced frequency, and $C(k)$ is the Theodorsen function, which governs the interaction between the shed vorticity in the wake and the circulatory lift force. While such models can be quantitatively accurate for cases of attached flow where viscous effects are negligible, they quickly lose validity when dealing with separated flows, which are often encountered in the extreme motions that are possible for birds, insects, and micro and unmanned aerial vehicles (MAV and UAV). It is precisely in these extreme cases that accurate predictive models are essential to prevent catastrophic failure and ensure ongoing successful flight. While more accurate predictions can be attained from high-fidelity simulations, the computational cost typically prohibits the direct use of such simulations for real-time prediction and control.

Biological examples such as insects [19, 130, 170] and birds [165] have seemingly evolved to take advantage of the high transient lift force that can be generated due to the formation of a leading edge vortex (LEV) during rapid pitch-up motion, for example. While these give motivating examples of the advantages of accurate understanding of unsteady aerodynamic effects, the preferred wing kinematics arising from evolution is highly specific and coupled to the geometry and other physiological features of the animal. Indeed, the characteristics of unsteady aerodynamic effects, particularly for separated flows, seem to be quite sensitive to both the geometry [86] and Reynolds number [178] of the airfoils. Studies into low Reynolds number flow over stationary symmetric airfoils [166, 1, 140, 32, 87] reveal complex, Reynolds number dependent effects associated with flow separation and reattachment, Kelvin-Helmholtz shear layer instabilities, transition to turbulence, and von Kármán vortex shedding, with the dynamics of pitching airfoils influenced by these phenomena as well as dynamic stall vortex formation, growth and detachment, which are again sensitive to Reynolds number [166, 1, 140, 32, 87]. These observations motivate the development of general modeling procedures that can be easily applied to a range of parameter cases. In addition, it is desirable for such methods to be sufficiently general such that they can be

applied to more realistic aircraft configurations, rather than just airfoils. As an example, such data driven modeling was considered for the case of accurate prediction and control of lift for a low Reynolds number pitching airfoil [23, 24], using the eigensystem realization algorithm [79] (ERA) and observer/Kalman filter identification [82] (OKID). There has also been a significant amount of work in terms of nonlinear modeling, ranging from low order state- space models formulated from theoretical considerations [60], to Volterra series models that have been used to model a range of unsteady aerodynamic and aeroelastic phenomena [141, 94, 12].

1.2 Data-driven modeling

As has already been mentioned, researchers in fluid dynamics can now generate more high-fidelity data than ever before. The presence of increasingly large data sets calls for appropriate data analysis techniques, that are able to extract tractable and physically relevant information from the data. In particular, a much-desired goal in fluid mechanics, and indeed many other fields, is to obtain simple models that are capable of predicting the behavior of seemingly complex systems. Low-dimensional models can not only improve our fundamental understanding of such systems, but are often required for purpose of efficient and accurate prediction, estimation and control. Broadly speaking, one can obtain low-dimensional information about a system (whether it be in the form of a reduced-order model, or simply spatial modes corresponding to certain energetic or dynamic characteristics) in numerous ways, potentially using some combination of data collected from simulations and experiments, and theoretical knowledge of the system, such as the governing partial differential equations (PDEs).

Purely data-driven methods can include those developed particularly for fluids applications, such as the dynamic mode decomposition (DMD) [136, 134], or those which are appropriated from other communities, such as the eigensystem realization algorithm (ERA) [89, 79], which was first applied to study spacecraft structures, but has more recently been appropriated to model a wide range of fluids systems [26, 2, 75, 76, 23, 24, 16, 74, 54].

Dynamic mode decomposition allows for the identification and analysis of dynamical features of time-evolving fluid flows, using data obtained from either experiments or simulations. In contrast to other data-driven modal decompositions such as the proper orthogonal decomposition (POD), DMD allows for spatial modes to be identified that can be directly associated with characteristic frequencies and growth/decay rates. Following its conception, DMD was quickly shown to be useful in extracting dynamical features in both experimental and numerical data [136, 134]. It has subsequently been used to gain dynamic insight on a wide range of problems arising in fluid mechanics [129, 137, 135, 109, 138, 47, 64, 77, 90, 103, 61, 133, 49] and other fields [65].

DMD has a strong connection to Koopman operator theory [88, 98], as exposed in Rowley et al. [129], and further reviewed in Mezić [99], which can justify its use in analyzing nonlinear dynamical systems. Since its original formulation, numerous modifications and extensions have been made to DMD. Chen et al. [28] highlights the connection that DMD shares with traditional Fourier analysis, as well as proposing an optimized algorithm that recasts DMD as an optimal dimensionality reduction problem. This concept of finding only the dynamically

important modes has also been considered in subsequent works of Wynn et al. [176] and Jovanović et al. [78]. All of these works are motivated, in part, by the fact that by default DMD will output as many modes as there are pairs of snapshots (assuming that the length of the snapshot vector is greater than the number of snapshots), which is arbitrary with respect to the dynamical system under consideration. In reality, one would prefer to output only the modes and eigenvalues that are present (or dominant) in the data. When the data is corrupted by noise (as will always be the case to some degree, especially for experimental data), this process becomes nontrivial, since noisy data might have a numerical rank far larger than the dimension of the governing dynamics of the system. Further to this, one cannot expect to have a clean partition into modes that identify true dynamical features, and those which consist largely of noise.

Simple ways of achieving this objective can involve either first projecting the data onto a smaller dimensional basis (such as the most energetic POD modes) before applying DMD, or by choosing only the most dynamically important DMD modes after applying DMD to the full data. One could also truncate the data to a dimension larger than the assumed dimension of the dynamics, and then apply a balanced truncation to the resulting dynamical system to obtain the desired reduced-order model. This approach is sometimes referred to as *over-specification* in the system identification community [e.g., 81]. Keeping a higher dimension of data than that of the assumed dynamics can be particularly important for input-output systems that have highly energetic modes that are not strongly observable or controllable [122]. Ideally, any algorithm that restricts the number of DMD modes that are computed should also additionally be computationally efficient. A fast method to perform DMD in real time on large datasets was recently proposed in Hemati et al. [67], while a library for efficient parallel implementation of number of common modal decomposition and system identification techniques is described in Belson et al. [17]. Sayadi and Schmid [132] also gives an explicit implementation of DMD for parallelized computation. One can also achieve computational speedup by incorporating efficient methods to compute singular value decompositions, typically the computational bottleneck in DMD, to speed up the computation [53].

An notable limitation of the methods mentioned so far is that (when considered in the context of data-driven reduced-order modeling techniques) they are linear, in the sense that the reduced-order model that is identified is in the form of a linear system of ordinary differential equations (ODEs). While there have been a number of examples of nonlinear data-driven modeling techniques used in fluids applications [94, 112, 115, 58, 12, 39, 84, 68, 41, 25, 70, 6], their widespread use has been more limited, and the underlying theory is less established, than linear techniques. More details concerning the application of data-driven modeling techniques in fluid mechanics can be found in recent review articles [22, 123].

Perhaps the most common method to obtain a nonlinear reduced-order model for fluids systems comes via a projection of the governing equations onto a low-dimensional basis that is optimal for capturing the energy of the data, i.e., the proper orthogonal decomposition (POD) [95, 18, 73], which is a procedure referred to as Galerkin projection. Galerkin projection (GP) has been used to extract models for many different fluids systems, a non-exhaustive list includes flow past a cylinder at low Reynolds number [44, 105, 106], grooved channels [44] the wall region of turbulent boundary layers [7, 114], flat plate boundary layers [121], turbulent plane Couette flow [101, 144], turbulent pipe flow [20] cavity oscillations [126, 125],

mixing layers [120, 162, 13], and compressible flows [128]. One significant drawback of GP models is that they ignore modes that are low in energy, but are required for the dissipation of energy in the full system. A number of modifications have been proposed to address this concern, as well as other issues with such models. Aubry et al. [7] and Podvin [114] use an eddy viscosity term that accounts for energy dissipated into unmodeled modes, Östth et al. [108] investigate a hierarchy of eddy viscosity formulations, while Wang et al. [168, 169] incorporate LES closure modeling strategies. Cordier et al. [37, 38] summarize a number of calibration techniques that can be used to improve the accuracy of Galerkin models, and also discuss the various ways in which the error of such models can be quantified. Balajewicz et al. [14] employs a subspace rotation technique to stabilize the models, which, unlike other calibration techniques, maintains consistency with the original governing equations. Balajewicz et al. [13] imposes constraints to balance the turbulent kinetic energy of the resulting model. All of these modifications of Galerkin projection heighten the “data-driven” nature of the method. Noack et al. [107] gives an in-depth summary and analysis of many issues, variations, progress, and open problems on the topic of Galerkin projection models, while Luchtenburg et al. [92] gives a clear expository introduction of the main ideas in Galerkin projection, with examples.

While it was mentioned above that DMD could be classified as a “linear” method, connections between the DMD algorithm and the Koopman operator [129, 99] give promise that it can ultimately be used to model and understand nonlinearities. The Koopman operator [88] gives a means of representing a finite-dimensional, nonlinear system as an infinite dimensional linear system, and DMD gives a finite-dimensional approximation to this operator.

In particular, an extension of DMD that potentially allows for better representation of nonlinear data has also recently been proposed [172], and although the computational costs increase dramatically with the dimension of the system, a kernel method described in Williams et al. [173] reduces the cost to be comparable to standard DMD.

1.3 Organization and contributions

Following this introductory chapter, Chapter 2 presents a summary of key concepts and techniques in data-driven modeling of fluid systems, as well as presenting some results that, besides being of some independent interest, will motivate the research directions taken in the subsequent chapters. Broadly speaking, Chapters 3 and 4 will focus specifically on the application pitching airfoils, while Chapters 5 and 6 focus on data-driven modeling techniques, particularly extensions and improvements to DMD and their application to fluids systems. A more detailed description of the contents of this thesis follows. As well as acknowledging the published and forthcoming literature related to this thesis, this section also details the contributions of coauthors to these publications. Unless otherwise mentioned, the conception, execution and writing up of all research is my own, with guidance from Clarence Rowley across all phases. The acknowledgments section of this thesis gives a more comprehensive account of the many additional people who provided valuable assistance for my research.

Chapter 3 obtains models to predict the pressures and forces on a rapidly pitching airfoil. This is one of the first applications to experimental data of a recently developed

variant of DMD by Proctor et al. [116] to allow for the identification of systems that contain inputs (e.g., systems that are being forced or controlled externally in some manner). We show that this modeling approach is convenient for constructing “switched models”, whereby one can predict the behavior of a nonlinear system by switching between a family of linear models. In particular, the “DMD with inputs” method gives models for which the coordinates of the models remain consistent with each other, which eliminates complications and ill-conditioning that has been observed when using alternate methods [40]. This modeling approach allows for the formulation of a switched model that remains accurate over a wide range of angles of attack, ranging from attached to fully separated flow. The experiments were performed at the Andrew Fejer Unsteady Flow wind tunnel at the Illinois Institute of technology under the supervision of Professor David Williams. The experiments and subsequent analysis were performed with the assistance of Nicole Schiavone, an undergraduate working in Professor Rowley’s lab in summer 2014. Material in this chapter is based on the conference paper [41]:

- S. T. M. Dawson, N. K. Schiavone, C. W. Rowley, and D. 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.

Chapter 4 explores a phenomena that is identified in Chapter 2: that airfoils undergoing low-amplitude sinusoidal pitching motion generate enhanced lift when pitching at preferred frequencies. A systematic parameter sweep over the pitching frequency, amplitude, and base angle of attack is conducted, with the mean and frequency content of the forces analyzed. In addition, the flow fields are studied using DMD. Daniel Floryan and Maziar Hemati provided assistance in the conception of this project, and in editing the resulting paper. Material in this chapter is based on the conference paper [42]:

- S. T. M. Dawson, D. C. Floryan, C. W. Rowley, and M. S. Hemati. Lift enhancement of high angle of attack airfoils using periodic pitching. In *54th AIAA Aerospace Sciences Meeting*, page 2069, 2016.

Chapter 5 shows how a recently developed extension to DMD can be utilized to obtain nonlinear reduced-order models for fluids systems. We modify the extended DMD (EDMD) algorithm to include a Tikhonov regularization step, which is found to give improved results for the purposes of nonlinear system identification. The method is demonstrated on the canonical example of flow past a circular cylinder, for data starting near the unstable equilibrium solution and converging to the periodic vortex shedding limit cycle. It is demonstrated that this approach can be superior to classical POD-Galerkin projection, particularly in cases where the data is noisy, is from a limited spatial region, is not spatially resolved, or is only collected over a short time window. Material in this chapter is contained in the paper:

- S. T. M. Dawson and C. W. Rowley, Nonlinear reduced-order models of fluids systems using extended dynamic mode decomposition, *In preparation for Theoretical and Computational Fluid Dynamics*, 2016.

Some of the results presented in Chapter 5 are also used in an upcoming review paper [123]:

- C. W. Rowley and S. T. M. Dawson. Model reduction for flow analysis and control. *Annual Review of Fluid Mechanics*, 49(1), 2017.

I was the primary contributor to the sections of this paper that feature research presented in this thesis.

Chapter 6 analyzes the effect of noise of DMD. As well as giving an explanation for a previously identifies sensitivity to noisy data, three variants of the DMD algorithm are proposed, all of which give improved performance on noisy data. Maziar Hemati and Matthew Williams assisted in some of the derivations of some of the mathematical results, and Hemati conceived the original idea for a total least-squares variant of DMD. Additionally, one example in this chapter makes use of experimental data provided by Jessica Shang. This work is published in the following paper [43]:

- S. T. M. Dawson, M. S. Hemati, M. O. Williams, and C. W. Rowley. Characterizing and correcting for the effect of sensor noise in the dynamic mode decomposition. *Experiments in Fluids*, 57(42):1–19, 2016.

The introductory and background chapters (Chapters 1 and 2) also use material from the aforementioned publications. In addition to these publications, Chapter 2 also uses feedback control ideas that are published in the following paper [24]:

- S. L. Brunton, S. T. M. Dawson, and C. W. Rowley. State-space model identification and feedback control of unsteady aerodynamic forces. *Journal of Fluids and Structures*, 50:253–270, 2014.

Steven Brunton conceived the research idea to design feedback controllers using reduced-order models that he had identified, and assisted me in their implementation in numerical simulations.

In addition to the these publications, elements of this thesis have also been presented in the following conference presentations:

- S. T. M. Dawson, M. S. Hemati, M. O. Williams, and C. W. Rowley. Improved Dynamic Mode Decomposition Algorithms for Noisy Data. *SIAM Conference on Uncertainty Quantification*, Lausanne, Switzerland, April 2016 (invited speaker).
- S. T. M. Dawson and C. W. Rowley. Nonlinear models for fluids systems using extended dynamic mode decomposition. *68th Annual Meeting of the APS Division of Fluid Dynamics*, Boston, MA, November 2015.
- M. S. Hemati, M. O. Williams, S. T. M. Dawson, and C. W. Rowley. Koopman Operator Theoretic Methods in Systems and Control, *SIAM Conference on Control and its Applications*, Paris, France, July 2015.
- S. T. M. Dawson, M. S. Hemati, M. O. Williams, and C. W. Rowley, Improving the Accuracy of Dynamic Mode Decomposition in the Presence of Noise. *SIAM Conference on Applications of Dynamical Systems*, Snowbird, UT, May 2015 (invited speaker).

- C. W. Rowley, M. O. Williams, M. S. Hemati, and S. T. M. Dawson. Reduced-order Models using Dynamic Mode Decomposition. *SIAM Conference on Computational Science and Engineering*, Salt Lake City, UT, March 2015.
- S. T. M. Dawson, M. S. Hemati, M. O. Williams, and C. W. Rowley. Characterizing and correcting for the effect of sensor noise in the dynamic mode decomposition. *67th Annual Meeting of the APS Division of Fluid Dynamics*, San Francisco, CA, November 2014.
- N. K. Schiavone, S. T. M. Dawson, C. W. Rowley, and D. R. Williams. Modeling unsteady forces and pressures on a rapidly pitching airfoil. *67th Annual Meeting of the APS Division of Fluid Dynamics*, San Francisco, CA, November 2014.
- S. T. M. Dawson, S. L. Brunton, and C. W. Rowley. Nonlinear switched models for control of unsteady forces on a rapidly pitching airfoil. *66th Annual Meeting of the APS Division of Fluid Dynamics*, November 2013.
- S. T. M. Dawson, S. L. Brunton, and C. W. Rowley. Feedback control of a pitching and plunging airfoil via direct numerical simulation. *65th Annual Meeting of the APS Division of Fluid Dynamics*, November 2012.
- S. L. Brunton, S. T. M. Dawson, and C. W. Rowley. Feedback control of a pitching airfoil based on unsteady lift models. *42nd AIAA Fluid Dynamics Conference*, June 2012.

Effort is made to keep notation consistent throughout this thesis, however on occasion notation changes between chapters, in attempt to uphold existing conventions in the relevant fields. While each of these chapters is largely self contained, we present in Chapter 2 underlying preliminaries that are broadly relevant across all sections of this thesis. Due to the distinct subject areas covered, we include additional introductory material in each chapter, and in some cases include appendix sections within each chapter, for material that is relevant to the chapter but is tangential to the main ideas and results. With this structure readers should be able to move from Chapter 2 to any chapter of particular interest.

Chapter 2

Background and motivating results

This chapter will introduce some concepts that will be used and referred to on a number of occasions throughout this thesis. Section 2.1, presents a range of algorithms that have seen common use for identifying models and features of fluids systems. Relevant literature and applications of such methods were discussed in detail in Section 1.2. The method that is utilized more predominantly in this thesis is DMD (Section 2.1.3), which is featured to various extents in each of Chapters 3–6. The description of DMD is preceded by a discussion of POD in Section 2.1.1. This serves as background both for both DMD and Galerkin projection (GP, Section 2.1.2). GP, which utilizes POD modes as an efficient low-dimensional basis for approximating the governing equations, is used in Chapter 5 as a comparison to the EDMD method of nonlinear system identification. We include a presentation of ERA both since it is used in Section 2.2, and to highlight the similarities with DMD; similarities which exist between numerous data-driven linear modeling/system identification algorithms.

2.1 Data-driven modeling of fluids systems

2.1.1 Proper orthogonal decomposition

The goal of the proper orthogonal decomposition (POD) is to obtain a set of empirical spatial modes that optimally represent a given dataset from an energetic standpoint. Assume that we can decompose the dynamics of some system $\mathbf{u}(x, t)$ (which could be the time-varying velocity field of a fluid, say) by

$$\mathbf{u}(x, t) = \mathbf{u}_0(x) + \sum_{i=1}^{\infty} \mathbf{u}_i(x) a_i(t), \quad (2.1.1)$$

where $\mathbf{u}_0(x)$ is some fixed (often average) data, and $\{\mathbf{u}_i(x)\}_{i=1}^{\infty}$ are a set of orthonormal basis functions (modes). POD takes these modes to be those which successively capture the most energy of the velocity field. Each POD mode \mathbf{u}_i satisfies the integral

$$\int_{\Omega} \mathbb{E}[\mathbf{u}_i(x) \mathbf{u}_j(x')] \mathbf{u}_i(x') dx' = \lambda \mathbf{u}_i(x),$$

where $\mathbb{E}[\mathbf{u}_i(x)\mathbf{u}_j(x')]$ denotes the expectation. As indicated by Equation (2.1.1), POD is normally performed after first subtracting the mean (or perhaps an equilibrium point) from the data. This approach has the advantage that \mathbf{u}_0 satisfies the required non-homogeneous boundary conditions, meaning that all other modes \mathbf{u}_i will satisfy homogenous boundary conditions, so any linear combination of modes of the form given by Equation (2.1.1) will automatically satisfy the correct boundary conditions of the problem at hand. In discrete terms, if we arrange finite-dimensional data collected from a simulation or experiment into a matrix \mathbf{Y} , with each column representing a snapshot of data at a given time, then the POD modes are the columns of \mathbf{U} in the singular value decomposition $\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$. Here the i^{th} entry of the diagonal matrix $\mathbf{\Sigma}$ corresponds to the energy contained in the i^{th} POD mode. In this discrete formulation, for simplicity we are omitting any rescaling of the data that should be performed so that the modes are orthonormal with respect to the usual inner product. That is, if \mathbf{u}_i and \mathbf{u}_j are columns of \mathbf{U} , then we really should have

$$\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \int_{\Omega} \mathbf{u}_j^*(x') \mathbf{u}_i(x') dx' \approx \sum_{k=1}^n \mathbf{u}_j^*(x_k) \mathbf{u}_i(x_k) dx_k = \mathbf{u}_j^* \mathbf{W} \mathbf{u}_i = \delta_{ij},$$

for an appropriate weight matrix \mathbf{W} , rather than $\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \mathbf{u}_j^* \mathbf{u}_i = \delta_{ij}$.

The original data \mathbf{Y} can be represented in terms of POD coefficients by $\tilde{\mathbf{Y}} = \mathbf{U}^* \mathbf{Y}$. If we wish to reduce the dimension of this data, we may do so in an optimal way (with respect to energy content) by simply truncating the columns of \mathbf{U} beyond a certain point, which corresponds to removing POD modes that are of sufficiently low energy. Doing this gives a reduced-order approximation of the data $\tilde{\mathbf{Y}}_r = \mathbf{U}_r^* \mathbf{Y}$, where \mathbf{U}_r contains the first r columns of \mathbf{U} . Note that there are alternative truncation techniques that may be more effective than energy maximization for certain applications, for example balanced POD [122] gives a reduced-order linear state space model that is optimal with respect to a given set of sensors and actuators.

2.1.2 Galerkin projection

The idea behind GP is to approximate the governing PDEs that describe a given system with a low-dimensional set of ODEs. This is accomplished by projecting the equations onto spatial POD modes identified using the methods described in Section 2.1.1. We begin with the incompressible Navier–Stokes equations:

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} &= -(\mathbf{u} \cdot \nabla) \mathbf{u} + \nu \Delta \mathbf{u} - \nabla p \\ \nabla \cdot \mathbf{u} &= 0. \end{aligned} \tag{2.1.2}$$

If we take the (spatial) inner product of Equation (2.1.2) with a given mode \mathbf{u}_j , we obtain

$$\left\langle \frac{\partial \mathbf{u}}{\partial t}, \mathbf{u}_j \right\rangle = -\langle (\mathbf{u} \cdot \nabla) \mathbf{u}, \mathbf{u}_j \rangle + \nu \langle \Delta \mathbf{u}, \mathbf{u}_j \rangle - \langle \nabla p, \mathbf{u}_j \rangle. \tag{2.1.3}$$

Substituting in the finite-dimensional approximation of Equation (2.1.1), we obtain

$$\dot{\mathbf{a}} = \mathbf{L}\mathbf{a} + \mathbf{Q}(\mathbf{a}, \mathbf{a}) + \mathbf{f}, \quad (2.1.4)$$

where \mathbf{L} is a linear operator (i.e., a matrix), \mathbf{Q} is a bilinear operator, and \mathbf{f} is a vector, each defined based on the identified spatial POD modes by

$$\mathbf{L}_{ij} = \nu \langle \Delta \mathbf{u}_j, \mathbf{u}_i \rangle, \quad \mathbf{Q}_{ijk} = - \langle (\mathbf{u}_j \cdot \nabla) \mathbf{u}_k, \mathbf{u}_i \rangle, \quad \mathbf{f} = - \langle \nabla p, \mathbf{u}_i \rangle. \quad (2.1.5)$$

This gives a means of approximating the Navier-Stokes equations by a set of nonlinear ODEs. As mentioned in Section 6.1, there are many modifications that have been proposed for this general procedure, most typically to account for the energy transfer to unmodeled modes (i.e., the energy cascade to finer spatial scales). For cases where spatial symmetries exist (e.g., in the streamwise and azimuthal directions for circular pipe flow), one can show that the POD modes must become Fourier modes, which can simplify their computation. It is also possible to “factor out” such symmetries by using an optimally chosen moving frame of reference [124, 127].

2.1.3 Dynamic mode decomposition

DMD has undergone a number of formulations, interpretations and modifications since its inception. Common to all methods is the requisite collection and arrangement of data, which is now summarized. Suppose we collect snapshots of data \mathbf{y}_i , which we assemble as columns in the data matrix \mathbf{Z} . For fluids systems \mathbf{y}_i will typically be a velocity field snapshot, but more generally it is a vector of observations of an evolving dynamical system at a given time. From \mathbf{Z} , we select all pairs of columns that are sampled at a time difference Δt apart, and place them into the matrices \mathbf{Y} and $\mathbf{Y}^\#$ (where the data in a given column of $\mathbf{Y}^\#$ was collected Δt after the equivalent column of \mathbf{Y}). Note that if \mathbf{Z} consists of a sequential time-series of data, then \mathbf{Y} and $\mathbf{Y}^\#$ are simply \mathbf{Z} with the last and first columns excluded, respectively. Let \mathbf{Y} and $\mathbf{Y}^\#$ each be n by m matrices, so we have m pairs of snapshots, each of size n . By not explicitly requiring a single time-series of data, we allow for larger or irregular time gaps between snapshot pairs, the concatenation of data from multiple time-series, and for the removal of any corrupted or spurious data. Recently, Tu et al. [161] proposed an interpretation of DMD modes and eigenvalues as the eigendecomposition of the matrix

$$\mathbf{A} = \mathbf{Y}^\# \mathbf{Y}^+, \quad (2.1.6)$$

where \mathbf{Y}^+ denotes the Moore-Penrose pseudoinverse of a matrix \mathbf{Y} . While this is a succinct interpretation, and one which will be useful in the ensuing analysis, it is typically not an efficient (or even feasible) means of performing DMD (as discussed in Tu et al. [161]). This is especially true when $n \gg m$, which is often the case in high-dimensional fluids systems. Instead, since \mathbf{Y} and $\mathbf{Y}^\#$ have rank at most $\min(m, n)$, it is typically more efficient to first project the data onto a subspace that is (at most) of this dimension. One way to do this is by projecting the original snapshots onto the POD modes of the data, which is implicitly done in all DMD algorithms. Note that the POD modes of \mathbf{Y} are the columns of \mathbf{U} in the singular value decomposition $\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$ (though typically POD is performed after first

subtracting the temporal mean of the data, which is not done here). We present here a typical algorithm to compute DMD, that is most similar to that proposed in [161] as *exact* DMD.

Algorithm 1 (DMD).

1. Take the reduced singular value decomposition (SVD) of \mathbf{Y} , letting $\mathbf{Y} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$.
2. (Optional) Truncate the SVD by only considering the first r columns of \mathbf{U} and \mathbf{V} , and the first r rows and columns of $\mathbf{\Sigma}$ (with the singular values ordered by size), to obtain \mathbf{U}_r , $\mathbf{\Sigma}_r$, and \mathbf{V}_r .
3. Let $\tilde{\mathbf{A}} := \mathbf{U}_r^* \mathbf{Y}^\# \mathbf{V}_r \mathbf{\Sigma}_r^{-1}$.
4. Find the eigenvalues μ_i and eigenvectors w_i of $\tilde{\mathbf{A}}$, with $\tilde{\mathbf{A}}w_i = \mu_i w_i$.
5. Every nonzero μ_i is a DMD eigenvector, with a corresponding DMD mode given by $\varphi_i := \mu_i^{-1} \mathbf{Y}^\# \mathbf{V}_r \mathbf{\Sigma}_r^{-1} w_i$.

This method is similar to the original formulation of DMD [134], but for the fact that in step 5 the DMD modes are no longer restricted to lie within the column space of \mathbf{Y} . We also explicitly provide the optional step of truncating the SVD of \mathbf{Y} , which might be done if the system is known to exhibit low dimensional dynamics, or in an attempt to eliminate POD modes that contain only noise. We note that this is not the only means to reduce the dimension of the identified system dynamics, nor is it necessarily optimal. Indeed, Wynn et al. [176] develop a method that optimizes the projection basis in parallel while performing a DMD-like eigendecomposition. Jovanović et al. [78] take a different approach, seeking a small number of nonzero modes from the full eigendecomposition that best approximate the system dynamics. An empirical comparison between these various dimensionality-reduction techniques will be given in Section 6.3.3. Note that the continuous eigenvalues λ_{ci} of the system are related to the discrete time eigenvalues identified via DMD via $\lambda_{ci} = \log(\mu_i)/\Delta t$. The growth rate γ_i and frequency ω_i associated with DMD mode φ_i are then given by $\lambda_{ci} = \gamma_i + i\omega_i$.

The matrix $\tilde{\mathbf{A}}$ is related to \mathbf{A} in Equation (2.1.6) by $\tilde{\mathbf{A}} = \mathbf{U}_r^* \mathbf{A} \mathbf{U}_r$. While \mathbf{A} can be viewed as an approximating linear propagation matrix in \mathbb{R}^n (i.e., the space of original data vectors), $\tilde{\mathbf{A}}$ is the equivalent propagation matrix in the space of POD coefficients, which we will sometimes refer to as POD space. Another interpretation of $\tilde{\mathbf{A}}$ is that it is the spatial correlation matrix between the POD modes \mathbf{U}_r , and the same POD modes shifted by the assumed dynamics \mathbf{A} [134]. If we let $\tilde{\mathbf{x}}_k = \mathbf{U}_r^* \mathbf{x}_k$ be the representation of a given snapshot \mathbf{x}_k in the POD basis and let $\tilde{\mathbf{Y}} = \mathbf{U}_r^* \mathbf{Y}$ and $\mathbf{Y}^\# = \mathbf{U}_r^* \mathbf{Y}^\#$, then it is easy to verify that the equivalent of Equation (2.1.6) in POD space is

$$\tilde{\mathbf{A}} = \tilde{\mathbf{Y}}^\# \tilde{\mathbf{Y}}^+. \quad (2.1.7)$$

Equation (2.1.7) will be useful for the analysis performed in this thesis, particularly Chapter 6.

As well as the inherent empirical usefulness of being able to extract spatial modes corresponding to certain frequencies and growth or decay rates, DMD has the potential to give a

deep understanding of the system dynamics, due to a connection with the Koopman operator [129]. For a discrete dynamical system $\mathbf{x}_{k+1} = F(\mathbf{x}_k)$, suppose we have scalar-valued observables $g(\mathbf{x})$, which reside in some function space \mathcal{G} . The Koopman operator \mathcal{K} is defined to act on this space of observables, such that we have

$$\mathcal{K}(g(\mathbf{x}_k)) = g(F(\mathbf{x}_k)) = g(\mathbf{x}_{k+1}). \quad (2.1.8)$$

In other words, \mathcal{K} maps from a scalar-valued function of the state at a given time to the value of that function one timestep in the future. The main advantage of studying a dynamical system through consideration of the Koopman operator is that, even for nonlinear dynamical systems, the Koopman operator is linear. The catch is that, instead of a working in a finite dimensional space, \mathcal{K} is infinite dimensional, due to the fact that the function space \mathcal{G} is (uncountably) infinite-dimensional. While any numerical approximation of \mathcal{K} will necessarily be finite-dimensional, one might hope to at least obtain accurate estimations of its eigendecomposition. For DMD, it is implicitly assumed that the observable functions are simply the data collected at each spatial location. Extended DMD (EDMD) [172] considers the case where these observables may be defined as any transformations of the data, with the idea being to use a richer set of observable functions to find a better finite-dimensional approximation to \mathcal{K} . For further discussion of DMD, EDMD, and connections to the Koopman operator, see Rowley et al. [129], Mezić [99], Tu et al. [161], Williams et al. [172], and Rowley and Dawson [123]. EDMD will be used and discussed in more detail in Chapter 5.

2.1.4 Eigensystem realization algorithm

The eigensystem realization algorithm (ERA) is a method that extracts a linear state space model from impulse response data. As mentioned in section 1.2, it was first conceived for analyzing the structural dynamics of spacecraft in Juang and Pappa [79], but also shares close similarities with a number of previously proposed techniques [e.g., 72, 89]. More details about a range of similar methods and their potential applications can be found in [163, 164, 119].

The output of ERA is a discrete-time linear state-space system, taking the form.

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{A}_d \mathbf{x}_k + \mathbf{B}_d \mathbf{u}_k \\ \mathbf{y}_k &= \mathbf{C}_d \mathbf{x}_k + \mathbf{D}_d \mathbf{u}_k, \end{aligned} \quad (2.1.9)$$

The ERA algorithm proceeds as follows:

Algorithm 2 (ERA).

1. *Collect output data from an impulse response of the form $\{\mathbf{y}_0, \mathbf{y}_P, \mathbf{y}_{2P}, \dots, \mathbf{y}_{mP}\}$ and $\{\mathbf{y}_1, \mathbf{y}_{P+1}, \mathbf{y}_{2P+1}, \dots, \mathbf{y}_{mP+1}\}$.*
2. *Assemble the block Hankel matrices*

$$\mathbf{H} = \begin{bmatrix} \mathbf{y}_0 & \mathbf{y}_P & \mathbf{y}_{2P} & \cdots & \mathbf{y}_{m_c} \\ \mathbf{y}_P & \mathbf{y}_{2P} & \mathbf{y}_{3P} & \cdots & \mathbf{y}_{(m_c+1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{y}_{m_o P} & \mathbf{y}_{(m_o+1)P} & \mathbf{y}_{(m_o+2)P} & \cdots & \mathbf{y}_{(m_o+m_c)P} \end{bmatrix},$$

$$\mathbf{H}^\# = \begin{bmatrix} \mathbf{y}_1 & \mathbf{y}_{P+1} & \mathbf{y}_{2P+1} & \cdots & \mathbf{y}_{m_c P+1} \\ \mathbf{y}_{P+1} & \mathbf{y}_{2P+1} & \mathbf{y}_{3P+1} & \cdots & \mathbf{y}_{(m_c+1)P+1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{y}_{m_o P+1} & \mathbf{y}_{(m_o+1)P+1} & \mathbf{y}_{(m_o+2)P+1} & \cdots & \mathbf{y}_{(m_o+m_c)P+1} \end{bmatrix},$$

where m_c and m_o are chosen such that $m_c + m_o \leq m$.

3. Compute the (reduced) SVD $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$.
4. Truncate the SVD by only considering the first r columns of \mathbf{U} and \mathbf{V} , and the first r rows and columns of $\mathbf{\Sigma}$ (with the singular values ordered by size), to obtain \mathbf{U}_r , $\mathbf{\Sigma}_r$, and \mathbf{V}_r , where r is the desired model order.
5. The matrices of the reduced-order model of a system with p inputs and q outputs are given by

$$\begin{aligned} \mathbf{A}_r &= \mathbf{\Sigma}_r^{-1/2} \mathbf{U}_r^T \mathbf{H} \mathbf{V}_r \mathbf{\Sigma}_r^{1/2}, \\ \mathbf{B}_r &= \text{the first } p \text{ columns of } \mathbf{\Sigma}_r^{1/2} \mathbf{V}_r^T, \\ \mathbf{C}_r &= \text{the first } q \text{ rows of } \mathbf{U}_r \mathbf{\Sigma}_r^{1/2}, \\ \mathbf{D}_r &= \mathbf{y}_0. \end{aligned} \tag{2.1.10}$$

Note that when $p = 1$, the data pairs in step 1 can just be taken from an impulse response sequence with its last and first entries removed. This more general formulation allows for the skipping of data when assembling \mathbf{H} , which can reduce computational costs, while still allowing data to be used across a large total time window.

In general, input-output data might not be available in the form of an impulse response, in which case other more general subspace methods may be used (e.g., Verhaegen and Dewilde [163]). Another approach is to use a technique such as observer/Kalman filter identification [82] to compute an impulse response from input-output data with random inputs, before applying ERA. We note that there are close similarities between DMD and ERA. Indeed, the two algorithms become equivalent when $m_0 = 1$, in the sense that the \mathbf{A} matrices from either method are only different by a similarity transform [161].

2.2 Motivating results

In this section, we present some preliminary results that will serve as motivation for the research in the following chapters. The main idea will be to use ERA to identify models for a pitching airfoil system, and use these models to design feedback systems that allow for the lift to be controlled. In other words, a controller will be used to determine the airfoil kinematics required to achieve a desired (possibly time-varying) lift coefficient. Since these results are largely provided for motivation, we omit a number of technical details from this section. The core methodology behind these results is given in Brunton et al. [24]. For brevity, we also defer a detailed discussion of the experimental and numerical methods to Chapters 3 and 4 respectively.

We consider an airfoil undergoing simple pitching motion about the quarter chord, with the system input being the kinematics of the airfoil (captured by the angular acceleration, $\ddot{\alpha}$)

and the output the lift coefficient, $C_L = \frac{F_L}{0.5\rho AU_\infty^2}$. Models (in the form of low order state-space realizations of the system) are identified by applying the eigensystem realization algorithm to discrete-time impulse response data. This occurs after first extracting the components of the lift that are directly proportional to the angle of attack and its derivatives (α , $\dot{\alpha}$, and $\ddot{\alpha}$), described in further detail in [24]. For direct numerical simulations, such impulse response data is directly simulated, while in experiments it is acquired by applying the observer/Kalman filter identification algorithm (OKID) [82] to the input/output data from pseudo-random, frequency rich maneuvers. Feedback controllers are designed using H_∞ loopshaping, with a controller designed to meet a desired “loopshape” (i.e., the transfer function describing the dynamics of the closed-loop system). See Skogestad and Postlethwaite [143] for further details about such methods.

This procedure can be applied using both theoretical models (e.g., the Theodorsen model, Equation 1.1.1) and the previously discussed reduced-order models, identified in both direct numerical simulations (DNS) and wind tunnel experiments. Simulations are performed on a two-dimensional flat plate airfoil at a Reynolds number of 100, using an immersed boundary projection method [149, 36]. Experiments were conducted using a NACA 0006 airfoil of chord length 0.245 m with a free stream velocity of 3 m/s, giving a Reynolds number of approximately 50,000. Gusting conditions are generated in the wind tunnel by using a series of shutters downstream of the test section. In spite of the differences in parameters and conditions between the DNS and experimental work, we demonstrate that the same control methodology can be successfully applied in both cases.

Feedback control is implemented for tracking reference lifts of a range of magnitudes, both with and without the presence of gusting disturbances. To begin with, we consider tracking step changes in reference lift in DNS. Figure 2.1 shows accurate tracking of the desired lift over a range of magnitudes, even significantly beyond the maximum value that can be held in steady conditions ($C_L = 0.97$, for this system). This highlights one of the major benefits of using feedback control: even though the system is clearly non-linear, we are able to make the output (which is often what we care most about) behave linearly. Here the nonlinear effects are compensated for by the controller modifying the input to the system. Figure 2.2 shows lift tracking step responses for experimental conditions. The presence of noise and time lags in the system degrades the performance. Nonlinear effects further limit the performance for high amplitude steps.

Having validated that feedback control can be successfully implemented on simple step maneuvers, we proceed to investigate more complex desired lift profiles. Figure 2.3 shows the performance of the DNS system in tracking a sinusoidally varying reference lift. Interestingly, we note that the addition of a periodic component to the reference lift at certain frequencies allows for successful tracking of a higher average lift.

While these results show one use for reduced-order models, they also suggest a few limitations of such an approach. Firstly, accurate lift tracking was only possible when an accurate, real-time lift measurement was available. This is primarily due to the fact that the system is nonlinear. This motivates the development of nonlinear modeling procedures considered in Chapters 3 and 5.

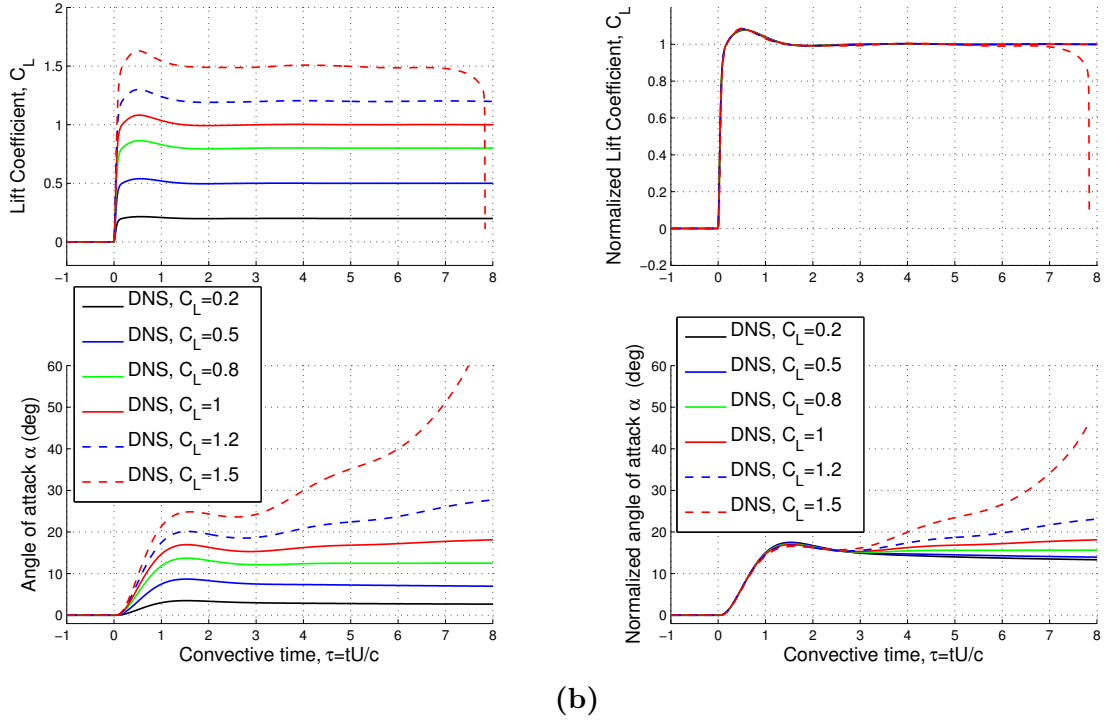


Figure 2.1: **(a)** Step responses of various magnitude for the closed-loop system in DNS, and **(b)** normalized step responses. Linear behavior in the output (lift coefficient) is observed, despite large nonlinearities present in the system, which is evident from the variation in the angle of attack plots.




Figure 2.2: Step responses of various magnitude for the closed-loop system in wind tunnel experiments. The controller performance degrades for high-amplitude steps.

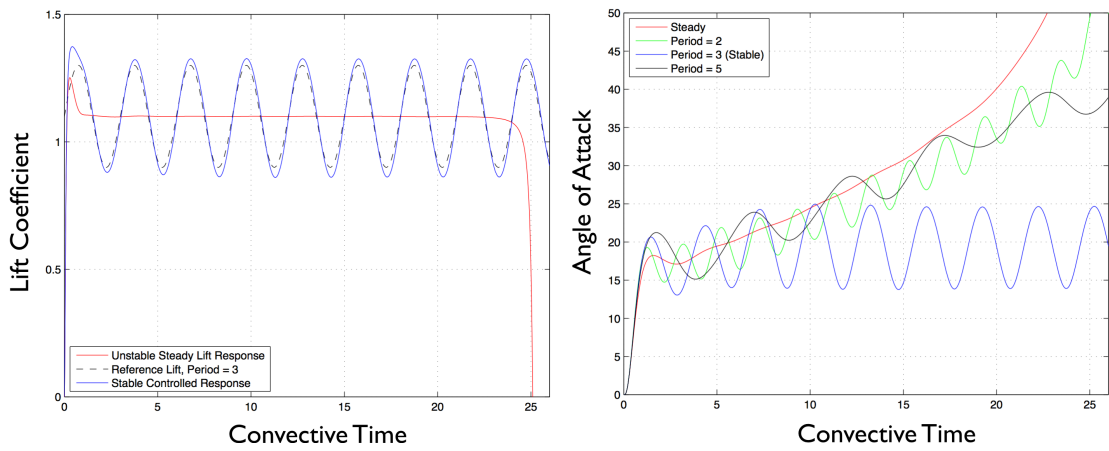


Figure 2.3: Controller performance in tracking sinusoidally varying reference lift coefficient.

Figure 2.3 exposes some interesting flow physics for this particularly pitching airfoil system, where higher average lift coefficients are found when pitching at a preferred frequency (corresponding to oscillations of period $\frac{tU}{c} = 3$). This phenomena is explored systematically in Chapter 4.

We also observe a degradation of performance when using experimental, rather than numerical data. Aside from issues with implementing feedback control in experimental systems that contain time lags between command and actuation, one cause of this degradation is the presence of noise in the data. This is particularly important when using data-driven modeling approaches. We study the presence of noise in the dynamic mode decomposition in Chapter 6, and propose modified algorithms that can give improved performance with noisy data.

Chapter 3

A data-driven modeling framework for predicting forces and pressures on a rapidly pitching airfoil

This chapter formulates a switched linear modeling procedure to understand and predict the unsteady aerodynamic forces arising from rapid pitching motion of a NACA 0012 airfoil, at a Reynolds number of 50,000. The system identification procedure applies a generalized dynamic mode decomposition algorithm to time-resolved wind tunnel measurements of the lift and drag forces, as well as the pressure at six locations on the suction surface of the airfoil. Linear state space models are identified for 5° pitch-up and pitch-down maneuvers within an overall angle of attack range of 0° – 20° . The identified models accurately capture the effects of flow separation and leading-edge vortex formation and convection. It is shown that switching between different linear models can give accurate prediction of the nonlinear behavior that is present in high-amplitude maneuvers. The models are accurate for a wide range of motions, including pitch-and-hold, sinusoidal, and pseudo-random pitching maneuvers. Providing the models access to a subset of the measured data channels can allow for improved estimates of the remaining states via the use of a Kalman filter, which could be of use for aerodynamic control applications.

3.1 Introduction

Airfoils undergoing pitching motion can exhibit a range of distinct features, including boundary layer separation and reattachment, Kelvin-Helmholtz shear layer instabilities, dynamic stall vortex formation and detachment, and von Kàrmàn wake instabilities. Indeed, the dynamics of the system at low angles of attack, where the flow may be fully attached to the upper surface of the airfoil, often bears little resemblance to the behavior at high angles of attack, where the airfoil can more closely resemble a bluff body, with a large recirculation region and periodic vortex shedding. For this reason, if one were to attempt to formulate a reduced-order model that captures the dynamics of a pitching airfoil restricted to one of these regimes, not only would the resulting models have substantial differences, but the entire modeling approach chosen could be fundamentally different. For example, a typical

model at small angles of attack might be derived using classical unsteady aerodynamic theory as discussed in Section 1.1, whereas a high angle-of-attack model might take the form of a Goman-Khrabrov model [60], for example. Such a model is motivated by the movement of the separation point on the airfoil undergoing (relatively) slow pitching, and thus only makes sense flow which is at least partially separated. Indeed, such a model was formulated for a system similar to that considered in this chapter (though with a narrower range of motions and angles of attack than considered here) in Williams et al. [171]. Therefore, it can be challenging to find a reduced-order modeling framework that is both sufficiently simple to explicitly formulate and rapidly simulate, and general enough to capture the full range of phenomena that emerge under different conditions. As discussed in Section 1.1, this complexity can motivate the use of data-driven modeling methods, which allow for the identification of dynamical systems that accurately capture the features present in the data, and are generalizable to a wide range datasets and systems of varying complexity.

The overriding goal of this chapter is to formulate a modeling framework that allows for the prediction of the system dynamics across wide range of angles of attack, ranging between fully attached and fully separated flow, for arbitrary pitching maneuvers. For the purposes of this work, we specifically seek accurate prediction of the time-varying forces and pressures incident on an airfoil as it is pitching. One downside of using data-driven methods is that one can lose or bypass the physical understanding that explains the behavior of the system. For this reason, we make an effort to not just formulate models that accurately represent the data itself, but also to interpret the dominant features that such models are representing.

The modeling approach used in this chapter utilizes a recently-proposed variant of dynamic mode decomposition, which allows for the identification of systems with inputs [116]. One advantage of the present modeling approach is that, unlike those generated using the eigensystem realization algorithm, for example, the model states can be directly related to measurements. This can allow easy switching between neighboring linear models, which subsequently permits the formation of a switched linear model that is capable of predicting nonlinear behavior. The algorithm is described in Section 3.2, which is followed by a description of the experimental setup in Section 3.3. Section 3.4 demonstrates that the obtained models are accurate for a range of high-amplitude pitching maneuvers. Section 3.5 contains a more general discussion of the results and subsequent conclusions of this study. Section 3.6 contains appendices to this chapter, which consists of more detailed results pertaining to the forces and pressures on a stationary airfoil, as well as a short description of the Kalman filter used in Section 3.4.

3.2 System identification method

We use a variant of dynamic-mode decomposition to identify models describing the pitching airfoil, which is briefly described here. The goal will be to identify a family of linear discrete-time systems of the form

$$\mathbf{x}_{i+1} = \mathbf{A}_k \mathbf{x}_i + \mathbf{B}_k \mathbf{u}_i, \quad (3.2.1)$$

where \mathbf{x}_{i+1} is the state of the system one timestep after the system was in state \mathbf{x}_i . By identifying different \mathbf{A}_k and \mathbf{B}_k matrices for different angles of attack and directions of pitching, we can assemble a family of linear models $\{\mathbf{A}_k, \mathbf{B}_k\}_{k=1}^N$ that can subsequently

be pieced together to allow for accurate prediction of maneuvers spanning a wider range of angles of attack than any single linear model would be capable of. We now describe in general terms the identification procedure for a linear model $\{\mathbf{A}, \mathbf{B}\}$. Further details concerning the specific data chosen to constitute the state \mathbf{x} will be given in Section 3.4.1.

Suppose we collect a time-series of measurements \mathbf{x}_i , which we assemble as columns into a matrix \mathbf{Z} , as described in Section 2.1.3. From \mathbf{Z} , we select all pairs of data that are separated by some nominal time Δt , which we assemble into matrices \mathbf{Y} and $\mathbf{Y}^\#$. If \mathbf{Z} consists of uniformly sampled data, then \mathbf{Y} and $\mathbf{Y}^\#$ are \mathbf{Z} with the last and first columns removed, respectively. Standard DMD can be characterized as finding the eigendecomposition of a matrix \mathbf{A} satisfying (or approximately satisfying) $\mathbf{Y}^\# = \mathbf{A}\mathbf{Y}$ [161]. Depending on the size of \mathbf{Y} and $\mathbf{Y}^\#$, \mathbf{A} is either the (Frobenius) minimum-norm solution (if the data matrices have more rows than columns), or the least-squares solution (otherwise). The usefulness and validity of this approach relies assumption that the system is autonomous, and not greatly affected by external inputs. If we have known inputs \mathbf{u}_i assembled into a matrix \mathbf{U} , then it is possible to modify DMD [116] to instead seek the matrices \mathbf{A} and \mathbf{B} satisfying

$$\mathbf{Y}^\# = [\mathbf{A} \ \mathbf{B}] \begin{bmatrix} \mathbf{Y} \\ \mathbf{U} \end{bmatrix}. \quad (3.2.2)$$

Provided that the size of the state m is not excessive, we may compute the augmented system matrices $[\mathbf{A} \ \mathbf{B}]$ through

$$[\mathbf{A} \ \mathbf{B}] = \mathbf{Y}^\# \begin{bmatrix} \mathbf{Y} \\ \mathbf{U} \end{bmatrix}^+, \quad (3.2.3)$$

where $+$ denotes the Moore-Penrose pseudoinverse. Since fluids systems are, in general, nonlinear, the ability of the identified linear system to accurately model all of the data may be limited. However, an intelligent selection of state variables \mathbf{x} can go a considerable way towards factoring out much of the nonlinearity in the system. To begin with, rather than directly using force and pressure measurements, we can instead consider deviations from the equilibrium (or mean) values at a given angle of attack. This allows for the resulting linear model to be accurate despite nonlinear static behavior.

3.3 Experimental method

Experiments were conducted at the Andrew Fejer wind tunnel at the Illinois Institute of Technology, with a diagram of the airfoil mounting shown in Figure 3.1. A NACA 0012 airfoil of chord length $c = 0.245$ m was used in a test section of length 3 m and cross-section 0.6 m by 0.6 m. The airfoil spanned the width of the test-section, thus minimizing three-dimensional effects. The airfoil was mounted upon a six-axis ATI nano17 force transducer in its spanwise center, which allowed for the measurement of time-resolved forces and moments. This, in turn, was mounted upon two pushrods actuated by Copley servo tubes, allowing for pitching motion to be commanded. For the results presented here, only the rear pushrod was actuated, which resulted in pitching about an axis $0.11c$ from the leading edge. The presence of the airfoil mounting and pushrod mechanism creates a disturbance to the flow, which influences the force measurements in particular. This is the main cause of the nonzero




Figure 3.1: Schematic of the experimental setup.

lift at an angle of attack $\alpha = 0^\circ$, which will be observed in Section 3.4. This effect has been previously documented for this wind tunnel setup [171]. Six pressure taps were installed along the chord of the airfoil at one spanwise location, located at distances of $0.050c$, $0.217c$, $0.385c$, $0.552c$, $0.720c$, and $0.887c$ aft of the leading edge. The freestream velocity U was measured using a pitot tube and remained constant at a nominal value of 3 m/s, giving a Reynolds number $\frac{cU}{\nu}$ of approximately 50,000 (ν is the kinematic viscosity of air), and a convective time $t_c = \frac{c}{U} = 0.0817$ s. Note that some blockage effects meant that the freestream velocity changed by a small amount as the angle of attack changed (approximately 3% when pitching between 0° and 20°). All forces and pressures were nondimensionalized using the averaged velocity at the relevant phase of the maneuver. (Note that this neglects unsteady gusting effects associated with the changing velocity, but since the changes are small, these effects should be negligible.)

Force and pressure data was acquired at a frequency of 1000 Hz. For each maneuver, data was phase-averaged over at least 50 cycles to reduce the effect of measurement noise. All maneuvers were also performed with the wind tunnel off before and after data was collected with the tunnel switched on. These results were also phase-averaged, and subtracted from the tunnel-on data. This eliminates (for the force readings) the effects of the mass of the wing, the added-mass terms associated with accelerating the surrounding air, and also any other effects on the measurement equipment resulting directly from the maneuver performed. By eliminating added-mass terms, we isolate the circulatory fluids forces arising from a given pitching maneuver.

3.4 Results

Here results are presented for the identification (Section 3.4.1) and performance of the suite of identified models. To test the performance of the family of models that have been identified, we analyze their ability to predict a range of other maneuvers. These range from compositions of similar individual maneuvers (Section 3.4.2), to sinusoidal (Section 3.4.3) and pseudo-random (Section 3.4.4) pitching maneuvers. The latter two of classes of maneuver bear little

similarity to the maneuvers used for identification. In this sense, we will be able to show the generality of these models, which highlights that the identified models represent more than simply fits to the data, and have predictive capabilities.

3.4.1 System identification results

Models were identified separately from pitch-up and pitch-down maneuvers between 0° – 5° , 5° – 10° , 10° – 15° , and 15° – 20° , with model states obtained from the 6 pressure readings and the lift and drag measurements. The prescribed maneuvers take the canonical form [51]

$$\alpha(t) = \frac{MG}{\max(G)}, \quad G = \log \left[\frac{\cosh(a(t - t_1)) \cosh(at_2)}{\cosh(a(t - t_2)) \cosh(at_1)} \right]. \quad (3.4.1)$$

Nondimensionalizing time by $\frac{U}{c}$, we take $t_2 - t_1 = 2$, and $a = \frac{10U}{c} = 122.4$ (along with $M = 5^\circ$). With these parameter values, a becomes the main governing parameter that determines the rate of the step.

The duration of the pitch was approximately 4 convective time units ($4\frac{c}{U}$). Static data at the corresponding angles of attack was first subtracted from all measurements, and all data was nondimensionalized (forces by $\frac{1}{2}\rho c U^2$ and pressures by $\frac{1}{2}\rho U^2$). To identify models, the DMD-type algorithm described in Section 3.2 was used, which allows for data with inputs (which in this case was taken to be measurements of either α and $\dot{\alpha}$, or just $\dot{\alpha}$). Using this method, we arrive at models of the form

$$\mathbf{x}_{i+1} = \mathbf{A}\mathbf{x}_i + \mathbf{B}\mathbf{u}_i,$$

as described in Section 3.2. Here we let $\mathbf{x} = [\hat{C}_{p1} \ \hat{C}_{p2} \ \dots \ \hat{C}_{p6} \ \hat{C}_l \ \hat{C}_d]^T$, and $\mathbf{u} = [\alpha \ \dot{\alpha}]^T$. Note that while we treat α as an input for convenience, the fact that it is entirely dependent on $\dot{\alpha}$ means that we could also treat it as an additional system state. Here $\hat{\cdot}$ represents the deviation from an equilibrium condition, $\hat{C}_i = C_i - C_i^e(\alpha)$, where $C_i^e(\alpha)$ is the equilibrium (or mean in the case of an unsteady state) value of the coefficient at a given angle of attack. We first attain this equilibrium data for angles of attack in the range $\alpha \in [0^\circ, 22^\circ]$, which is shown in Figure 3.2. The behavior of the system for fixed angles of attack can also contain significant complexity. We provide further data and analysis of the fixed airfoil configuration in an appendix, Section 3.6.1. To motivate the development of unsteady models, we also show in Figure 3.2 how data acquired for a pitching airfoil deviates from these equilibrium values. Considering just the static data, we observe that the lift coefficient increases close to linearly (with a slope of approximately 1.7π) between 0° and 8° , before the lift curve reaches its peak and then plateaus between 10° and 15° , before again increasing beyond 15° . The lift plateau corresponds to the airfoil stalling, with the flow over the suction surface becoming separated. Further evidence for this comes from examining both the drag curve, which sees a large increase in drag beyond $\alpha = 8^\circ$, and in the first two pressure coefficients, which give a sharp drop in pressure beyond this angle. Prior to full separation, there is evidence for partial separation towards the rear of the airfoil. Pressure sensors 3–6 all measure a drop in




Figure 3.2: Force and pressure coefficient data for a static airfoil at angles of attack between 0° and 22°, as well as for an airfoil sinusoidally pitching between 5° and 15° at a rate $k = \frac{\pi fc}{U} = 0.051$. Pitching data was phase-averaged over 200 cycles, while static data was obtained while keeping the airfoil fixed at angles of attack between 0° and 22°, in 1° increments.

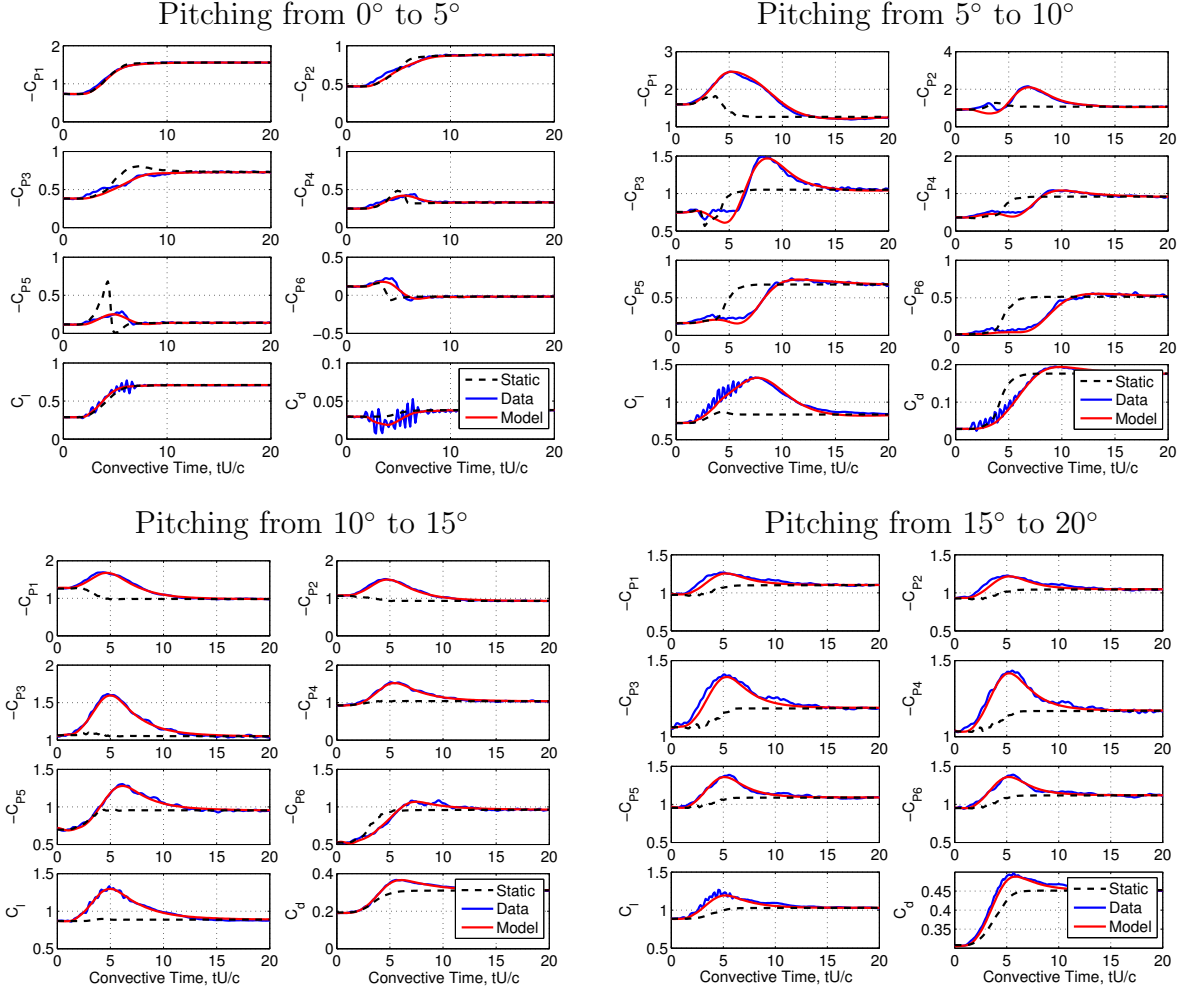


Figure 3.3: Force and pressure data for 5° pitch-up maneuvers for starting angles of 0° , 5° , 10° and 15° , which were used for system identification. In all cases, the identified models accurately replicate the experimental data. Also shown is the static data at the relevant instantaneous angle of attack.

pressure at a critical angle between 2° and 6° , which appears to signify the separation point moving upstream of the given sensor.

Returning now to the system identification procedure, Figures 3.3 and 3.4 show the performance of each model in predicting the pressure and force coefficients for the maneuver upon which they were identified. For reference, the static pressure and force coefficients at the instantaneous angle of attack are also shown in Figures 3.3 and 3.4. Rather than subtracting the full static curves before system identification, we found improved results by assuming linear variation in the static values throughout the maneuver. This avoids issues with separation-related “jumps” occurring at different angles of attack in for the static and moving airfoil, which makes the static-subtracted data less smooth.

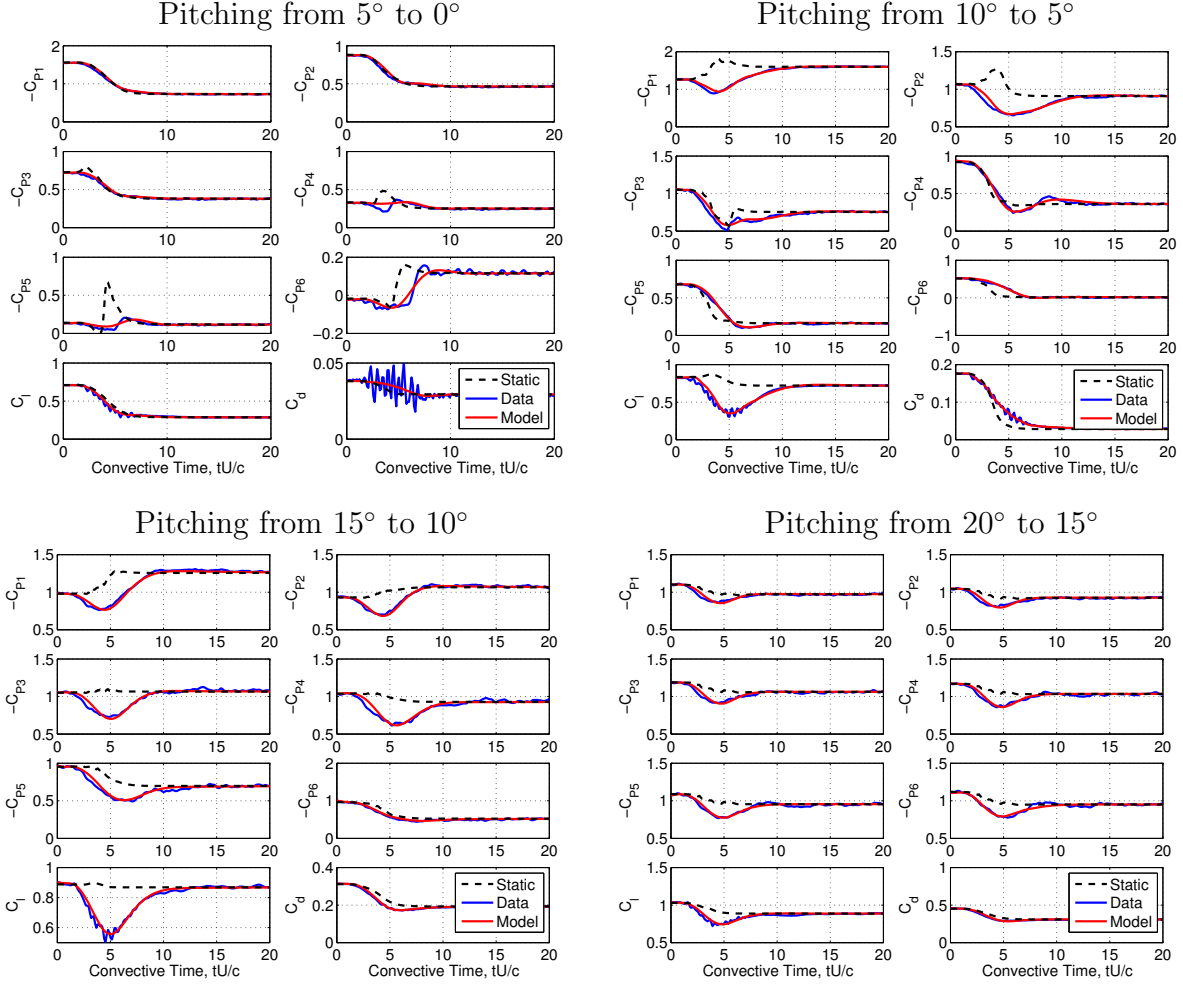


Figure 3.4: Force and pressure data for 5° pitch-down maneuvers for starting angles of 5° , 10° , 15° and 20° , which were used for system identification. In all cases, the identified models accurately replicate the experimental data. Also shown is the static data at the relevant instantaneous angle of attack.

We can identify two dominant features of the pitch-up and pitch-down behavior. A temporary rise in C_l, C_d and $-C_p$ is observed which is consistent with the formation and convection of a leading-edge vortex (see, e.g., all measurements for pitching between 15° and 20°), and a time-lag in reaching the steady state value, most likely due to the boundary layer requiring time to reach its new equilibrium configuration (see, e.g., C_{p6} when pitching between 0° and 5°). We qualitatively summarize the presence of each of these features in Table 3.1.

Maneuver	$-C_{p1}$	$-C_{p2}$	$-C_{p3}$	$-C_{p4}$	$-C_{p5}$	$-C_{p6}$	C_L	C_D
0° to 5°	N	D	D	D	D	D	N	N
5° to 10°	V	VD	VD	VD	VD	VD	V	VD
10° to 15°	V	V	V	V	V	VD	V	V
15° to 20°	V	V	V	V	V	V	V	V
20° to 15°	V	V	V	V	V	V	V	V
15° to 10°	V	V	V	V	V	N	V	N
10° to 5°	V	V	N	N	D	D	-V	D
5° to 0°	N	N	V	D	D	N	N	N

Table 3.1: Qualitative features observed during pitch-up and pitch-down maneuvers for each pressure and force measurement. V refers to effects of a vortex, D refers to a time lag, while N indicates that neither of these effects are significant.

We finally note that we obtain quite different models for pitch-up and pitch-down maneuvers. To show this explicitly, Figure 3.5 shows the inaccuracy of the prediction of a 5° – 10° pitch-up model for a 10° – 5° pitch-down maneuver, which arises primarily because the 5° – 10° pitch-up model predicts the existence of a time-delay, which is not present in the 10° – 5° pitch-down data. This has important implications for the use of pseudo-random system identification maneuvers, which necessarily incorporate both pitching up and pitching down motion.

3.4.2 Multiple pitch-up and pitch-down maneuvers

We now consider a maneuver consisting of two pitch-ups followed by two pitch-downs, each in rapid succession. We attempt to predict the maneuver by switching between the relevant models for each pitch-up and pitch-down. For this maneuver, four different models are used. Given that the state of each model consists of the same variables, this is simply a matter of switching the \mathbf{A} and \mathbf{B} matrices used to propagate the system.

The results for this procedure in predicting C_l, C_d , as well as two of the pressure coefficients are shown in Figure 3.6, where we have considered double pitch-up/down maneuvers between 10° and 20° with different pitching rates. We vary the a parameter from equation 3.4.1 to modify the pitching rate (halving and doubling it from the value used in system identification). In all cases, we switch between sub-models at $t = 10, 20$ and 30 convective times, using the final predicted state from one sub-model as the initial condition for the




Figure 3.5: Example of inaccurate prediction of a pitch-up model ($5^\circ - 10^\circ$ model, red curves) on pitch down maneuver ($10^\circ - 5^\circ$ maneuver, blue curves).

next. To give some basis for comparison, we show the performance of a single linear model (that identified from a $5^\circ - 10^\circ$ pitch-up) in predicting this maneuver in Figure 3.7. We note that the only section of this maneuver that this model accurately predicts is that which is most similar to its identification maneuver.

Figure 3.8 shows a quadruple pitch-up and -down maneuver, which switches between all models. From all of these results, we find that switching between models generally works well, though sometimes it can induce “jumps” immediately after switching, particularly when switching between the pitch-up and pitch-down models between 5° and 10° . It is possible that these could be eliminated or reduced with further refinements to the system identification and/or switching procedure.

3.4.3 Sinusoidal pitching

Next, we consider high-amplitude sinusoidal pitching maneuvers, pitching between 0° and 20° at rates $f = 0.2$ Hz and 0.4 Hz, giving a reduced frequencies $k = \frac{\pi f c}{U} = 0.051$ and 0.103 . In Figure 3.9 we show the predicted pressures and forces when using a single model (arbitrarily taken to be the pitch-up model from 5° to 10°), a switched model, and a switched equipped with a Kalman filter (that gives access to the first and sixth pressure measurements). Details concerning the design of this Kalman filter are given in an appendix, Section 3.6.2. We find that the switched model performs better than any single linear model, and that improved accuracy in all measurements can be achieved when using the Kalman filter. The latter observation demonstrates that, even if the models themselves have some inaccuracies in

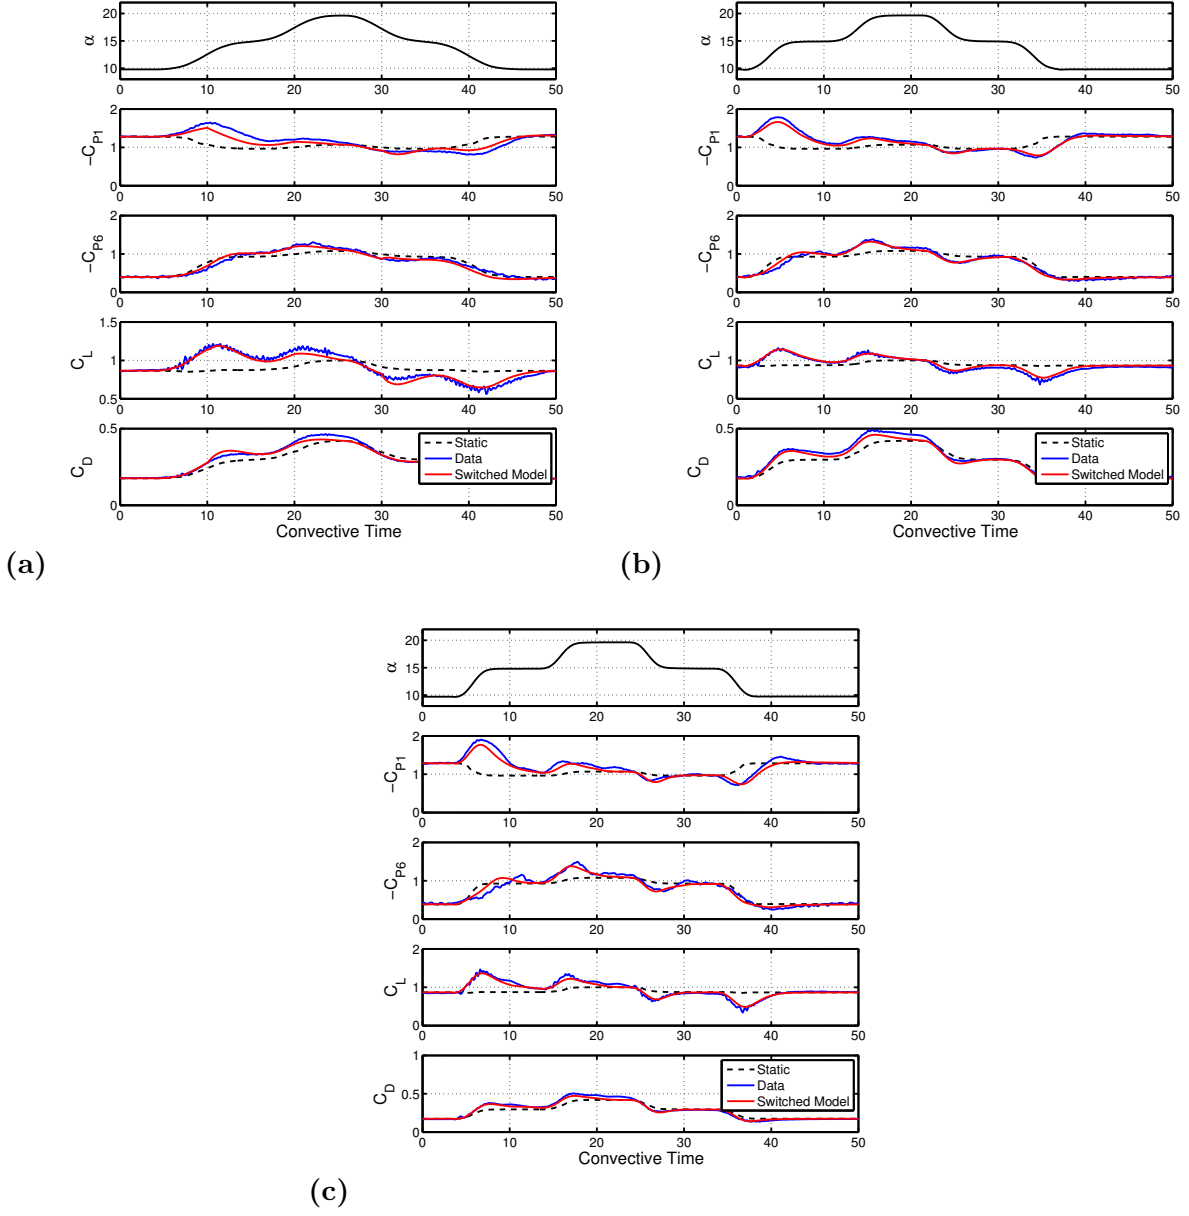


Figure 3.6: Switched model performance in predicting pressure and force coefficients for double pitch up/down maneuvers with different pitch-rates (increasing from (a) to (c)), between 10° and 20° . Switching between sub-models occurs at $t = 10, 20$ and 30 convective times, using the final predicted state from one sub-model as the initial condition for the next. Subplot (b) uses the same pitch-rate as the maneuvers used for system identification.

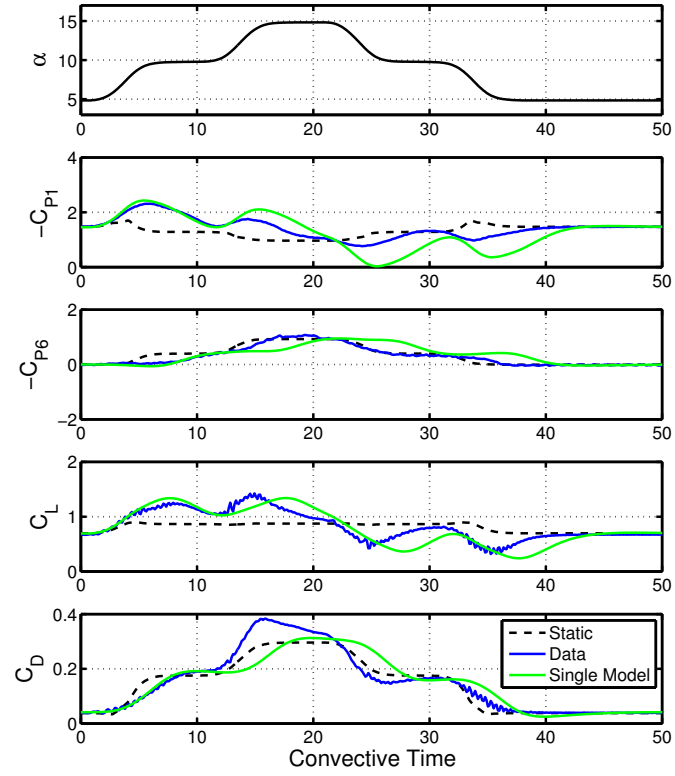


Figure 3.7: Performance of a single (5° – 10° pitch-up) model in predicting pressure and force coefficients for double pitch up/down maneuver.




Figure 3.8: Performance of switched model in predicting pressure and force coefficients for quadruple pitch up/down maneuver. Switching between models occurs at multiples of 10 convective times between $t = 10$ and $t = 70$, where the final predicted state from one sub-model as the initial condition for the next.

predicting the outputs, access to measurements of a subset of these outputs can improve the prediction of all outputs. This is relevant for the use of such models for real-time control, where, for example, we may seek to attain a desired lift force using only pressure measurements.

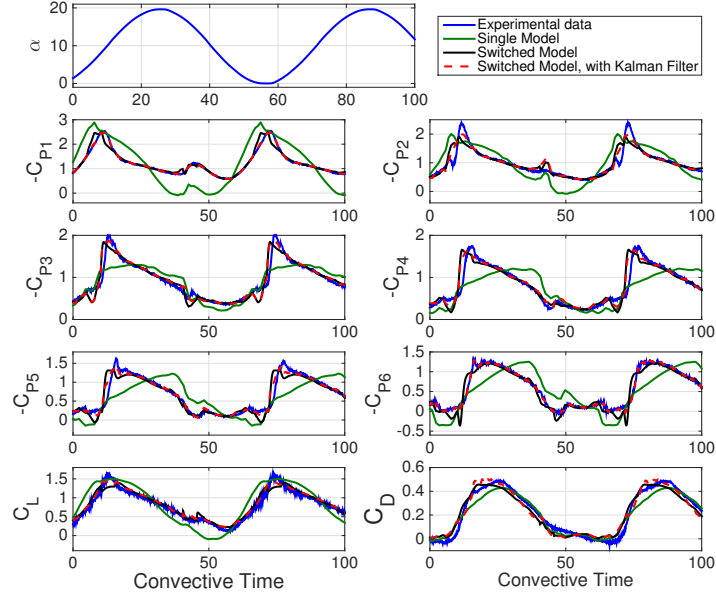
3.4.4 Pseudo-random pitching

We finally consider the case where the angle of attack varies in a pseudo-random manner. Figure 3.10 shows the performance of a switched model equipped with a Kalman filter in predicting the pressures and forces for a pseudo-random pitching maneuver. Again, we observe close agreement between the measured and predicted results. This close agreement highlights the full generality of the switched model, as it is capable of accurately predicting the behavior of the airfoil forces and pressures for arbitrary high-amplitude pitching motions.

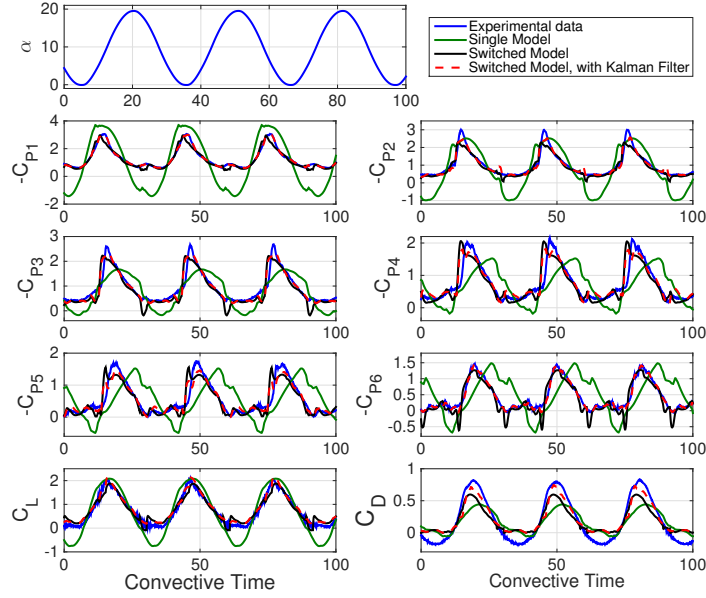
3.5 Discussion and conclusions

The results presented in Section 3.4 demonstrate that the system identification technique described in Section 3.2 can be of use for unsteady aerodynamic modeling applications. The fact that accurate models were attained from very simple pitch-up and pitch-down maneuvers gives the procedure an advantage over the OKID algorithm, which typically requires a concatenation of a variety of motions to obtain accurate models [23]. The absence of internal states in the resulting models mean that they are naturally suited for piecing together for the formation of a global switched model. This process is difficult for ERA models, where the internal states are not directly associated with physical measurements. Having measurements directly associated with model states means that the dimension of the observables must be at least as large as the dimension of the underlying dynamics (or their approximating model), though this restriction could be relaxed if we were to concatenate the data with time-shifted measurements (as is done in ERA), or by using transformations of the original data [172]. Conversely, the fact that the models are accurate suggests that 8th order linear models are sufficient to capture the phenomena present in the maneuvers considered. Indeed, in many cases it was found that it was possible to apply balanced truncation to reduce the dimension of the identified models without significant degradation of predictive accuracy.

In general, linear modeling techniques are appealing due to the simplicity of their identification and formulation, and the ease of use in simulation and controller design. Their accuracy in the prediction of nonlinear dynamics, however, will typically be fundamentally limited to a region in phase-space that is near to the identification maneuver. Scheduling between a family of linear models can go some way to incorporating nonlinear effects into a global model, thus increasing the region in phase-space where such models are accurate. This work demonstrated that, in the α direction of phase-space, such an approach can work between 0° and 20°, which includes the regimes where the flow over the suction surface is completely attached, partially separated, and fully separated. This range was the maximum



(a)



(b)

Figure 3.9: Predicted and actual pressure and force coefficients for high-amplitude sinusoidal pitching at dimensionless frequencies $k = \frac{\pi fc}{U} = 0.051$ and 0.103 . Predictions are made using a model both with and without a Kalman filter. When a Kalman filter is used, the model is given access to the pressure measurements nearest and furthest from the leading edge (C_{p1} and C_{p6}).

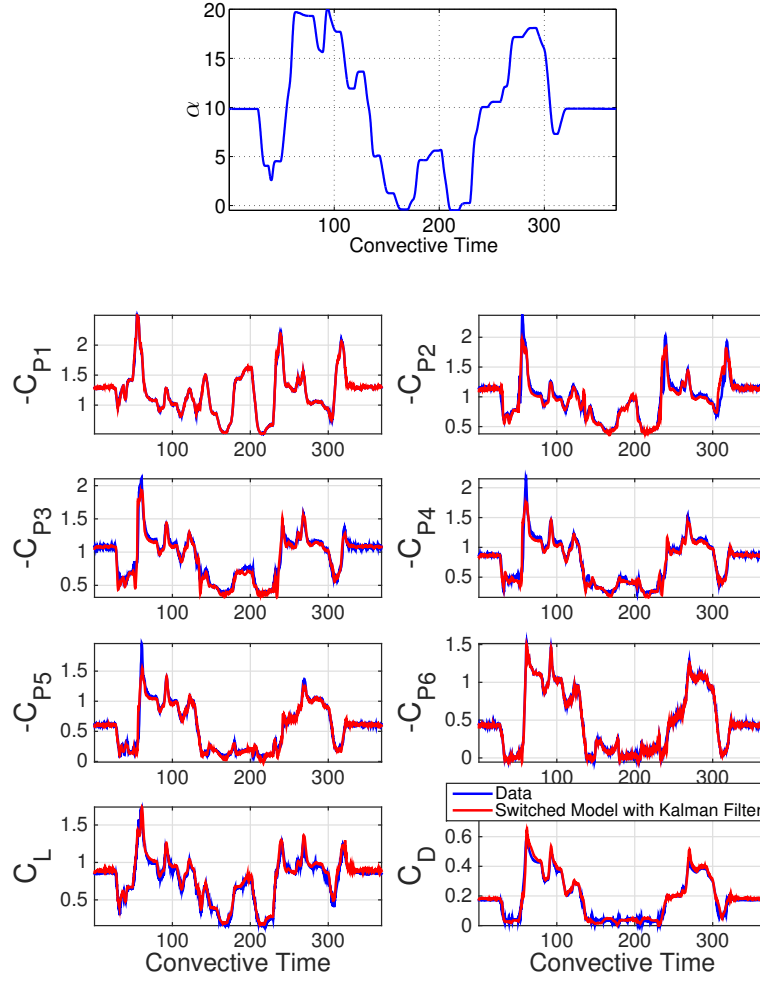


Figure 3.10: Performance of the switched model, equipped with a Kalman filter and measurements of the first and last pressure coefficient, in predicting pressure and force coefficients for high-amplitude pseudo-random pitching maneuver.

available given physical limitations of the airfoil mounting, but we imagine that separated flows at higher angles of attack should also be able to be accurately modeled, given that they are phenomenologically similar (in terms of being fully separated) to those near 20° .

In terms of the applicability of the model for maneuvers with different pitch-rates (i.e., the $\dot{\alpha}$ direction of phase space), we see in Figure 3.6 that the models remain accurate for a range of values of pitch rates. Looking at Figure 3.9, however, it is observed that switching between models while pitching at a relatively fast rate can lead to some degradation of model accuracy. This is a known limitation of gain scheduling models in general [139].

The fact that the acquired data was phase-averaged over a number of cycles means that any unsteady phenomena that are not phase-locked with the pitching motions will be averaged out of the identified models. Particularly for separated boundary layers, such effects (which can occur on a faster timescale to the pitching motions) can be significant, even if they are not directly controllable by pitching motion. Further work could, for example, incorporate such dynamics into state estimators, which could improve the real-time predictive power of such models.

The data that is obtained for the cases of a pitching and stationary airfoil is also of fundamental fluid mechanical interest, which will could further investigated by investigating the time-varying velocity field using particle image velocimetry (PIV). Specifically, it would be interesting to explore whether a small number of measurements could be used to accurately predict not only the pressures and forces (as was done in the present work), but also the entire velocity field in the vicinity of the airfoil.

3.6 Appendices

3.6.1 Further static airfoil results and analysis

The system identification method used in this chapter (described in Section 3.2) relies on an accurate knowledge of the characteristics of the stationary airfoil across the range of angles of attack considered in this investigation. In Section 3.4.1, the general behavior of the lift and pressure curves was analyzed, with the lift plateau pressure drop related to the increasing amounts of flow separation on the suction surface of the airfoil as the angle of attack increases. To analyze this more quantitatively, we show in Figure 3.11 the angle of attack at which the (negative) pressure at each location on the airfoil reaches its first local maximum. For angles of attack beyond this critical angle, we reason that the flow is at least partially separated.

To better understand the flow physics involved in the separation of the flow, we perform a spectral analysis of the data collected at constant angles of attack. Figures 3.12 and 3.13 show the spectra (power spectral density) for each output for the wind tunnel off and on, respectively. With the tunnel off, the plots are very similar for each angle of attack, so only one spectrum is shown (being the average of tunnel-off spectra at all angles of attack). We highlight a few important frequencies in the spectra. Figure 3.12 shows a peak at a frequency of 13.2 Hz (as well as its first harmonic) for all pressure measurements, which might be due to measurement noise in the output voltages. The force and angle measurements (again with

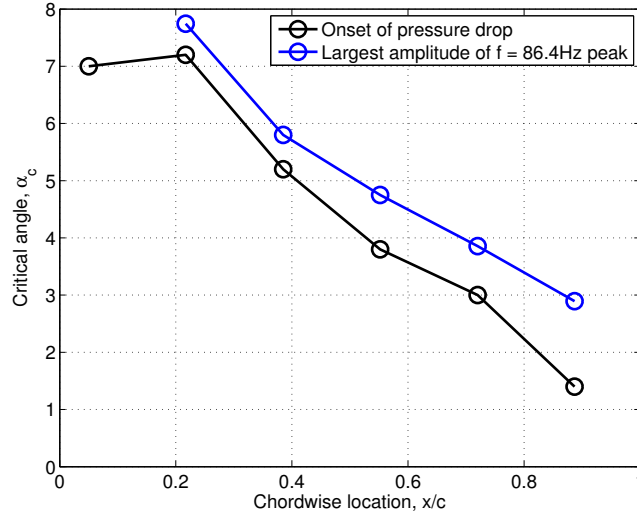


Figure 3.11: Location of static separation point (determined by pressure maxima) as function of α

the tunnel off) have peaks at 7.3 Hz and 37.6 Hz. Since these peaks (as well as some smaller ones in between the two frequencies) are the same for the measured angles and forces, it is likely that that they are due to mechanical vibration of the pushrod servos.

With the tunnel on, there are numerous additional features worth mentioning. A peak at 7.8 Hz appears in all measurements, and there is also see a peak at 3.8 Hz in some of the data. One might think that these correspond to the same resonance peak identified in the tunnel-off measurements. However, there are a couple of reasons to suggest that these are fluids related: they do not appear in the measured angle of attack spectra, nor does the 3.8 Hz peak appear when the tunnel is turned off, and this peak only seems to appear at higher angles of attack (when the surrounding frequencies also have larger spectral densities). Indeed, it seems that this frequency exhibits period-doubling behavior beyond a certain angle of attack.

We can estimate numerous characteristic frequencies related to fluids effects. Beyond a certain critical angle of attack, the flow should exhibit periodic vortex shedding. This should occur at a frequency corresponding to a Strouhal number ($\frac{f c \sin(\alpha)}{U}$) of 0.15–0.2, which gives a frequency range (for angles between 10° and 20°) of 5–14 Hz. In general, there is larger energy at all frequencies less than about 15 Hz for higher angles of attack, which is due to unsteadiness in the separated flow over the suction surface.

At lower angles of attack, the wake may also be unsteady, with the shear layer susceptible to Kelvin-Helmholtz instabilities [178, 177, 87]. These effects have been reported for similar airfoils and similar (low) Reynolds numbers, and occur at higher frequencies than that of the von Kármán vortex street. The frequencies for this instability is more sensitive to Reynolds number than the vortex shedding frequency. Note that Yarusevych et al. [178] also observes

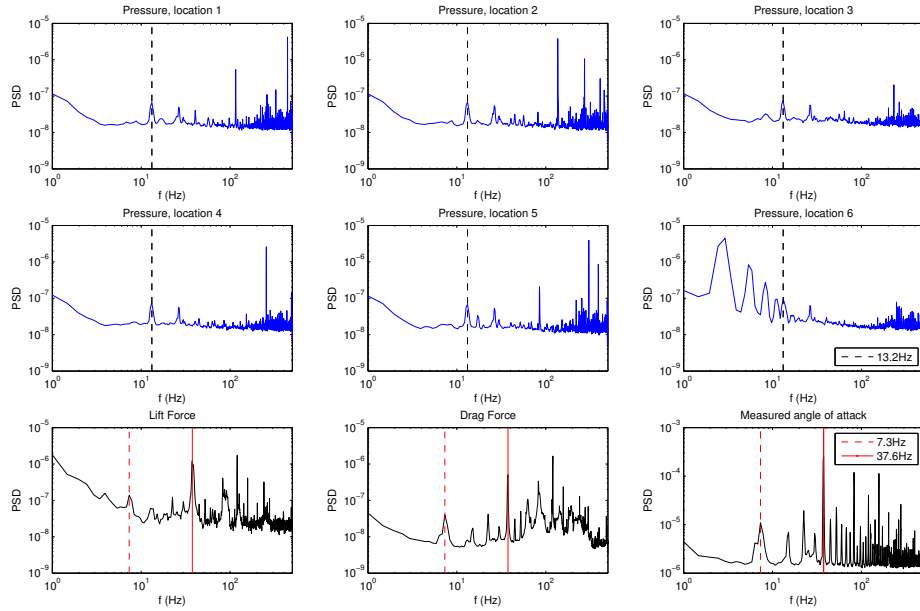


Figure 3.12: Power spectral densities of the pressures, forces, and angle of attack measurements made with the tunnel turned off.

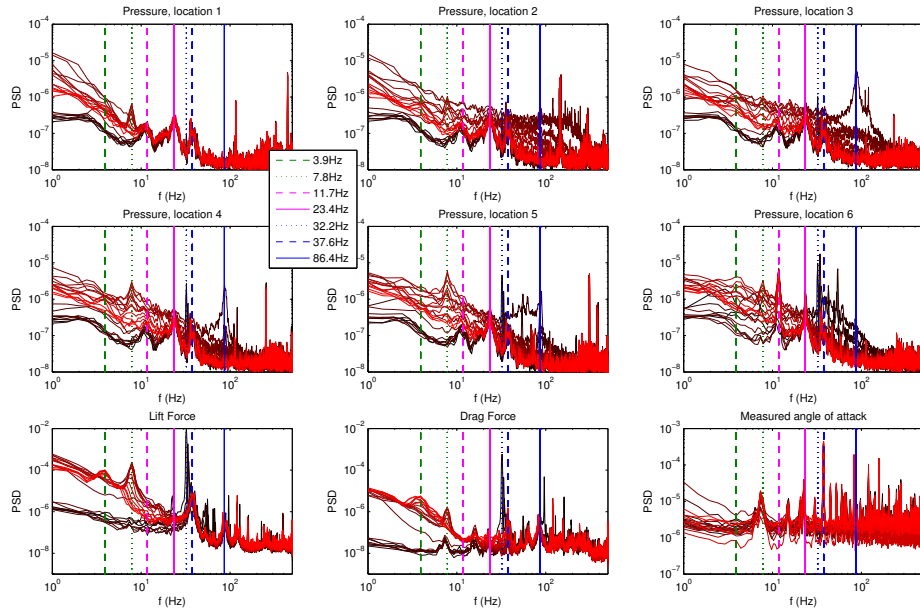


Figure 3.13: Power spectral densities of the pressures, forces, and angle of attack measurements for angles of attack ranging between 0° (black) and 21° (red), with pertinent frequencies identified by vertical lines.

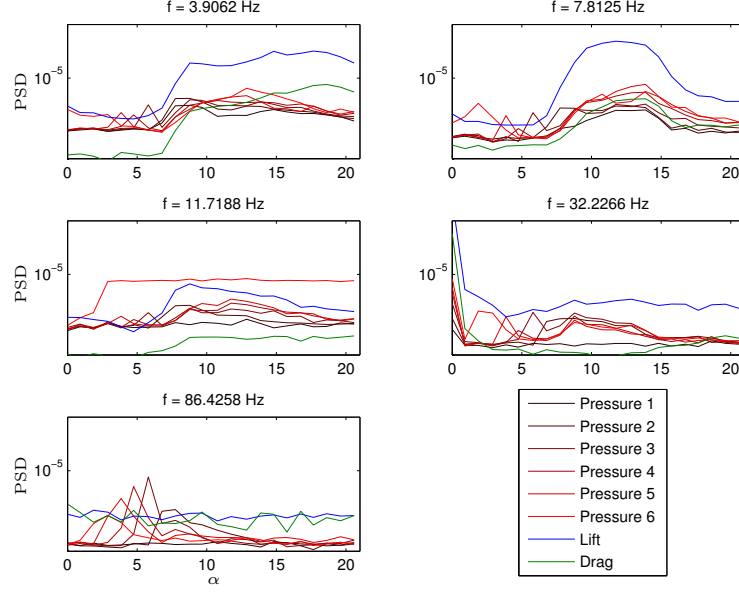


Figure 3.14: Power spectral density of pressure and force signals at selected frequencies as a function of α .

a drop in $-C_p$ at $\alpha = 5^\circ$ at a location $x/c \approx 0.65$, which is broadly consistent with our static pressure measurements. At $\alpha = 5^\circ$, Yarusevych [178] finds a spectral peak in the stream wise velocity at $x/c \approx 0.53$ at $f = 55\text{Hz}$, which corresponds to $\frac{f}{U} = 5.75$ for $Re = 55 \times 10^3$, which in turn corresponds to a frequency of approximately 70 Hz for our experiment. Note that this is of similar frequency to the peaks observed in pressure measurements for moderate angles of attack (indicated with blue vertical lines), suggesting that the same instability mechanism is responsible for the peaks recorded here.

In Figure 3.14 we show how the PSD magnitude varies with angle of attack for selected frequencies. We note that at 7.8 Hz, all measurements show largest amplitude between 10° and 15° , possibly indicating that coherent periodic shedding is most notable at this frequency in this range of angles of attack. While the vortex shedding frequency should decrease as the angle of attack increases beyond this range, we lose discernible peaks in the spectra. At the frequency of 86.4Hz, we find localized maxima at different angles of attack for each pressure sensor, with the location of the maxima moving upstream along the chord as the angle of attack increases. We suggest that this frequency is associated with Kelvin-Helmholtz vortex rollup, which occurs near the point at which the flow separates. Further evidence of this is seen in Figure 3.11, where it is shown that the angle at which the resonance peak at 86.4 Hz is largest occurs at an angle just beyond the pressure maximum, beyond which it was reasoned that the flow becomes separated.

3.6.2 Kalman filter design

The linear models (taking the form given in Equation 3.2.1) identified in this work give a prediction of forces and pressures given knowledge of the airfoil pitching kinematics. In some situations, it might be possible to supplement knowledge of these kinematics (i.e., the system inputs) with some number of additional measurements. This section briefly describes the setup and design of a Kalman filter [85], that is used to improve the estimate of the state of the system using such additional measurements. Suppose we have a state space system of the form

$$\begin{aligned}\mathbf{x}_{i+1} &= \mathbf{A}\mathbf{x}_i + \mathbf{B}\mathbf{u}_i + \mathbf{G}\mathbf{w}_i \\ \mathbf{y}_i &= \mathbf{C}\mathbf{x}_i + \mathbf{D}\mathbf{u}_i + \mathbf{H}\mathbf{w}_i + \mathbf{v}_i,\end{aligned}$$

which is a generalization of Equation 3.2.1 to include the influence of plant disturbances \mathbf{w} , as well as an output equation that includes sensor noise \mathbf{v} . For the purposes of this work, we will assume that $\mathbf{H} = 0$, $\mathbf{D} = 0$, and \mathbf{C} is a matrix that selects a subset of the states \mathbf{x} as outputs. The Kalman filter gives an estimate of the state $\hat{\mathbf{x}}$ from the system inputs \mathbf{u} and (potentially noisy) measurements \mathbf{y} , whose dynamics are governed by the equation

$$\hat{\mathbf{x}}_{i+1} = \mathbf{A}\hat{\mathbf{x}}_i + \mathbf{B}\mathbf{u}_i + \mathbf{L}(\mathbf{y}_i - \mathbf{C}\hat{\mathbf{x}}_i - \mathbf{D}\mathbf{u}_i).$$

Here the matrix \mathbf{L} gives the optimal state estimate for given disturbance and noise covariance matrices, $\mathbf{Q} = E(\mathbf{w}\mathbf{w}^T)$ and $\mathbf{R} = E(\mathbf{v}\mathbf{v}^T)$, respectively. Further details concerning Kalman filter design and the computation of \mathbf{L} may be found in standard optimal control textbooks [e.g., 146]. For this work, we take \mathbf{Q} and \mathbf{R} to be appropriately sized diagonal matrices, and set all diagonal entries to be equal aside from the entries of \mathbf{Q} corresponding to the lift and drag states, which we decrease by a factor of 10 to avoid excessive oscillations in the estimated force coefficients. We find that Kalman filter performance is relatively insensitive to changes in these weights.

Chapter 4

Lift enhancement at high angles of attack with periodic pitching

In this chapter, we study flow over a sinusoidally pitching, two-dimensional flat plate at a Reynolds number of 100, across a range of pitching amplitudes, frequencies, mean angles of attack, and pitch axis locations. We report on the lift, drag, and wake structures present in different regions of parameter space. We examine the average and spectral properties of the forces on the airfoil, and use dynamic mode decomposition to examine the structures and frequency content of the wake. Focus is given to a number of regions in parameter space where interesting behavior is observed. In particular, we find that in the regime where the flow on the upper surface of the airfoil is separated, but the steady wake is stable, pitching at a specific frequency excites a vortex shedding mode in the wake, leading to substantial increase in the lift and drag forces. This phenomena is insensitive to pitch-axis location and amplitude. At higher angles of attack where the wake for a steady airfoil exhibits periodic vortex shedding, we find that, in addition to this mean lift maxima, the interaction between the natural and forced modes gives rise to more complex behavior.

4.1 Introduction

The unsteady motion of airfoils at low Reynolds number and high angle of attack leads to a range of phenomena that cannot adequately be explained by classical aerodynamic theories. It is precisely these conditions that are encountered by small fliers, ranging from biological examples such as birds [165], insects [19, 130, 170], and bats, or manmade UAVs and MAVs. The vortex dynamics excited by airfoil motion, actuation, or indeed present in the natural flow at sufficiently high angles of attack, can significantly affect aerodynamic performance [97]. This motivates work that seeks to understand and control such phenomena, and indeed modeling the dynamics of pitching and plunging airfoils has attracted significant recent attention [e.g., 23, 24, 11, 63, 66, 41, 180]. That is not to say that such unsteady aerodynamic studies are only contemporary; the need to study unsteady aerodynamic effects was originally motivated through the study of helicopter aerodynamics [91, 167, 154].

This chapter will investigate the interaction between periodic vortex shedding that can occur for bluff bodies (or airfoils at sufficiently large angle of attack), and imposed sinusoidal

pitching motion. In the case of plunging motion, it has been observed that lock on [131] can occur between plunging frequency and natural vortex shedding [179, 34, 33], while similar phenomena have been found for surging oscillations over a wide range of Reynolds number [30]. We will show that similar phenomena are observed in the case of pitching motion. In experimental conditions, plunging oscillations may also lead to a bifurcation of the wake direction [35].

It has been observed that optimal pitching trajectories maximize the circulation that is entrained in leading edge vortices [100], linking the concepts explored here to the general notion of a formation number that can be used to explain characteristic sizes and frequencies associated with vortex phenomena [57]. While this work will only consider two-dimensional airfoils, we note that three dimensional phenomena can give rise to additional complexity. For example, aspect ratio effects being significant on lift enhancement due to periodic forcing [150], and in the case of three dimensional airfoils that are free to exhibit rolling motion, forced pitching can lead to self-excited roll oscillations [156].

There are strong parallels between studying the effect on lift and drag in the context of lifting airfoils, and in the investigation of thrust-generating airfoils [45], where it is found that propulsive efficiency is maximized when flapping produces a reverse von Kàrmàn wake that excites the least stable spatial mode of the mean wake flow [158, 157]. We describe the numerical method and scope of the work in Section 4.2, before presenting results in Section 4.3. We will focus on presenting and analyzing results at parameters where the pitching motion triggers, strengthens, or interferes with vortex shedding.

4.2 Numerical method and scope of investigation

An immersed boundary projection method is used to perform direct numerical simulations of the incompressible Navier–Stokes equations, using the approaches of Taira and Colonius [149] and Colonius and Taira [36]. We summarize the main features of this approach here, and refer readers to the relevant papers for further details. In this method, the influence of a solid body immersed in a surrounding fluid is represented with a series of boundary forces, whose locations do not have to conform to the underlying grid. The origins of immersed boundary methods date back to Peskin [113] in the study of heart hemodynamics. The method presented in Taira and Colonius [149] formulates the solution for the flow field and boundary forces as a Lagrange multiplier problem, which gives it the same form as a projection (or fractional step) method [31, 153], which may be expressed as an LU decomposition of the discretized equations [111]. The fractional step method solves the governing equations by first solving the momentum equation without the pressure term, next solving a Poisson equation for the pressure, and lastly projecting the pressure-free intermediate in a manner that accounts for the now-known pressure. Rather than performing this procedure directly, it is possible to use a nullspace method to eliminate the need to solve for the pressure separately [36]. The use of a discrete sine transform allows for computational speedup, but at the cost of requiring a uniform computational grid. This restriction may be relaxed through the use of a series of nested, uniform grids of different size and resolution [36], which allows for sufficient spatial resolution near the body, and a large total domain to eliminate boundary effects.




Figure 4.1: Computational domain used for this study, with the size and location of the airfoil shown. Dashed lines represent the borders of each nested grid.

The domain for this work consists of four nested grids about a flat plate airfoil. A diagram of the computational domain is shown in Figure 4.1. Each of the four grids contains 600 by 300 grid points, with a total computational domain extending 96 and 48 chord lengths in the streamwise and transverse directions, respectively. The Reynolds number (based on chord length and freestream velocity) is fixed at 100 throughout.

Resolution studies were performed to ensure that the resolution and extent of the domain were sufficient. To give an example of this, we show in Figure 4.2 results from performing simulations using between one and five nested grids. To show that the extent of the innermost domain (i.e., that with the finest resolution) is sufficient, we also show the result obtained from using three larger nested grids, each centered on the airfoil, of size 1500 by 500 grid points (with the same spatial resolution as before). It is observed that increasing the size of the total domain beyond 4 grids has minimal effect on the forces on the airfoil, with the mean lift changing by less than 0.4% when a fifth outer grid is included.

Crank-Nicholson and third-order Runge-Kutta time steppers are used to evolve the linear and nonlinear terms respectively, with step size ranging between $\Delta t = 0.0005c/U$ and $\Delta t = 0.01c/U$, with the smaller range of Δt required to resolve pitching motions with larger frequencies and/or amplitudes. For the domain used, advancing the solver one timestep typically takes less than one second on an Intel Xeon E5 CPU. The relatively small computational cost of an individual simulation for such two-dimensional, low Reynolds number flows makes the thorough investigation of a high-dimensional parameter space feasible.

In all cases, we run the simulations for sufficiently long such that any limit cycle or long-time behavior is reached before the data that is used for analysis is collected. We consider airfoil kinematics of the form

$$\alpha(t) = \alpha_M + \alpha_A \sin(2\pi f^* t), \quad (4.2.1)$$

where $f^* = fc/U$ is a dimensionless frequency. We perform simulations with the mean angle of attack varying in 5° increments between 15° and 45° , with pitching amplitudes of 1° , 2° , 5° and 10° . We consider frequencies in the range $f^* \in [0.01, 2]$, with approximately 20 frequencies used for each α_M and α_A . For some cases, additional frequencies are added to improve local resolution in parameter space. Performing these simulations with for pitching

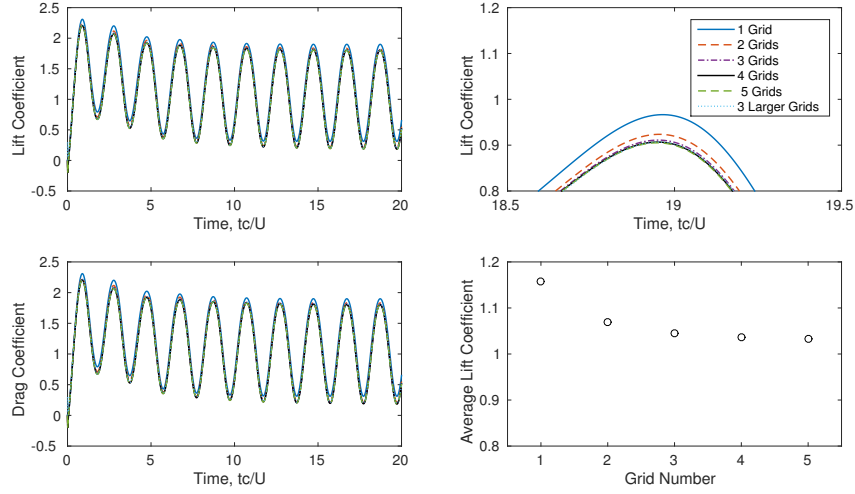


Figure 4.2: Resolution study to determine appropriate size of computational domain. The simulations start from zero initial conditions, with the airfoil performing sinusoidal pitching as defined by Equation 4.2.1 with $\alpha_M = 20$, $\alpha_A = 5$, $f^* = 0.5$. The left subplots show the time-evolution of the lift and drag coefficients, with the top right subplot showing a zoomed-in

about the leading edge, midchord, and trailing edge, this results in a total of approximately 1680 individual simulations.

4.3 Results

4.3.1 Static data

To give a sense for the behavior of the stationary airfoil, we show a typical lift curve in Figure 4.3, showing the lift when the flow is steady (i.e., the equilibrium wake, which can be stable or unstable depending on the angle of attack), and the maximum, minimum, and mean lift when the system is unsteady (corresponding to an unstable wake). We observe the expected linear relationship between the angle of attack, α , and the lift coefficient, C_L for low angles of attack. Once the angle of attack becomes sufficiently large ($\alpha > 10^\circ$), flow separation on the upper surface of the airfoil leads to a shallower lift slope. At $\alpha \approx 20^\circ$, the flow is separated and steady. Beyond a critical angle of attack $\alpha_c \approx 27^\circ$, the steady solution becomes unstable, and periodic vortex shedding is observed. This is an example of a supercritical Hopf bifurcation that is seen in the wake of bluff bodies as the Reynolds number is increased [145]. Note that in the case of an airfoil at an angle of attack, the projected area $c \sin(\alpha)$ is the effective length parameter for determining the location of the bifurcation. For $\alpha > \alpha_c$, the system exhibits higher lift (and drag) than would occur at the unstable equilibrium solution (which is computed using selective frequency damping [3]).




Figure 4.3: Lift curve for stationary airfoil, showing regions where the equilibrium is stable ($\alpha < 27^\circ$) and unstable ($\alpha > 27^\circ$), above which periodic vortex shedding occurs.

4.3.2 Force analysis

In this section, we study the lift and drag forces for pitching motion with various amplitudes, frequencies, and mean angles of attack. To begin with, only pitching about the midchord of the airfoil is considered.

Figure 4.4 shows the mean lift coefficient as a function of the dimensionless frequency, f^* , for pitching with a range of amplitudes, α_A , and mean angles of attack, α_M . We observe a distinct local peak in C_L for all values of α_M and α_A , aside from pitching with low amplitudes ($\alpha_A = 1^\circ$ or 2°) about $\alpha_M = 15^\circ$. The location of this lift peak moves slightly as α_M varies, from approximately $f^* = 0.3$ at $\alpha_M = 20^\circ$, to $f^* = 0.23$ at $\alpha_M = 45^\circ$. For $\alpha_M \geq 30^\circ$, we observe a second peak emerging at approximately twice the frequency of the dominant peak. This suggests that we excite dynamics that give enhanced lift when pitching at both a fundamental frequency and its first harmonic. We note also that the size of the lift increment seems to be largest for the intermediate base angles of 25° and 30° .

Figure 4.5 shows the mean drag coefficient for the same range of parameters as Figure 4.4. We find that there are increases in drag that show similar behavior to those for the lift. To compare the changes in lift and drag more explicitly, we plot the ratio between mean lift and drag coefficients in Figure 4.6. For larger pitching amplitudes, there is an increase in the lift-to-drag ratio at certain frequencies. The frequency of the lift-to-drag peak shows similar behavior to the peaks in both lift and drag (note that, unlike the lift and drag forces, lift over drag decreases with increasing α_M). This finding is potentially important for the effectiveness of such motions in flight (though other factors, such as the phase difference between lift and drag, would also be of importance). We additionally note that the frequency range at which these phenomena occur is well within that which is achievable by onboard actuators for centimeter-scale flapping-wing aerial vehicles [96, 152].

Turning our attention back to the lift coefficient, we plot in Figure 4.7 the increase in lift coefficient from the fixed-wing value at α_M , normalized by the amplitude of pitching.

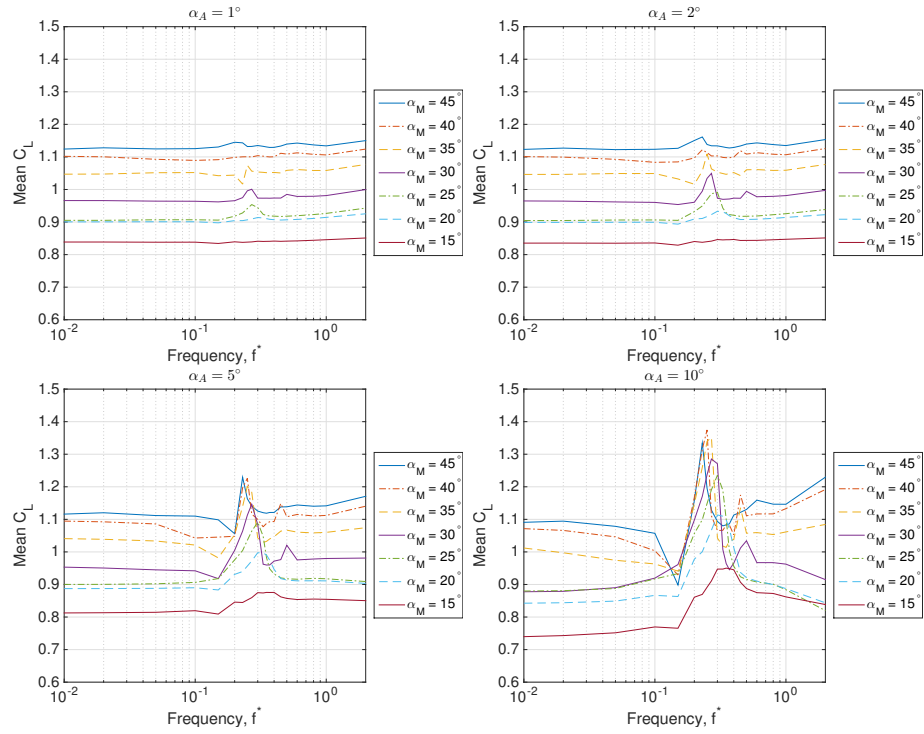


Figure 4.4: Mean lift coefficient for a range of pitching amplitudes α_A , frequencies, f^* , and mean angles of attack, α_M . Pitching is about the midchord.

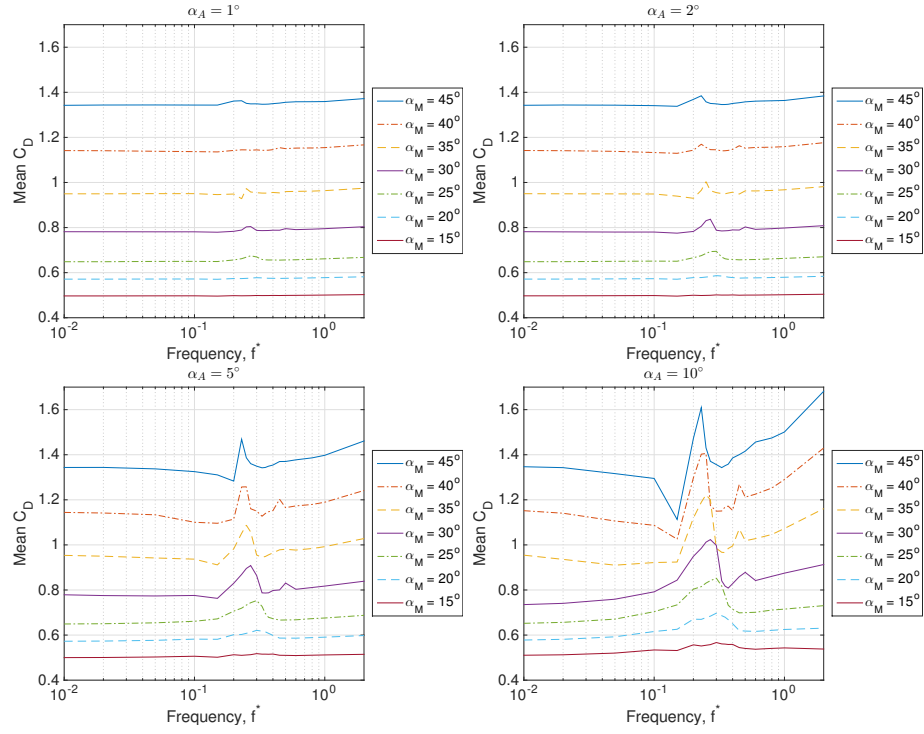


Figure 4.5: Mean drag coefficient for a range of pitching amplitudes α_A , frequencies, f^* , and mean angles of attack, α_M . Pitching is about the midchord.




Figure 4.6: Mean of lift-to-drag ratio for a range of pitching amplitudes α_A , frequencies, f^* , and mean angles of attack, α_M . Pitching is about the midchord.

For angles of attack below the critical angle (α_c) at which vortex shedding occurs, the lift increment is slightly larger for larger α_A , even after normalizing, despite the fact that the mean lift is lower than the fixed-wing value across other frequencies. This could be due to the fact that, for larger α_A , the maximum angle of attack attained in a pitching cycle is closer to or exceeds α_c , and thus better able to excite vortex shedding. Conversely, for $\alpha_M \geq 25^\circ$, the normalized lift increment is slightly larger for smaller pitching amplitudes, though in some cases this is in line with the larger average lift increments present for lower pitching amplitudes across all frequencies.

To analyze the effect of pitch axis location, we show in Figure 4.8 the normalized lift increment (i.e., the same quantity plotted in Figure 4.7) for pitching about the leading edge, rather than the midchord. We find very similar results to pitching about the midchord, with maximum increased lift at dimensionless frequencies between 0.25 and 0.3. The lift increment is larger for pitching about the leading edge, which could be due to the increased range of motion of the trailing edge for the same angular pitch amplitude. Note also that leading edge pitching will result in larger added mass forces, which for nonzero α_M will increase lift and decrease drag, particularly for high pitching frequencies. The remainder of this section will consider pitching about the midchord. The effect of pitch axis location will be investigated more thoroughly in the next section.

Figure 4.9 indicates how the mean lift compares to the maximum and minimum lift for $\alpha_A = 5^\circ$. Note that in some of the cases the lift is not periodic with the period of forcing, so these maximum and minimum values are the global extrema over many cycles. We observe (particularly clearly for lower mean angles of attack) that as the frequency increases from low values, the amplitude of the lift response increases, while the mean remains approximately constant. Above $f^* \approx 0.2$, the amplitude of the variation in lift decreases, but with the lift minimum rising more abruptly than the lift maximum falls. This asymmetry produces the higher average lift that is observed in the range $0.1 < f^* < 0.5$. For higher frequencies, larger added mass forces mean that the amplitude of the lift oscillations continue to increase, though the mean lift stays approximately constant.

To analyze the time-varying behavior in more detail, we take the discrete Fourier transform of the lift coefficient signal in time for each trial. The results for this are shown in Figure 4.10, for pitching amplitude $\alpha_A = 1^\circ$. For $\alpha_M \leq 25$, we observe one dominant frequency peak, corresponding to the pitching frequency f^* . Even though the undisturbed wake is stable at $\alpha = 25^\circ$, we still observe some frequency content near the almost-unstable vortex shedding mode across all pitching frequencies. For larger angles of attack, there is a second major peak in the spectra, corresponding to the vortex shedding frequency at the given α_M . In the region where these two frequencies are similar, there appear to be complex interactions between the dynamics associated with each frequency. As α_M grows larger, a distinct peak emerges at the first harmonic of the fundamental vortex shedding frequency, further complicating the frequency response of the system, which can now include, at very least, sums and differences of multiples of these frequencies.

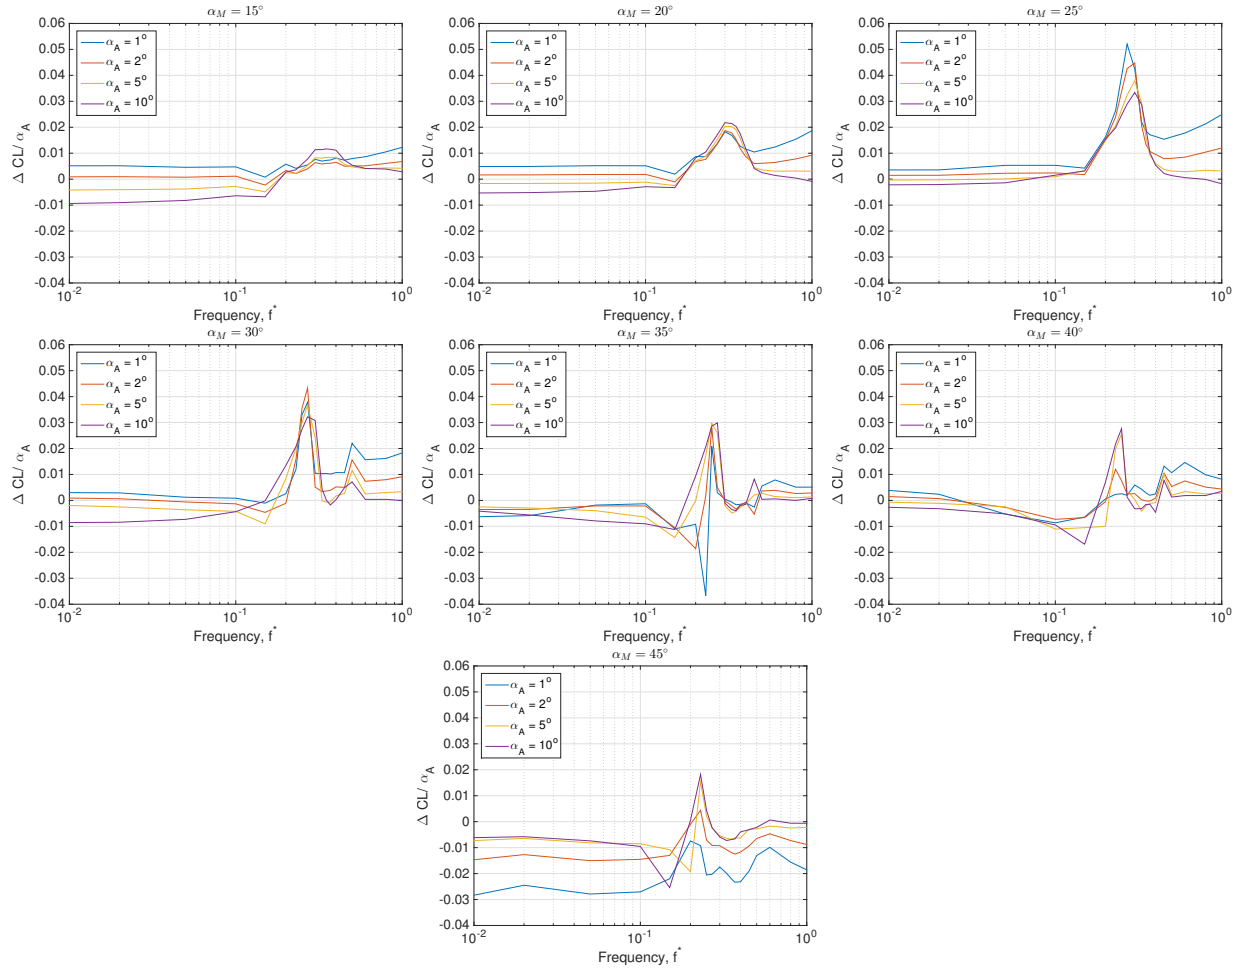


Figure 4.7: Mean lift coefficient increment over the fixed airfoil value, normalized by the amplitude of pitching α_A , across a range of pitching amplitudes α_A , frequencies, f^* , and mean angles of attack, α_M . Pitching is about the midchord.

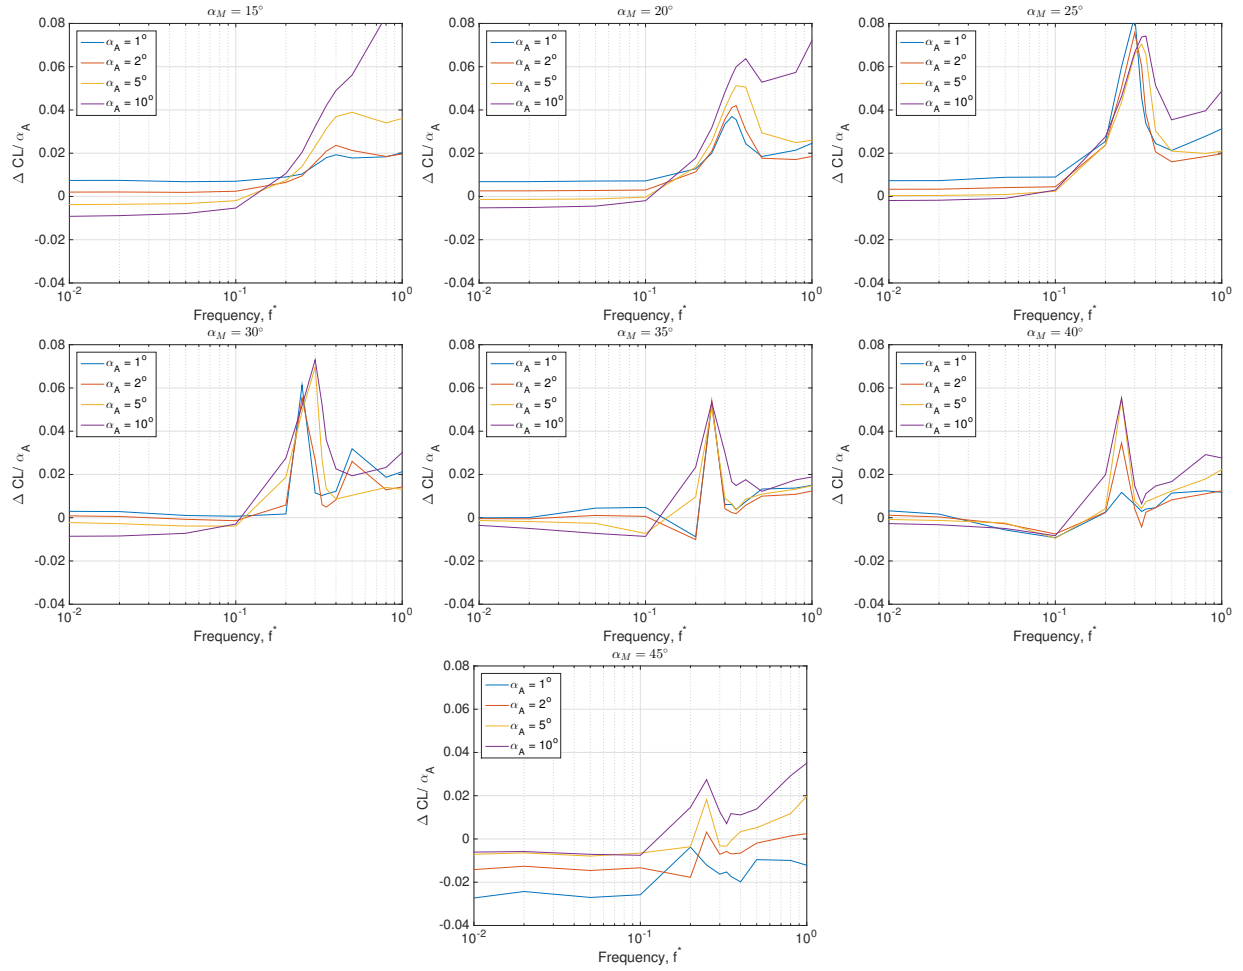


Figure 4.8: Mean lift coefficient increment over the fixed airfoil value, normalized by the amplitude of pitching α_A , across a range of pitching amplitudes α_A , frequencies, f^* , and mean angles of attack, α_M . Pitching is about the leading edge.

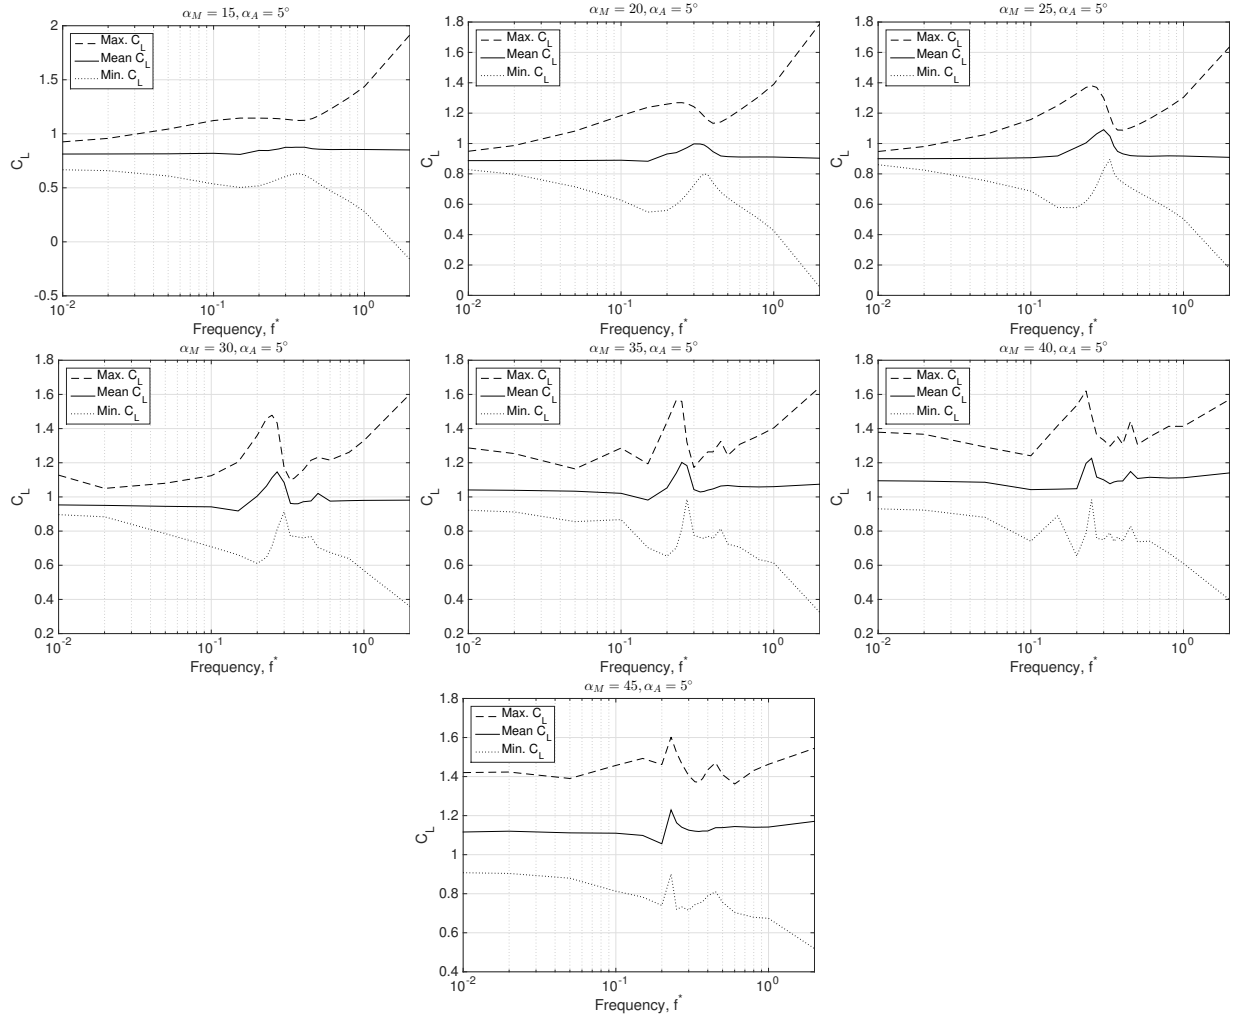


Figure 4.9: A comparison between the maximum, minimum, and mean lift coefficient attained for pitching with amplitude $\alpha_A = 5^\circ$, for a range of mean angles of attack, α_M , and frequencies, f^* . Pitching is about the midchord.

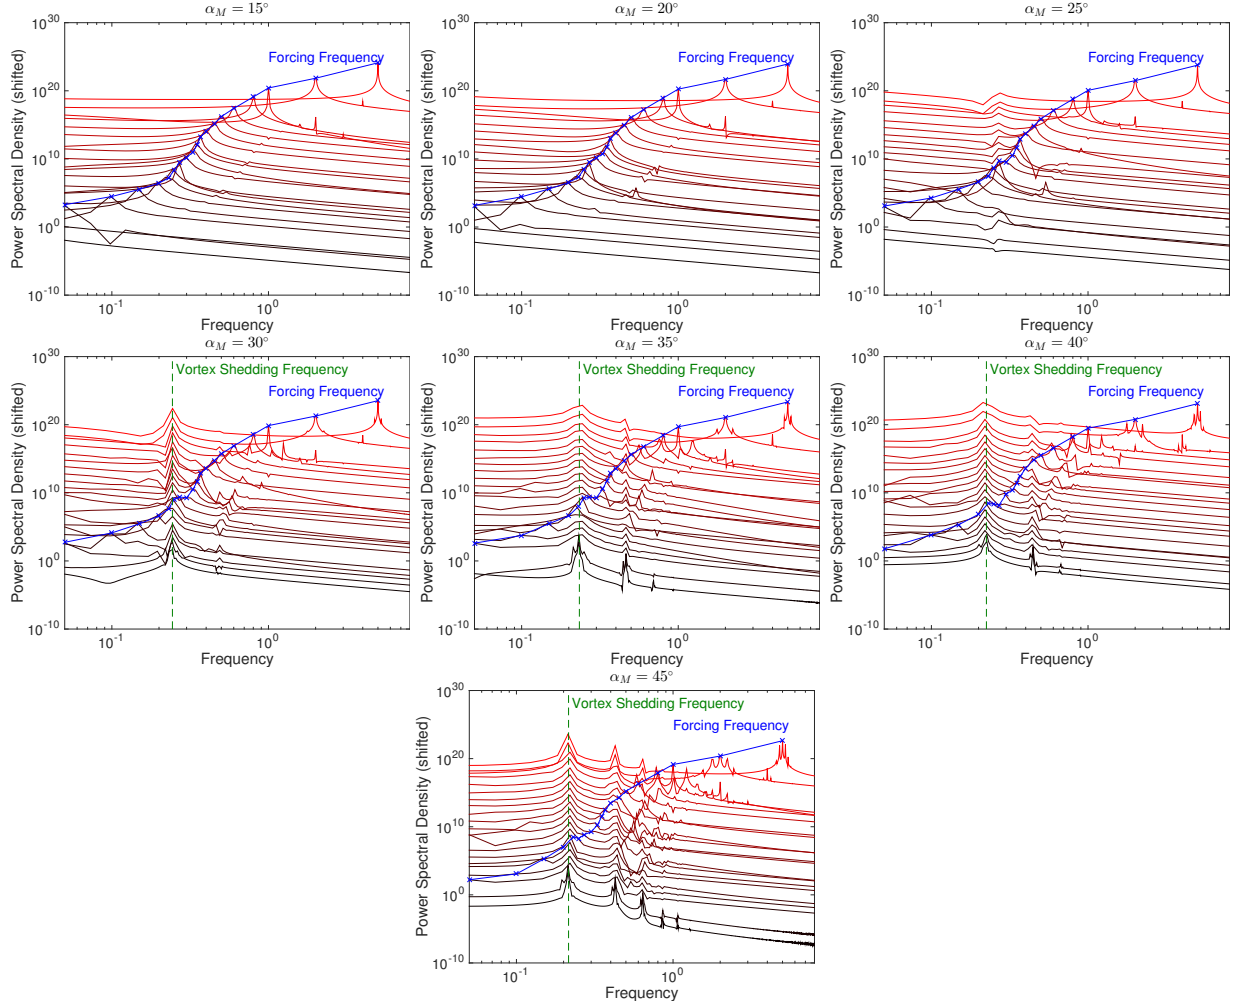


Figure 4.10: Power spectral densities of the lift force for pitching motion about the midchord with amplitude $\alpha_A = 1^\circ$. For clarity, the spectra for each forcing frequency is shifted by one order of magnitude, so the absolute scale of the vertical axis is not significant.

4.3.3 Wake analysis with dynamic mode decomposition

Here we analyze the flow field in the wake of the body, in an attempt to study the underlying physics behind the phenomena observed in Section 4.3.2. We will make use of DMD, which, as discussed previously (Sections 1.2 and 2.1.3), is a technique that can extract dynamical content from data, in the form of spatial modes and their associated growth/decay rates and frequencies of oscillation. Before applying DMD, we look at what insight can be gained through visualizations of snapshots of the wake. We begin by considering a base angle $\alpha_M = 20^\circ$, which is prior to the Hopf bifurcation at which unforced vortex shedding occurs, and thus has a stable equilibrium wake. Figure 4.11 shows vorticity field snapshots for a variety of forcing frequencies with amplitude $\alpha_A = 1^\circ$. We observe that at around $f^* \approx 0.3$, the vorticity field is qualitatively different, with periodic vortex shedding being excited by the pitching motion. This immediately suggests that the increased lift observed at these pitching frequencies arises due to enhanced vortex formation. We note that the vorticity fields for pitching about the leading edge (left) and midchord (right) show similar results, with leading-edge pitching generally leading to stronger, more distinct vortices forming closer to the airfoil.

The vorticity fields shown in Figure 4.12 have the same parameters as those for Figure 4.11, but with a larger pitching amplitude of $\alpha_A = 5^\circ$. We again see the same phenomena where the forcing excites vortex shedding, but here distinct vortices form and persist downstream over a wider range of frequencies. This is consistent with the results from Section 4.3.2, where for the $\alpha_M = 20^\circ$ case in Figure 4.4 we observe an enhanced lift over a wider range of frequencies for higher forcing amplitudes.

To investigate the behavior as we move closer to, and past, the critical angle of attack α_c , we show in Figure 4.13 snapshots of the vorticity field for mean angles of attack $\alpha_M = 25^\circ$, and 30° , keeping the pitching amplitude $\alpha_A = 5^\circ$. For $\alpha_M = 25^\circ$, the observed behavior is similar to that exhibited in Figures 4.11 and 4.12, with more distinct vortex shedding apparent for intermediate frequencies. For the $\alpha_M = 25^\circ$ case (left subplots of Figure 4.13), however, vortex formation appears to be promoted at a slightly lower frequencies than with $\alpha_M = 20^\circ$, which is consistent with the locations of the lift peaks for these values of α_M in Figure 4.4. When $\alpha_M = 30^\circ$ (right subplots of Figure 4.13), the fact that $\alpha_M > \alpha_c$ means that there would be vortex shedding without any pitching motion, which is consistent with the fact that clear vortex shedding is observed some distance downstream of the airfoil for all pitching frequencies. The interaction between the pitching and this natural instability appears to be more complex than (or at least, more difficult to interpret directly from vorticity snapshots), so we will make use of DMD to allow for more detailed and quantitative analysis of this case.

DMD is performed on the vorticity field for the same values of α_M and α_A as used in Figure 4.13 ($\alpha_M = 25^\circ$ and 30° , with $\alpha_A = 5^\circ$), for a variety of pitching frequencies. We use 401 snapshots for each case, with a time gap between snapshots of $\Delta t = 0.1c/U$. The snapshots are collected after first allowing the system to evolve from rest for 50 convective time units. We take data from a region downstream of the body as our domain for DMD, to avoid complications associated with having the moving body in the domain.

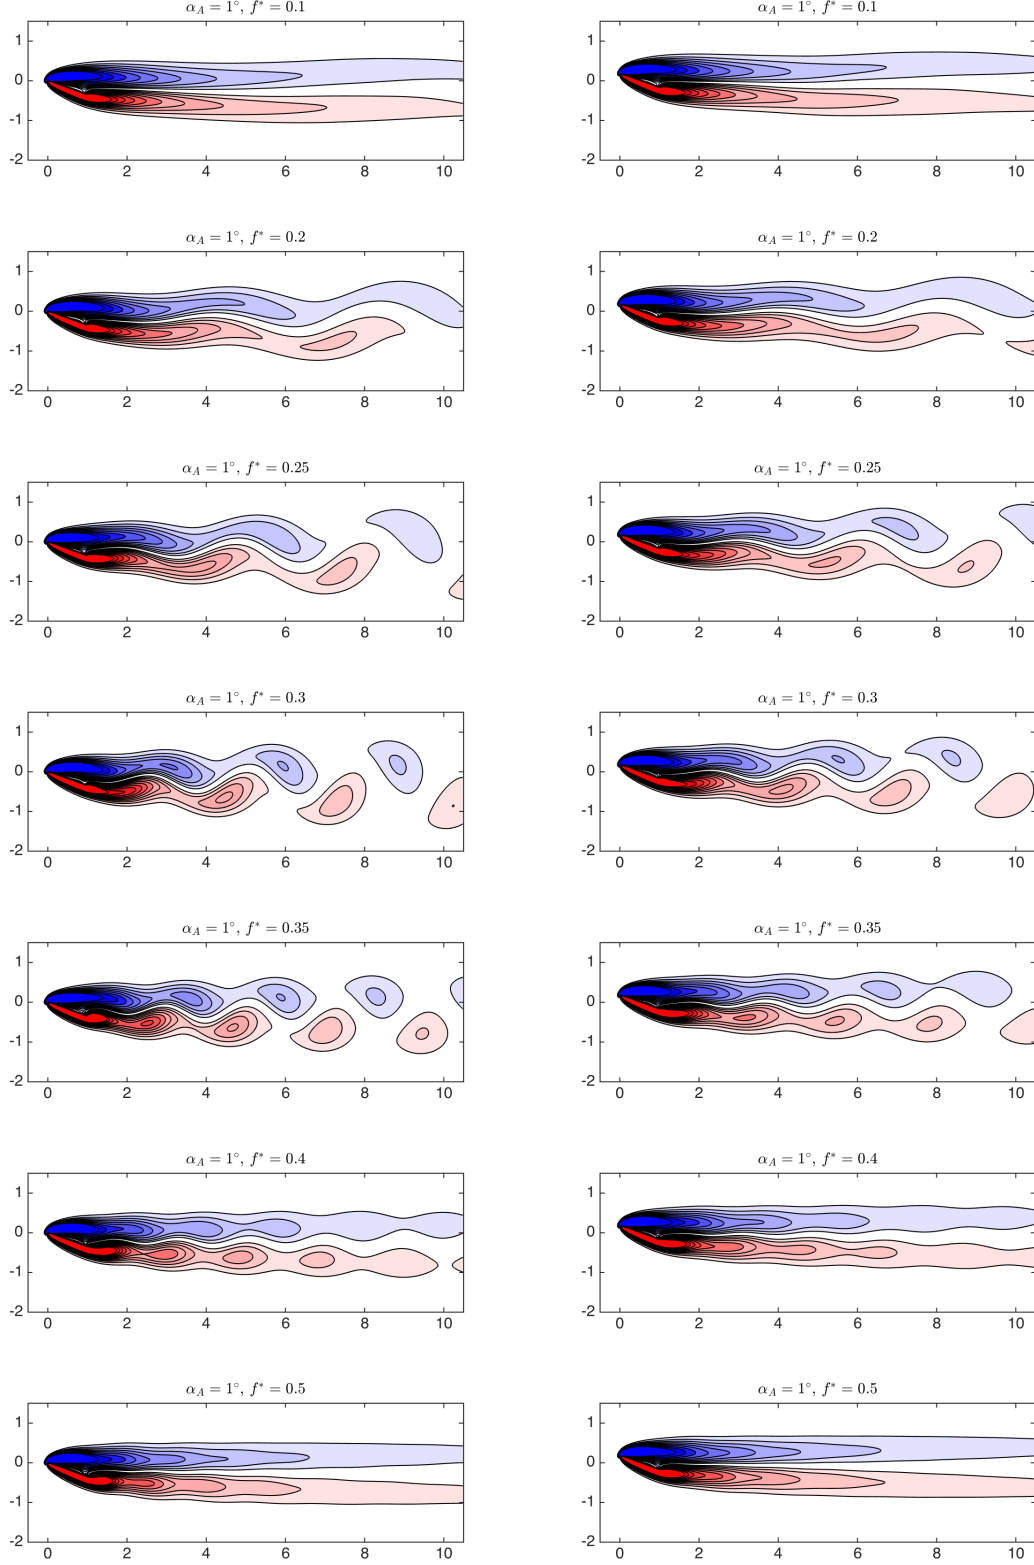


Figure 4.11: Instantaneous vorticity fields for pitching about the leading edge (left) and midchord (right) at a variety of frequencies, with $\alpha_M = 20^\circ$ and $\alpha_A = 1^\circ$.

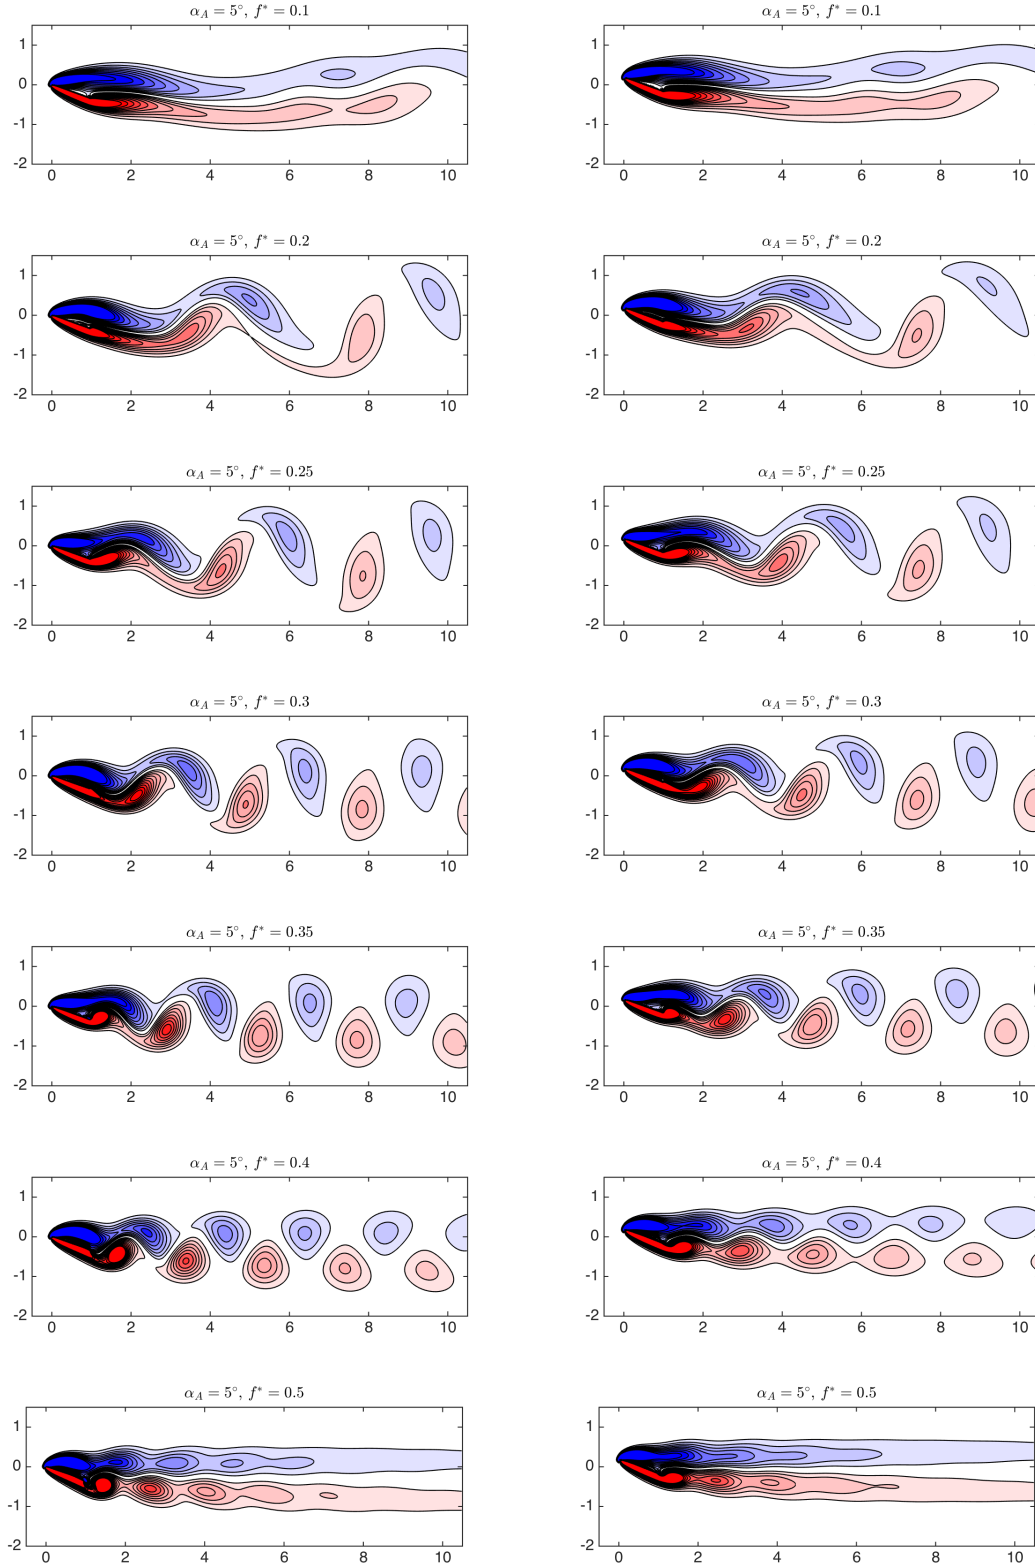


Figure 4.12: Instantaneous vorticity fields for pitching about the leading edge (left) and midchord (right) at a variety of frequencies, with $\alpha_M = 20^\circ$ and $\alpha_A = 5^\circ$.

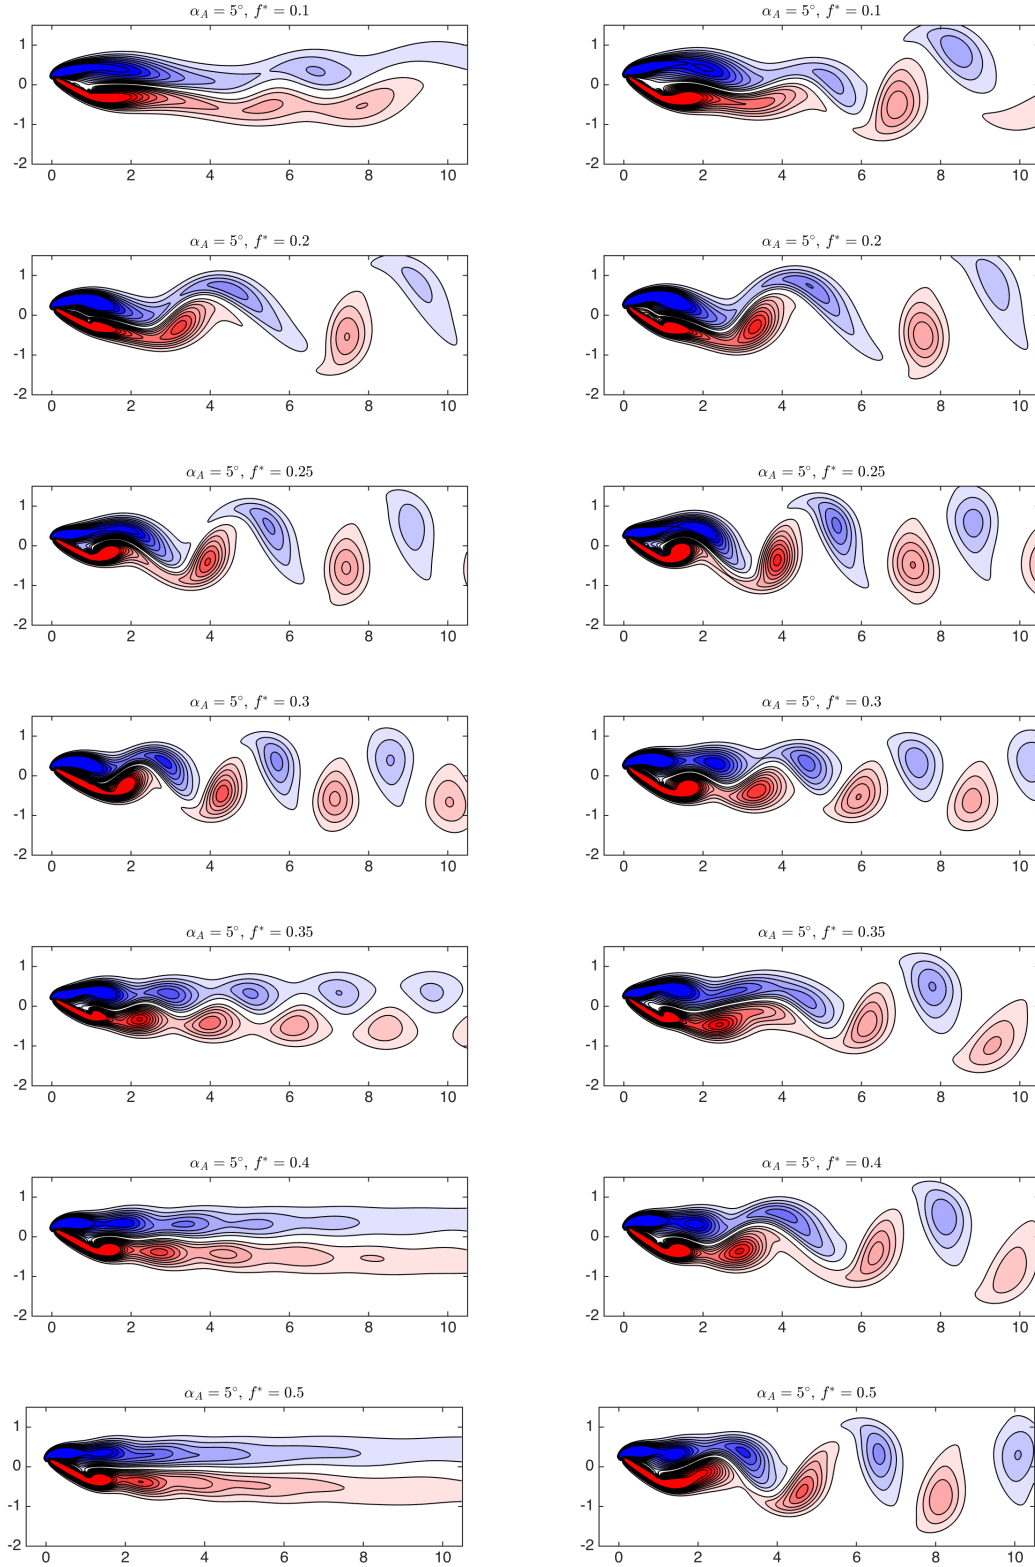


Figure 4.13: Instantaneous vorticity fields for pitching about the midchord (right) at a variety of frequencies, with $\alpha_A = 5^\circ$ and $\alpha_M = 25^\circ$ (left) and $\alpha_M = 30^\circ$ (right).

While it is most often applied to data from systems without external forcing, DMD can be used to examine forced systems [159, 116, 102], though one must take care in the interpretation of its outputs. In this context, we show that it can be an effective tool to determine both the frequency content of the wake, and also if the wake “locks on” to a vortex shedding mode. As an aside, we note that more rigorous connections between DMD and Fourier analysis can be made [28].

Figures 4.14 and 4.15 show results of applying DMD to the wake when $\alpha_M = 25^\circ$ for a variety of pitching frequencies. The amplitudes of each DMD mode are plotted against their corresponding frequency in the left subplots. Also plotted are the real components of the vorticity fields of the four modes of largest amplitude for each forcing frequency. Note that if one decomposes a time-varying field $\mathbf{u}(x, t)$ into DMD modes $\varphi_i(x)$ (with corresponding continuous-time eigenvalue λ_i) by

$$\mathbf{u}(x, t) = \sum_{i=1}^N a_i e^{i\lambda_i t} \varphi_i(x), \quad (4.3.1)$$

then we define the amplitude of mode φ_i by a_i , where it is assumed that each mode has been scaled to be of unit norm. See Tu et al. [161] for more details about other possible methods for defining DMD mode amplitudes. Note also that the frequency of a given DMD mode is $\frac{\text{imag}(\lambda_i)}{2\pi}$ (where the factor of 2π converts from radians to cycles, to be consistent with the definition of pitching frequency). The relative amplitudes of the zero-frequency mode (which will be referred to as the constant mode) and that at the mode with the frequency of forcing (henceforth the forcing mode) indicates the extent to which the forcing influences vortex shedding in the wake. Across all frequencies, most of the high-amplitude DMD modes occur at multiples of the forcing frequency. As the frequency increases from its smallest values, the amplitude of the forcing mode increases, to the extent where it is larger than the constant mode for $f^* = 0.25$ and 0.3 . At higher frequencies, the forcing mode again decrease in amplitude, indicating that forcing has a reduced influence on the wake. Aside from the higher amplitude of the forcing mode at $f^* = 0.25$ and 0.3 , the shape of the constant mode also changes with the forcing frequency. In particular, for $f^* = 0.2$ and 0.25 the constant mode gives a vorticity field that does not extend as far downstream of the body, and spreads out more in the direction normal to the flow. Referring back to Figures 4.12 and 4.13, the enhanced spreading of the mean flow is seemingly due to the larger vortices drifting above (for negative vorticity) and below (positive vorticity) the airfoil as they are shed. This phenomenon could also explain why the lift-to-drag ratio is often maximized at pitching frequencies slightly larger than those at which maximum lift occurs. The enhanced spreading at lower frequencies increases drag, so lift-to-drag benefits exist only at higher frequencies, where the wake remains thinner.

For pitching at $f^* = 0.05, 0.4$, and 0.5 , a DMD mode with frequency approximately 0.25 emerges, which is largest further downstream from other modes (about $8\text{--}10c$ downstream of the airfoil). This seems to be due to a marginally stable wake mode being excited, even

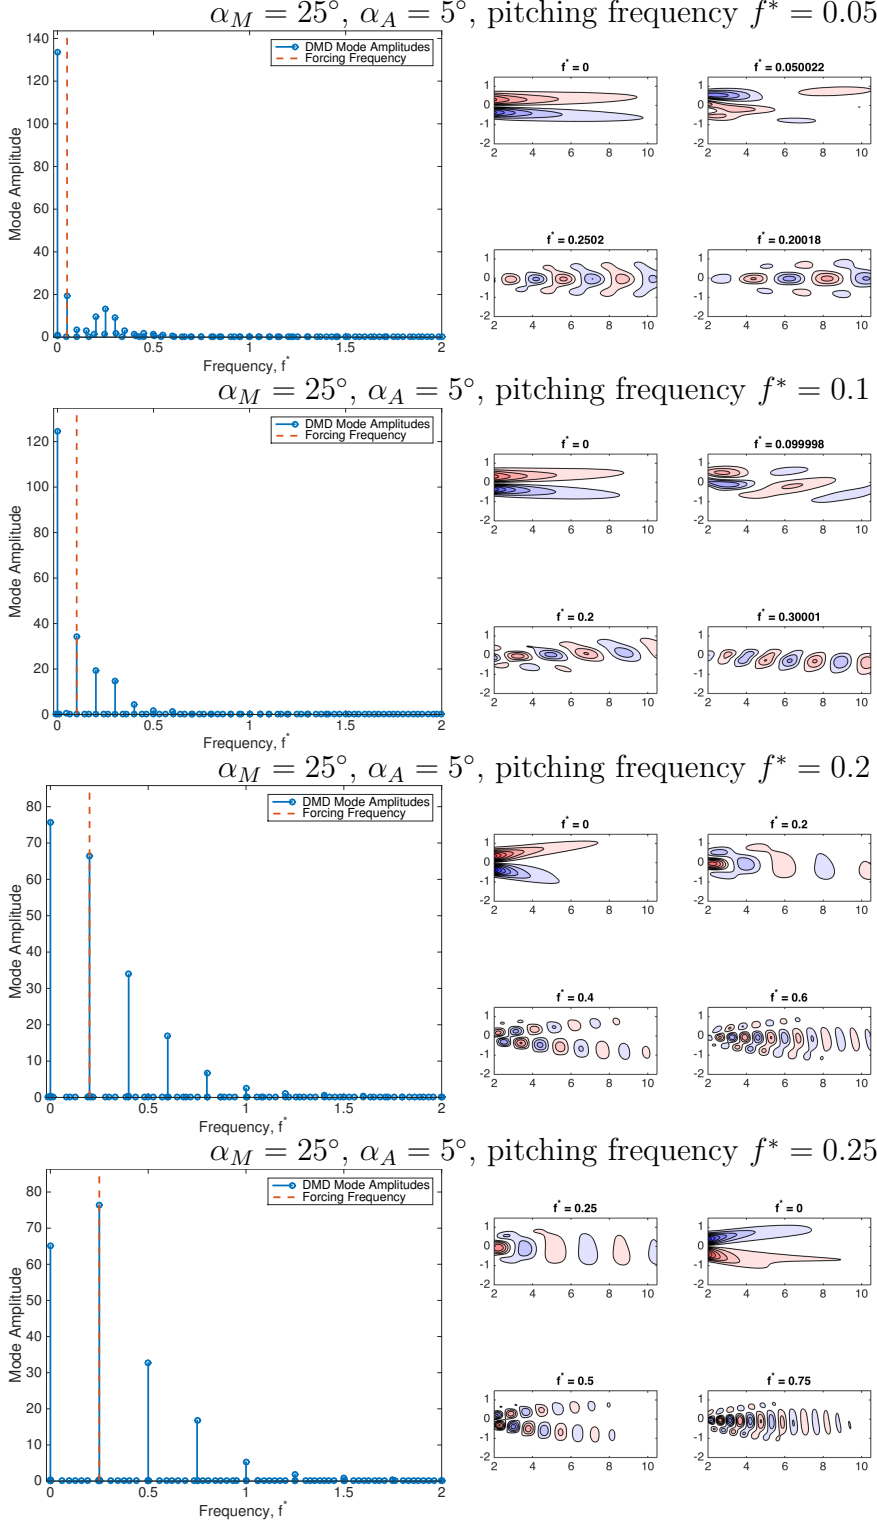


Figure 4.14: DMD eigenvalue amplitudes and frequencies (left) and eigenmodes corresponding to the four largest amplitude eigenmodes, for pitching with $\alpha_M = 25^\circ$, $\alpha_A = 5^\circ$, and $f^* = 0.05, 0.1, 0.2$ and 0.25 .

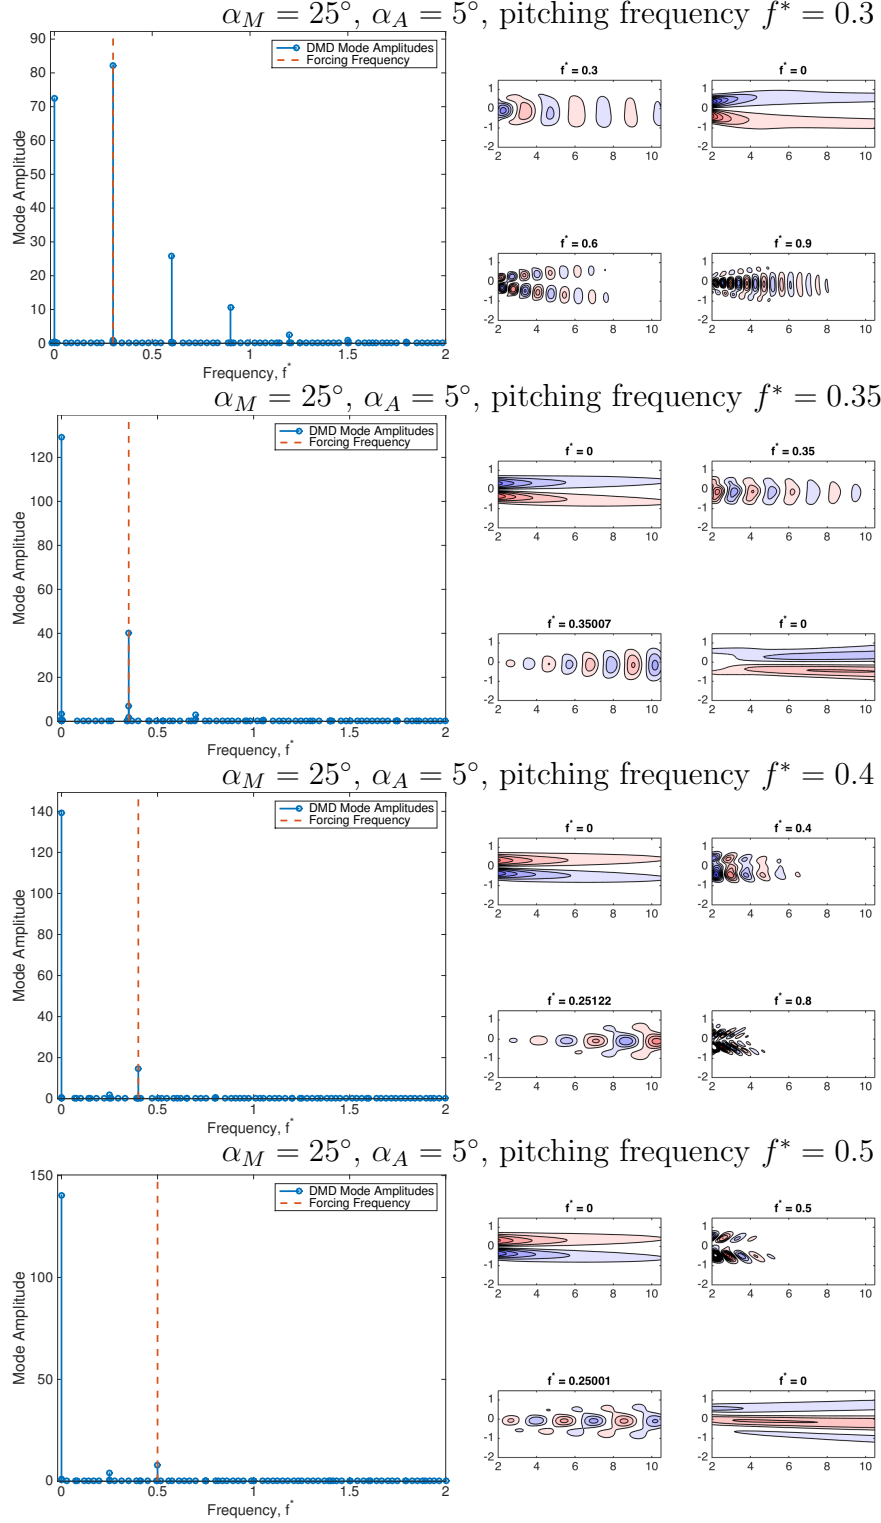


Figure 4.15: DMD eigenvalue amplitudes and frequencies (left) and eigenmodes corresponding to the four largest amplitude eigenmodes, for pitching with $\alpha_M = 25^\circ$, $\alpha_A = 5^\circ$, and $f^* = 0.3, 0.35, 0.4$ and 0.5 .

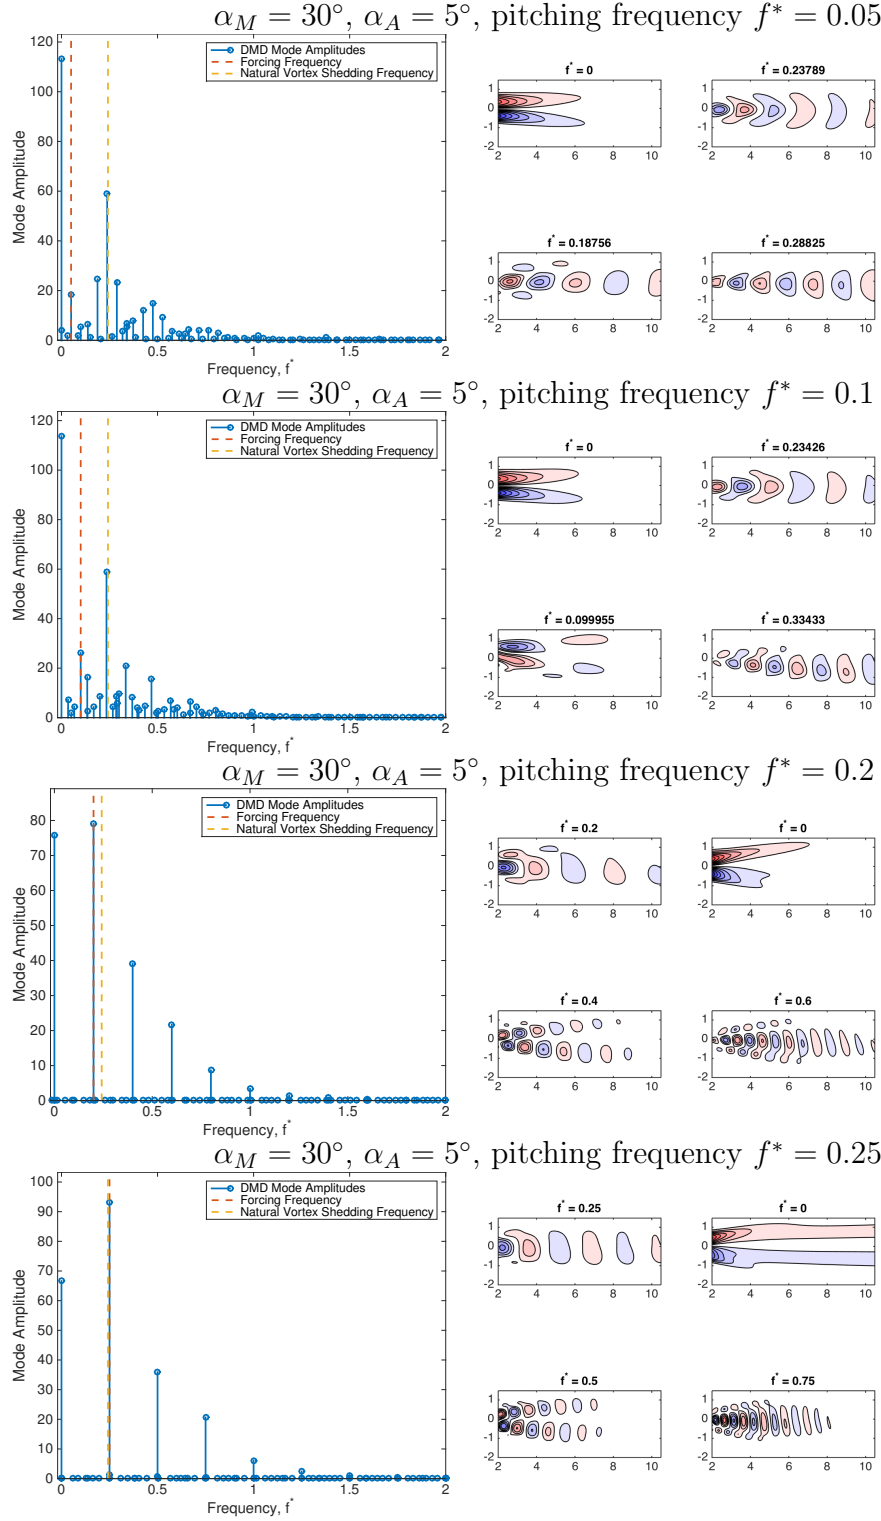


Figure 4.16: DMD eigenvalue amplitudes and frequencies (left) and eigenmodes corresponding to the four largest amplitude eigenmodes, for pitching with $\alpha_M = 30^\circ$, $\alpha_A = 5^\circ$, and $f^* = 0.05, 0.1, 0.2$ and 0.25 .

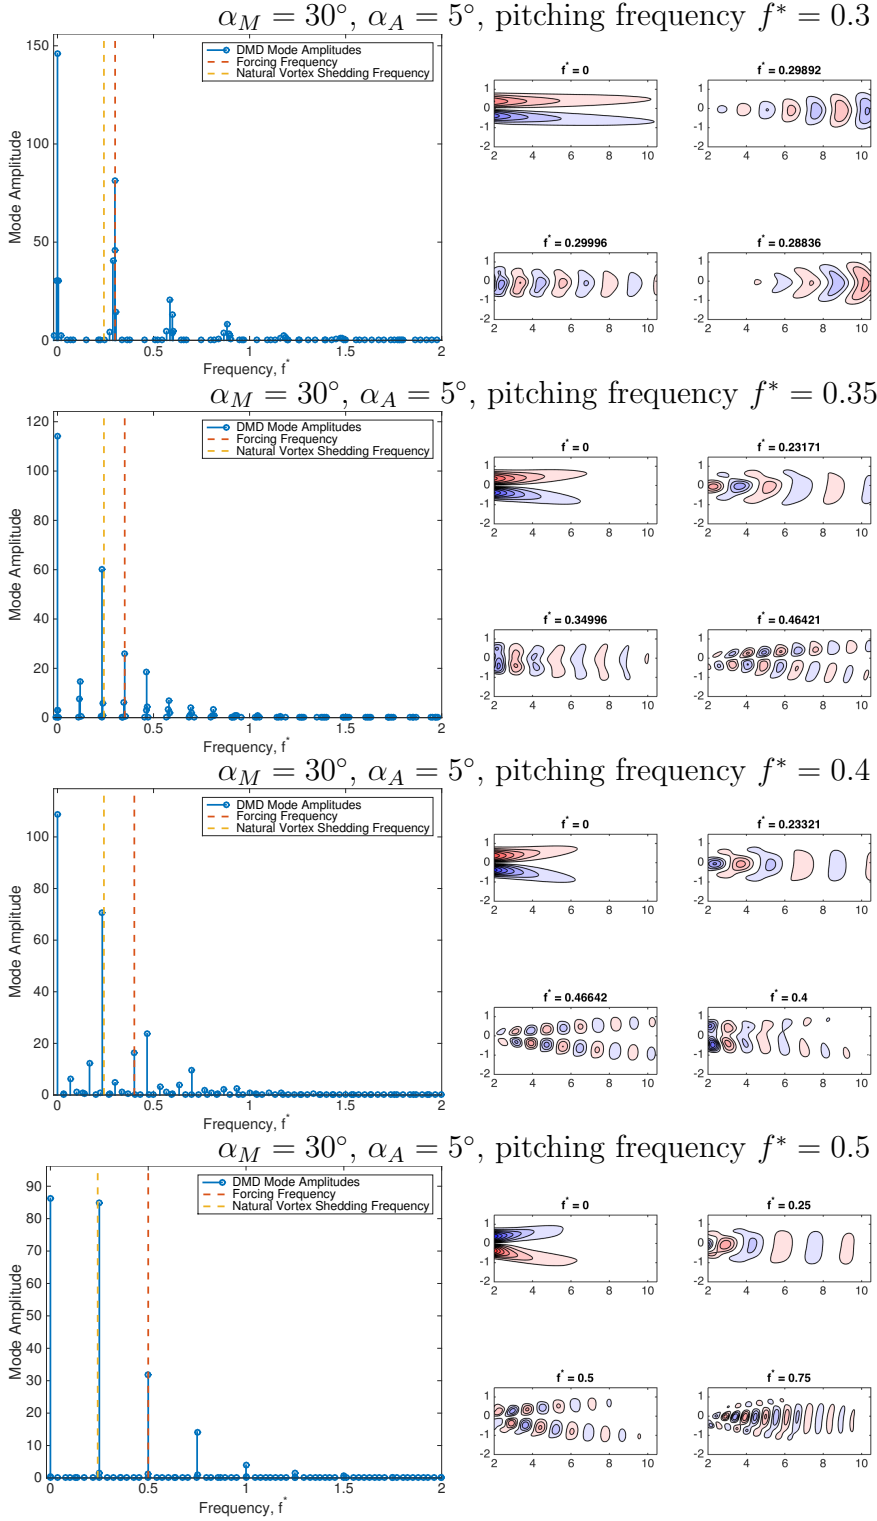


Figure 4.17: DMD eigenvalue amplitudes and frequencies (left) and eigenmodes corresponding to the four largest amplitude eigenmodes, for pitching with $\alpha_M = 30^\circ$, $\alpha_A = 5^\circ$, and $f^* = 0.3, 0.35, 0.4$ and 0.5 .

though it oscillates at a different frequency to the forcing frequency (though for $f^* = 0.5$, this represents period-doubling behavior). This is perhaps unsurprising, as for $\alpha_M = 25^\circ$ and $\alpha_A = 5^\circ$, the motion of the airfoil crosses over the critical α_c , so, in the absence of the wake locking-on to the frequency of pitching, this natural instability mode is excited. For $f^* = 0.05$, we additionally observe DMD modes at sums and differences of the forcing and natural frequencies.

Having analyzed the case where $\alpha_M < \alpha_c$, we show in Figures 4.16 and 4.17 the results of applying DMD to data collected with $\alpha_M = 30^\circ > \alpha_c$. For $0.2 \leq f^* \leq 0.3$, only the forcing frequency and its harmonics are seen in the dominant DMD modes, indicating that the wake instability is locking on to the forcing frequency. As before, the largest modes with nonzero frequency arise when pitching at $f^* = 0.2$ and 0.25 , for which enhanced lift was observed (see Figures 4.4 and 4.9). A small lift peak is also observed at $f^* = 0.5$, which we observe corresponds to a relatively large DMD mode at the natural frequency, which is half of the forcing frequency.

Sufficiently far away from the natural frequency of the wake, the natural instability mode does not lock-on to the frequency of forcing, and one obtains a wake with both frequencies, as well as harmonics and sums and differences of the natural and forcing frequencies, which is broadly consistent with observations from Figure 4.10, where the forcing and natural frequencies both emerge when they are sufficiently different.

4.3.4 Experimental Results

The data presented so far in this chapter has come entirely from two-dimensional direct numerical simulations, at the low Reynolds number of 100, with a simple flat plate airfoil geometry. It is thus natural to wonder how universal the phenomena explored in this chapter are. To partially address this question, we present some preliminary experimental data. Wind tunnel force data was acquired using the experimental setup described in Chapter 3, with a NACA 0009 airfoil. Limitations on the maximum force to which the force balance could be exposed meant that only very small amplitude motions ($\alpha_A \approx 0.5^\circ$) could be applied, up to a maximum pitching frequency of approximately $f^* = 0.41$. The actual amplitude of pitching motion decreased slightly as the frequency increased, due to the limitations of the pushrod dynamics. For this reason, we consider the normalized lift coefficient increment using the measured amplitude of pitching. Figure 4.18 compares the experimental observations of the normalized lift increment to the numerical data when pitching about the midchord at $\alpha_M = 20^\circ$ (previously shown in Figure 4.7).

Qualitatively, the same enhanced lift phenomena emerges in the experimental results, with a maximum lift increment observed when pitching at a frequency $f^* = 0.35$. The fact that the magnitude of the lift peak is different is unsurprising, given the differing Reynolds number, airfoil shape, and pitch axis (0.11c from the leading edge for the experimental data) between the numerical and experimental data.




Figure 4.18: Comparison between experimental data for a NACA 0009 airfoil at a Reynolds number of 50,000, with numerical results previously shown in Figure 4.7 for $\alpha_M = 20^\circ$.

4.4 Discussion and conclusions

For angles of attack below the critical angle where unforced vortex shedding first occurs (α_c), it was found that pitching at a certain frequency can excite vortex shedding in the wake, leading to higher mean lift. The magnitude and width in frequency range of the lift increment increases significantly as the forcing amplitude increases. This is perhaps a similar phenomenon to the widening of the “resonance horn” [21] observed by Choi et al. [30] for the case of a surging and plunging airfoil (though in that case, the system was above the critical bifurcation parameter).

For α_M above the critical angle for vortex shedding, there is a similar peak in the mean lift coefficient when the pitching frequency is close to the natural vortex shedding frequency or its first harmonic. When the natural and forcing frequencies are different, the interactions between the two frequencies can lead to complex frequency spectra in the forces and wakes. Note in particular that pitching some amount below the vortex shedding frequency can lead to a notable decrease in mean lift for $\alpha_A \geq 5^\circ$ at $\alpha_M \geq 35^\circ$.

While periodic pitching at the preferred frequency where vortex shedding is excited or enhanced also leads to an increase in drag, the differences in effect that pitching has on the two force components leads to an increase in the lift-to-drag ratio for frequencies slightly above the frequency for which the lift is maximized. It is interesting to note that while the frequency at which the maximum lift increment occurs decreases slightly as α_M increases (in agreement with the slight decrease in natural vortex shedding frequency), the frequency giving maximum lift-to-drag ratio increases with α_M . Indeed, for high α_M it seems that the maximum lift-to-drag ratio typically occurs between the peaks in lift and drag located at the vortex shedding frequency and its first harmonic, where there is a local lift minimum, and thus also a slightly more substantial drag minimum.

Beyond α_c , the interactions between the pitching and vortex shedding frequencies can lead to complex frequency spectra in both the forces (as seen in Figure 4.10) and wakes

(Figures 4.16-4.17). There are numerous methods by which one can analyze frequency content, and here we show how DMD can clearly distinguish between cases where all frequencies present are harmonics of the pitching frequency, and where there is a richer range of frequency content, as well as giving a sense of the relative amplitude of each frequency component of the wake.

There has been much effort in the past to understand, model, and predict unsteady aerodynamic forces, moments, pressures, and indeed many other quantities of interest for moving airfoils. Particularly in the case of separated flow, these dynamics are often highly nonlinear. One might seek to get around this by linearizing about a certain fixed point (say an angle of attack), in the hope that a linear model would at least be locally accurate. The findings presented here suggest that such an approach might be problematic, since a linear model (e.g., a transfer function) can only predict the magnitude and phase of a response to sinusoidal forcing, but not any change in the mean value. Thus, such effects must be accounted for separately, or a more complex modeling framework used.

Chapter 5

Nonlinear reduced-order models of fluids systems using extended dynamic mode decomposition

Data-driven approximations to the Koopman operator have much potential for capturing and illuminating the dynamics exhibited by the Navier-Stokes equations. In this chapter, we show that the elements of an identified finite dimensional approximation to the Koopman operator can be utilized for the construction of accurate nonlinear reduced-order models. We present a modification to the extended dynamic mode decomposition (EDMD) algorithm through the inclusion of a regularization parameter, which we find often gives more accurate models. The performance of models identified using our proposed method are compared to those found by performing a Galerkin projection of the governing equations onto proper orthogonal decomposition modes, for the canonical case of two-dimensional flow past a circular cylinder. We demonstrate that identifying nonlinear models using EDMD is particularly advantageous when the data available is noisy, or is only available within certain regions of space or time.

5.1 Introduction

A much-desired goal in fluid mechanics, and indeed many other fields, is to obtain simple models that are capable of predicting the behavior of seemingly complex systems. Low-dimensional models can not only improve our fundamental understanding of such systems, but are often required for purpose of efficient and accurate prediction, estimation and control. A brief review of the literature pertaining to a number of such approaches was given in Section 1.2, with relevant algorithms described in Section 2.1 We once more highlight the difference between approaches where the resulting reduced-order dynamics are obtained from a reduction of the full partial differential equations, as in Galerkin projection (GP), and those where the dynamics are inferred only from data, as in DMD or ERA.

In this chapter, we will propose a method for obtaining nonlinear reduced-order models that is based upon a recently-developed extension of DMD [172], referred to as extended DMD (EDMD), in which nonlinearities can be accounted for by an appropriate choice of observables. In particular, we will explore how this algorithm can be tailored to identify

nonlinear models that have the same (or similar) form as those that would be obtained through projection of the governing equations onto a low-dimensional basis obtained from data. In this sense, the nonlinear models that we identify in this work will come from EDMD, but will have similar form to those given by GP. We note that this is a different approach to using dominant DMD modes as an alternative to POD modes for a basis for projection [155]. A further recently-proposed approach uses DMD to efficiently approximate just the nonlinear component of the dynamics [4], as an alternative to the discrete empirical interpolation method [15, 27].

We lastly note that several previous studies have considered this dichotomy between using projection onto governing equations, or using time-resolved data, to identify models, most often for linear systems. For example, Illingworth et al. [76] considers models identified directly using ERA and models identified by considering each of the pertinent physical processes individually, while Hervé et al. [71] discusses the use of both GP models and ARMAX system identification methods with physically motivated terms to model (and subsequently control) flow over a backwards-facing step.

Section 5.2 describes EDMD, which is an extension of the DMD algorithm described in Section 2.1.3. We also introduce a regularized variant of this algorithm, which will be utilized for system identification purposes. Models will be identified and tested in Section 5.3 on data obtained from numerical simulations of flow past a circular cylinder, the results of which are discussed in further detail in Section 5.4.

5.2 Reduced-order models using extended dynamic mode decomposition

This section introduces the proposed modeling approach. Section 5.2.1 discusses the EDMD extension and how it may be used to obtain nonlinear models, while Section 5.2.2 gives a regularized modification to EDMD that we find to be advantageous when using EDMD for such purposes. When viewed as a method for reduced-order modeling, the main difference between the approach discussed here and that used for GP (see Section 2.1.2) is in how the temporal dynamics are identified: GP uses the governing equations, whereas DMD/EDMD uses only data to identify dynamics.

5.2.1 Extended dynamic mode decomposition and nonlinear models

A fundamental limitation to models extracted from DMD is their linearity, which can make them entirely unable to model fundamentally nonlinear features that arise in fluid flows. In the context of DMD being an approximation of the Koopman operator (see Section 2.1.3), this limitation amounts to Koopman eigenfunctions not lying within the span of the data. Continuing from the description of DMD in Section 2.1.3, rather than applying the DMD algorithm directly to the data \mathbf{y} , we may define a set of observables $\psi(\mathbf{y})$ (where $\psi: \mathbb{R}^n \rightarrow \mathbb{R}^k$ for some k) that span a space more conducive to approximating the true dynamics. Explicitly, we may proceed with the same DMD procedure described in Section 2.1.3, but take the

columns of the matrices \mathbf{Y} and $\mathbf{Y}^\#$ to be pairs $(\psi(\mathbf{y}_i), \psi(\mathbf{y}_i^\#))$. The EDMD matrix (which can be viewed as a finite-dimensional approximation to the Koopman operator) is then given by

$$\mathbf{A} = \mathbf{Y}^\# \mathbf{Y}^+ = \left[\psi(\mathbf{y}_1^\#) \ \psi(\mathbf{y}_2^\#) \cdots \psi(\mathbf{y}_m^\#) \right] \left[\psi(\mathbf{y}_1) \ \psi(\mathbf{y}_2) \cdots \psi(\mathbf{y}_m) \right]^+. \quad (5.2.1)$$

While $\psi(\mathbf{Y}) = \mathbf{U}_r^* \mathbf{Y}$ gives an optimal transformation of the data from an energetic perspective (and is what Equation (2.1.7) represents in the EDMD framework), it might not be a suitable transformation for correct identification of the dynamics. To this end, we use the form of Galerkin projection models to guide an alternative transformation. Expressing the POD coefficients as before by $\mathbf{a} = \mathbf{U}_r^* \mathbf{y}$, we let

$$\mathbf{q} = \psi(\mathbf{y}) = \begin{bmatrix} \mathbf{a} \\ \text{vec}(\mathbf{a} \otimes \mathbf{a}) \end{bmatrix}. \quad (5.2.2)$$

Here $\text{vec}(\mathbf{a} \otimes \mathbf{a})$ denotes a vector of all non-redundant quadratic couplings between POD coefficients, i.e. $(a_1)^2$, $a_1 a_2$, $(a_2)^2$, etc. If we keep r POD modes, then we have observables of dimension $k = r + r(r+1)/2$, with the possible addition of an additional constant observable to account for the mean of the data. We will not closely concern ourselves with how closely we may approximate the true Koopman operator using such a choice of observables, but will rather show that, in any case, the elements of the identified dynamics on this space of observables may be used to construct a nonlinear model of the system dynamics. To this end, we start by explicitly writing the identified dynamics using this approach by

$$\begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_r \\ (a_1)^2 \\ a_1 a_2 \\ \vdots \\ (a_r)^2 \end{bmatrix}_{t+\Delta t} = \begin{bmatrix} l_1^1 & l_1^2 & \cdots & l_1^r & q_1^{11} & q_1^{12} & \cdots & q_1^{rr} \\ l_2^1 & l_2^2 & \cdots & l_2^r & q_2^{11} & q_2^{12} & \cdots & q_2^{rr} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ l_r^1 & l_r^2 & \cdots & l_r^r & q_r^{11} & q_r^{12} & \cdots & q_r^{rr} \\ l_{11}^1 & l_{11}^2 & \cdots & l_{11}^r & q_{11}^{11} & q_{11}^{12} & \cdots & q_{11}^{rr} \\ l_{12}^1 & l_{12}^2 & \cdots & l_{12}^r & q_{12}^{11} & q_{12}^{12} & \cdots & q_{12}^{rr} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ l_{rr}^1 & l_{rr}^2 & \cdots & l_{rr}^r & q_{rr}^{11} & q_{rr}^{12} & \cdots & q_{rr}^{rr} \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_r \\ (a_1)^2 \\ a_1 a_2 \\ \vdots \\ (a_r)^2 \end{bmatrix}_t. \quad (5.2.3)$$

In this notation, the superscripts denote the given (linear or quadratic) term that is influencing the variable (i.e., POD coefficient, or quadratic monomial thereof) identified by the subscripts. Focusing on a single POD coefficient (i.e., one of the first r rows of Equation 5.2.3), we have that

$$a_k(t + \Delta t) = \sum_{i=1}^r l_k^i a_i(t) + \sum_{\substack{i,j=1 \\ j \leq i}}^r q_k^{ij} a_i(t) a_j(t), \quad (5.2.4)$$

which is the same form as Equation (2.1.5), but in discrete-time. Thus, by taking the first r rows of the \mathbf{A} matrix obtained from performing EDMD with observables given by Equation (5.2.2), we may obtain a system of nonlinear equations that can accurately model the evolution of the POD coefficients of the system in discrete time.

While we may expect that the first r rows of \mathbf{A} give an accurate model for the evolution of POD coefficients, in general we should not expect the same for the evolution of the quadratic monomials of POD coefficients. This is because the equations for the evolution of these terms should involve cubic terms, which are not spanned by our observables. Incidentally, this suggests that using any basis of polynomial observables for approximating the Koopman operator for the Navier-Stokes problems might be problematic. One possible alternative, not explored in this work, is to use the Kernel variant of EDMD [173] with an appropriately chosen kernel function that better spans the Koopman eigenfunctions.

5.2.2 A modification to DMD/EDMD

When in the least squares regime (i.e., when there are fewer observables than the number of snapshot pairs), Equations (2.1.6) and (5.2.1) can give solutions with large entries in \mathbf{A} . Empirically, these entries can be significantly larger in magnitude than those expected, say, from performing a Galerkin projection. It also appears that this “overfitting” can give models that lack stability. To mitigate these observations, we propose a simple modification to DMD/EDMD that penalizes the size of entries in \mathbf{A} . Note that the same process can be applied to both DMD and EDMD, whenever the data matrix \mathbf{Y} has more columns than rows. When this is the case, Equation (5.2.1) gives the solution to the minimization problem

$$\mathbf{A} = \operatorname{argmin}_{\mathbf{M}} \|\mathbf{M}\mathbf{Y} - \mathbf{Y}^\# \|_F^2. \quad (5.2.5)$$

We may add a penalization on the size of the entries of \mathbf{A} to formulate a joint least squares problem

$$\mathbf{A} = \operatorname{argmin}_{\mathbf{M}} (\|\mathbf{M}\mathbf{Y} - \mathbf{Y}^\# \|_F^2 + \beta^2 \|\mathbf{M}\|_F^2), \quad (5.2.6)$$

where β is a parameter that determines the extent to which large entries in \mathbf{M} are penalized. The element-wise nature of the Frobenius norm means that Equation (5.2.6) may be rearranged to give

$$\mathbf{A} = \operatorname{argmin}_{\mathbf{M}} \|\mathbf{M}[\mathbf{Y} \ \beta \mathbf{I}] - [\mathbf{Y}^\# \ \mathbf{0}]\|_F^2, \quad (5.2.7)$$

where \mathbf{I} and $\mathbf{0}$ are appropriately sized matrices. Note that this modification of DMD/EDMD is essentially a form of Tikhonov regularization. Equation (5.2.7) has an explicit solution, given by

$$\mathbf{A} = \mathbf{Y}^\# [\mathbf{Y} \ \beta \mathbf{I}]^+. \quad (5.2.8)$$

We remark that this regularization is equivalent as adding to the data set pairs of snapshots where each observable goes from some nonzero value β to 0 over a timestep of Δt . It should also be noted that Tikhonov regularization methods have been used previously as a method of calibrating GP models [37], and have also been used for other system identification techniques, such as finite impulse response models [83]. More generally, one can add a variety of different forms of penalty terms to obtain a desired balance between competing objectives. For example, in the context of EDMD, Williams et al. [174] uses a $L_{1,2}$ minimization to obtain what in our notation would be a solution with a small number of nonzero columns of \mathbf{A} .




Figure 5.1: Computational domain used for numerical simulation of flow past a circular cylinder. Dashed lines represent the borders of each nested grid. The gray shading denotes the region from which data was collected for modeling purposes

5.3 Example: flow past a circular cylinder

We test our proposed method using the much-studied example of 2D flow past a circular cylinder. Beyond a critical Reynolds number of approximately 47 [117], the equilibrium becomes unstable and the system will instead converge upon a limit cycle characterized by periodic vortex shedding. We take data for $Re = 60$, with the initial condition close to the unstable equilibrium. The data captures the initial growth of an instability near the unstable equilibrium, through to convergence to the limit cycle. This is an example where regular DMD will fail (in the sense of identifying an appropriate reduced-order model), since the process (which is a Hopf bifurcation) is fundamentally nonlinear. In particular, beyond the critical Reynolds number, the nonlinear terms must become non-negligible to balance the growth of the unstable linear dynamics, leading to the observed limit cycle behavior.

We will explore the performance of both EDMD and GP models on this system for a range of data, including that which is noisy and spatially truncated or sparse. The data was obtained from direct numerical simulation using an immersed boundary projection method [149, 36], which was described in more detail in Section 4.2. Selective frequency damping [3] was used to obtain the unstable equilibrium solution. To focus on the transitional region of the dynamics, the snapshots to be used were collected after first running the simulation from the unstable equilibrium for $250\frac{D}{U}$ time units. The simulations were performed on a domain consisting of five nested grids, as shown in Figure 5.1. Each grid is uniform, with the finest grid having a grid spacing of $0.02D$ in each direction, and each successively larger grid having double the grid spacing of the previous. The full domain has size $256D \times 64D$. This large domain was chosen so as to resolve both the flowfield on the region of the grid used for analysis (shown in gray), and the forces incident on the cylinder, to a high degree of accuracy.

This comparison between the performance of EDMD and GP models begins in a scenario where both models work relatively well: where clean data encompassing a large spatial domain is available across a window of time spanning all of the distinct dynamic regimes. While we find that GP models can indeed outperform the EDMD models that we identify

in favorable conditions, we will proceed to demonstrate that in the cases where the data is noisy or restricted, EDMD models can perform substantially better.

5.3.1 Data arrangement and selection

The data to be used was taken from the gray shaded region in Figure 5.1. As mentioned previously, we evolve the system from the equilibrium for some time ($250\frac{D}{U}$) before collecting data, then collect 1000 snapshots separated by a uniform timestep $\Delta t = 0.2\frac{D}{U}$. This time interval spans the growth of the instability from near the equilibrium, through to the convergence of the flow to a periodic limit cycle. We note that Galerkin projection in particular is quite sensitive to the resolution and extent of data chosen. For this reason, we do not claim that the GP models that we identify are the most accurate that can be obtained for such a system, but still serve as a basis for comparison to the EDMD models that are identified for the same choice of data.

The EDMD procedure requires a selection of observables. Despite narrowing down this choice substantially by choosing to work with linear and quadratic monomials of POD coefficients, there can still be some additional ambiguity that should be explicitly clarified. To begin with, one must decide how whether to subtract the mean flow from the data before applying POD. This step is almost always performed when performing POD, partly because then any reconstruction of the flow using a linear combination of POD modes will automatically satisfy the required boundary conditions [73]. Conversely, it is almost never done with DMD; doing so can lead to an undesirable equivalence to taking a discrete Fourier transform [28]. Furthermore, if one is to subtract a “mean”, for this flow one could conceivably take this to be any of the mean of the limit cycle, the mean of the data, or even the unstable equilibrium velocity field. To emphasize the fact that we are approaching this procedure from a data-driven perspective, and to be consistent with the subspace used for both procedures, here we first subtract the mean of the data before performing POD (for both GP and EDMD models). Note that this is different to what is typically done when constructing GP models for such a system, where the mean is most typically taken to be the mean of the limit cycle [e.g., 106]. The mean and POD modes identified in this manner are shown in Figure 5.2, with their relative energy content shown in Figure 5.3.

We further note that subtracting the mean of the velocity field before performing POD does not result in all of the observables used for EDMD having zero mean, as $(a_i)^2$ is always non-negative. When performing EDMD, we can also choose whether to explicitly include this mean mode as an observable, which amounts to allowing for a constant term in Equation 5.2.3. We will most often choose to do this, though will discuss this in more detail later. On top of all of these details, the regularization introduced in Section 5.2.2 introduces an additional parameter whose value must be set, which we will discuss in the next section.

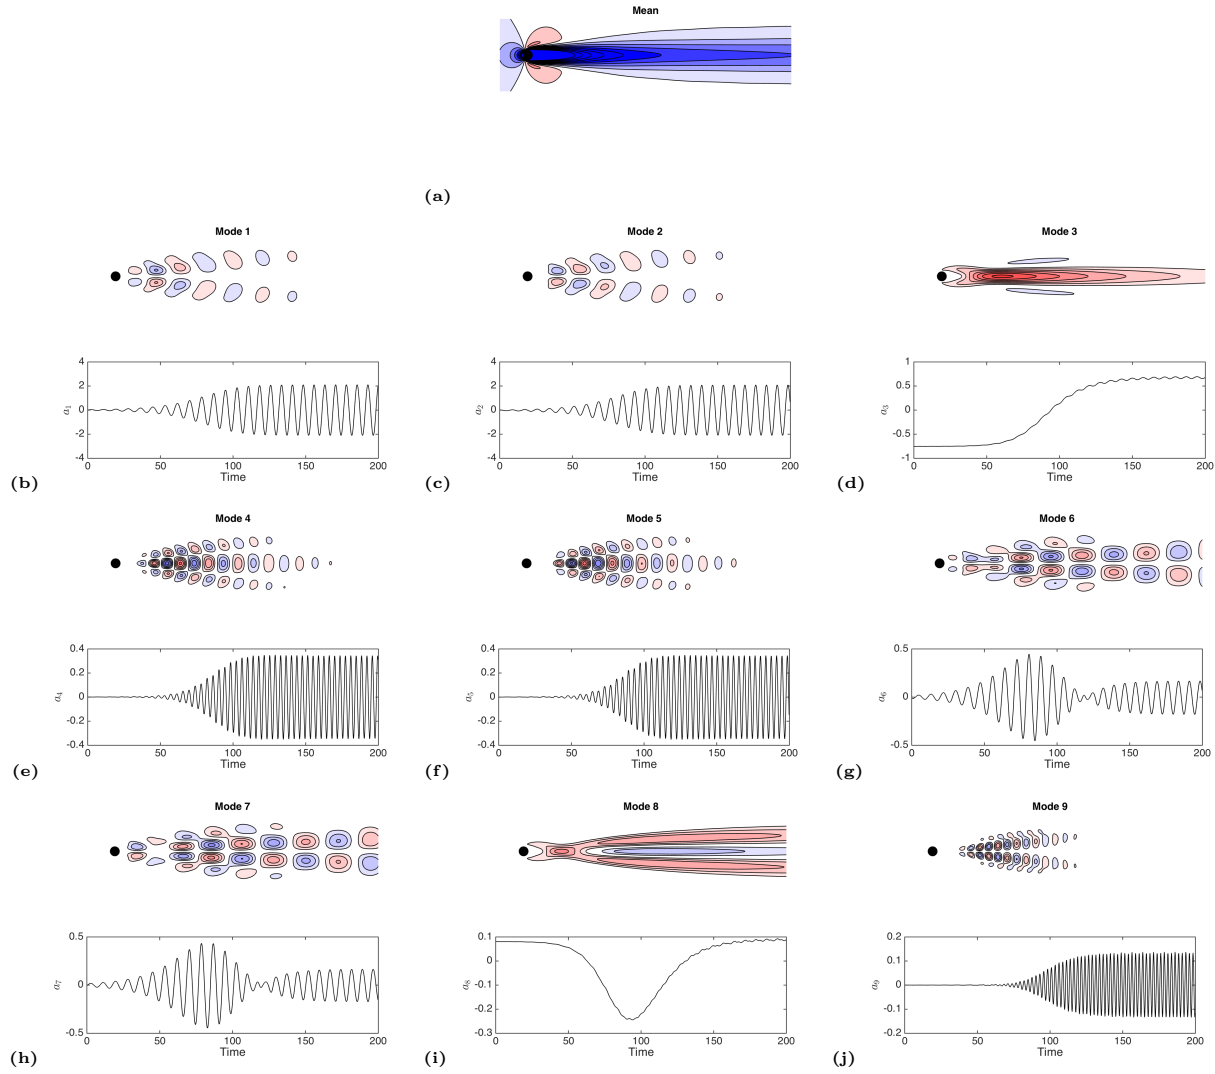


Figure 5.2: Contours of streamwise velocity for the (a) mean mode and (b)-(j) first nine POD modes, along with the corresponding time-varying coefficients of the POD modes identified from data collected as the flow transitions from near the unstable equilibrium to the limit cycle




Figure 5.3: **(a)** Relative and **(b)** cumulative energy content of the first 30 POD modes, for the same data used in Figure 5.2

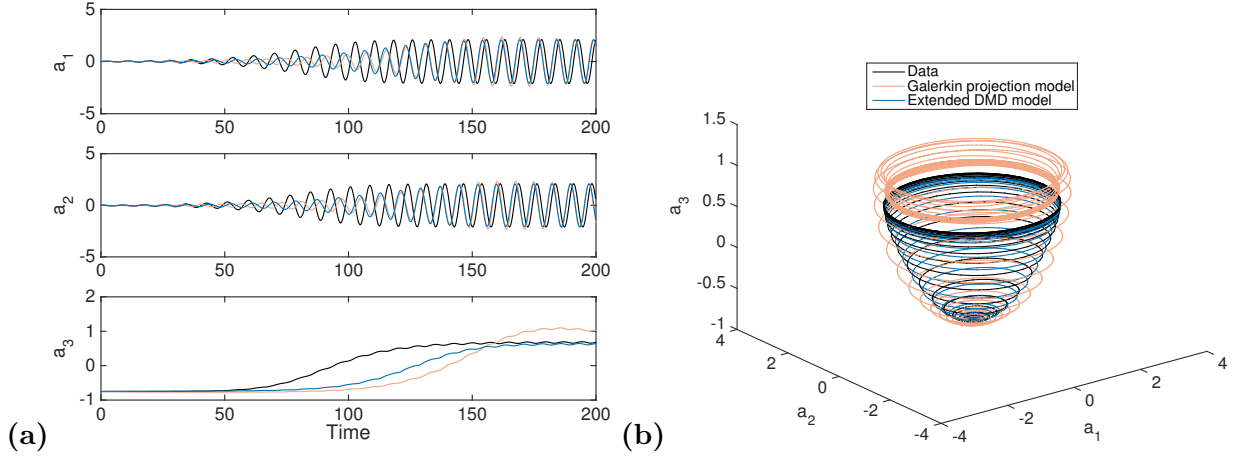


Figure 5.4: Performance of 3rd order GP and EDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, showing **(a)** time evolution and **(b)** phase portrait plots. The Galerkin and EDMD phase portrait models are allowed to evolve for 800 dimensionless time units to confirm limit cycle behavior

5.3.2 Comparison between EDMD and GP models, regularization, and model order dependence

This section analyzes the performance of both EDMD and GP models of a range of orders in modeling the dynamics of the flow past a cylinder. We begin with some sample results. Figure 5.4 shows the performance of both Galerkin projection, and the EDMD approach outlined in Section 5.2, in identifying a model that can predict the evolution of the first three POD coefficients. For this system, the dynamics of these coefficients are known to evolve on a paraboloid [106]. We observe that the EDMD model is more accurate than the GP model in terms of obtaining both the correct transient and limit cycle behavior. Figure 5.5 shows that the same findings hold when the dimension of the models increase to

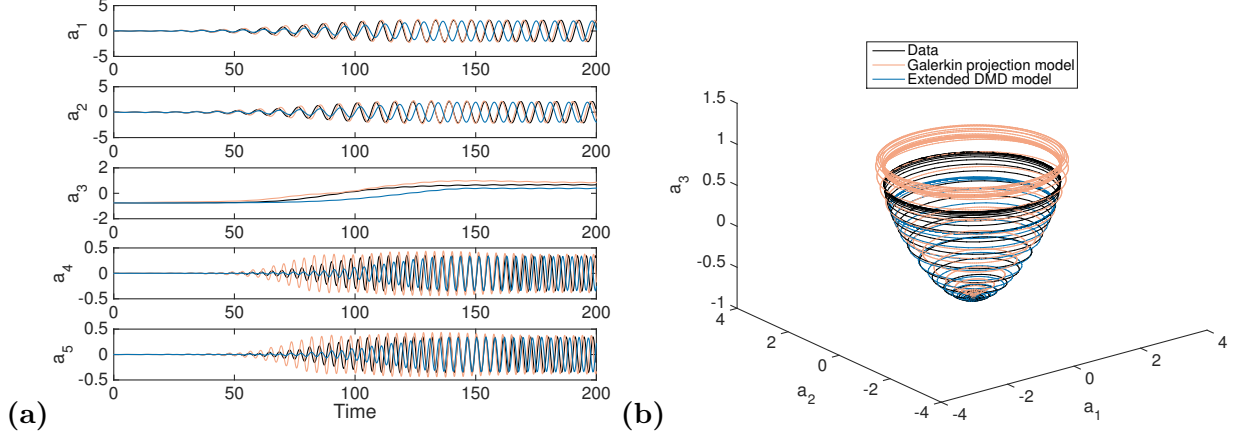


Figure 5.5: Performance of 9th order GP and EDMD nonlinear models in predicting the time-evolution of POD coefficients for transitional flow past a cylinder, showing (a) time evolution and (b) phase portrait plots. Only the first 5 modes are shown on the left, while the phase portrait shows the projection onto the first 3 modes. The Galerkin and eEDMD phase portrait models are allowed to evolve for 800 dimensionless time units to confirm limit cycle behavior

9. We keep β at a fixed value of 0.5 for both of these examples, which will be the default unless otherwise mentioned.

In order to study more systematically the accuracy of models of various order, we compare in Figure 5.6 the identified limit cycle amplitude (defined based on the coefficient of the first POD mode) and frequency for models of order 3–28. As well as showing results for EDMD models identified with $\beta = 0.5$, we additionally show the performance of models identified with an optimized value of β , where the optimal is found based on a direct search over the range $0 \leq \beta \leq 2$. This comparison shows that the results are relatively insensitive to β , and that the one initially chosen value tends to perform reasonably well across all model orders. It would be possible, however, to develop more sophisticated methods to tune β using the data available. To do this, one could apply standard extrema-seeking algorithms to find a value of β that allows the identified model to best match the data, by some chosen metric. One could also explore the possibility of replacing the $\beta \mathbf{I}$ term in Equation 5.2.7 with a diagonal matrix with different entries, though this then loses the direct connection to the joint least squares problem formulated in Equation 5.2.6.

In general, EDMD models are more accurate than GP models at predicting the limit cycle characteristics for models of order less than 8, while GP models are more accurate for high order, except for models of order 14, where the GP model performs uncharacteristically poorly. While for this simple system the high order models are not required to obtain an accurate representation of the system dynamics, it is important to verify that the proposed algorithm remains capable of identifying stable and accurate models as the model order increases, in order to be of use for more complex fluids systems.

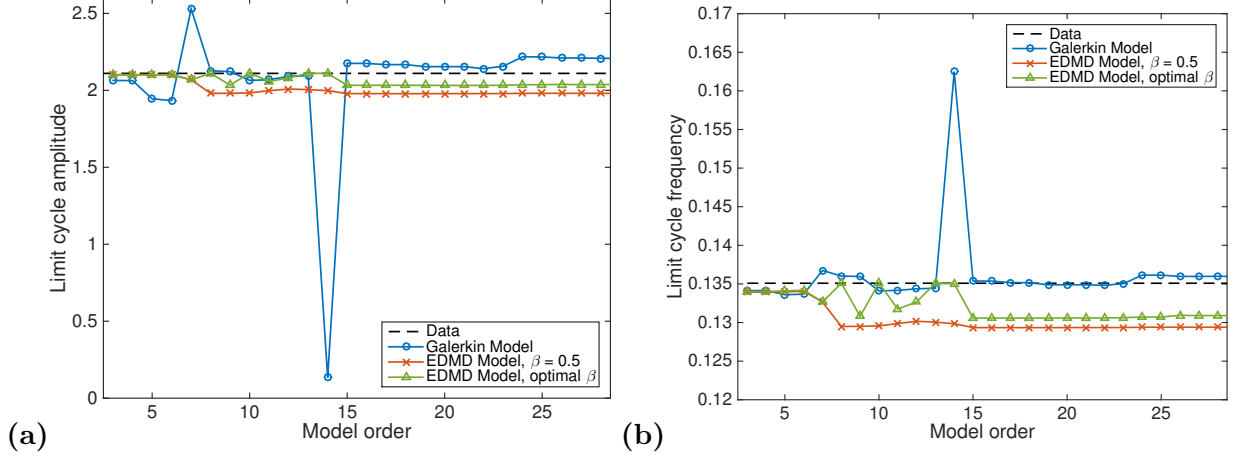


Figure 5.6: Prediction of the limit cycle amplitude (a) and frequency (b) of the first POD mode from both EDMD and GP models

5.3.3 Model prediction for untrained conditions

An important feature of any identified model is the ability to predict the behavior of the system along trajectories in phase space that are not contained in the data used for model identification. This ensures the practical utility of such a model, and increases confidence that the “correct” system dynamics are being captured. We show in Figure 5.7 the performance of 3rd order EDMD and GP models in predicting the evolution of the first 3 POD coefficients, starting at the mean of the limit cycle rather than near the unstable equilibrium. We observe that the GP model more accurately captures the initial transient as the flow approaches the slow manifold (paraboloid), though as before the EDMD model is more accurate in predicting the subsequent approach to the limit cycle. Note that there are two main reasons why this model might be somewhat inaccurate in the untrained region: not only are the dynamics in this region untrained, but the POD basis is additionally no longer energetically optimal (in terms of energy) for this dataset, so low order models in particular might not capture features that are both energetically and dynamically important.

5.3.4 Noisy data

An important quality for any modeling procedure to possess is a robustness to noisy data, such as that which might be acquired from experiments. In Figure 5.8, we show the results of identifying 3rd order models to the same data as considered previously, but corrupted by Gaussian white noise, with magnitude of 5% of the freestream velocity. It is observed that the EDMD model retains its accuracy, whereas the GP model captures none of the observed features of the system. This result is interesting, since the noise-corrupted data is actually being used to identify the dynamics in the case of EDMD, whereas GP only uses the data to obtain a spatial basis. As well as showing the robustness of the EDMD modeling procedure, these results highlight the sensitivity of GP to this basis selection.

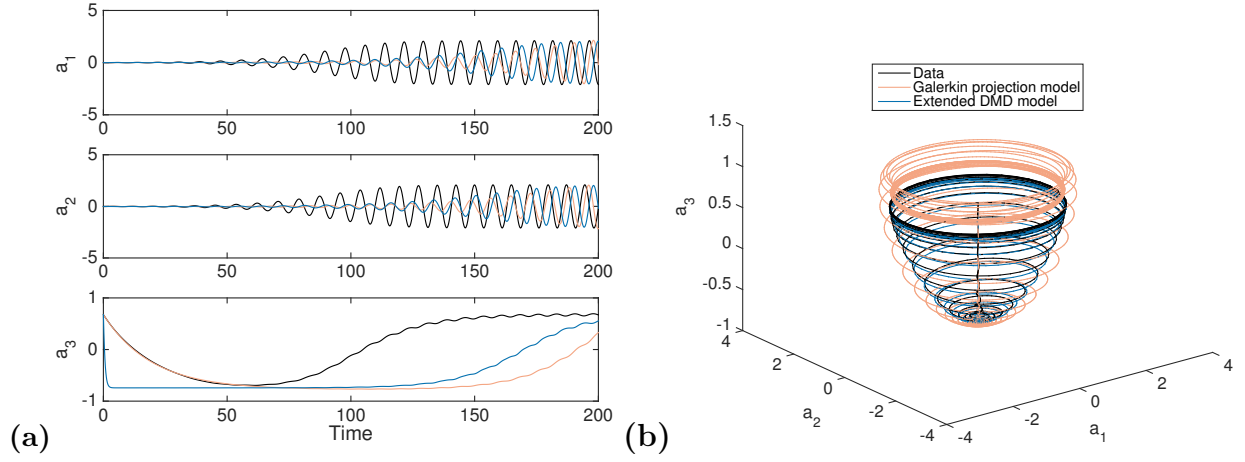


Figure 5.7: Performance of 3rd order GP and EDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, showing **(a)** time evolution and **(b)** phase portrait plots. The initial condition is taken to be the mean of the limit cycle

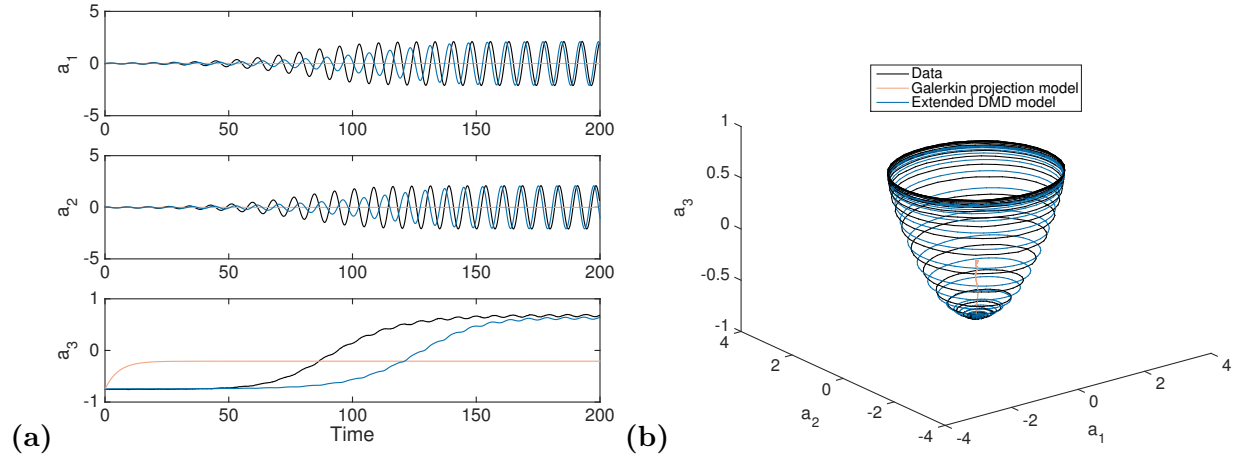


Figure 5.8: Performance of 3rd order GP and EDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, showing **(a)** time evolution and **(b)** phase portrait plots. The Galerkin and EDMD phase portrait models are identified from data that is corrupted by noise of standard deviation $\sigma = 0.05U$. $\beta = 0.5$

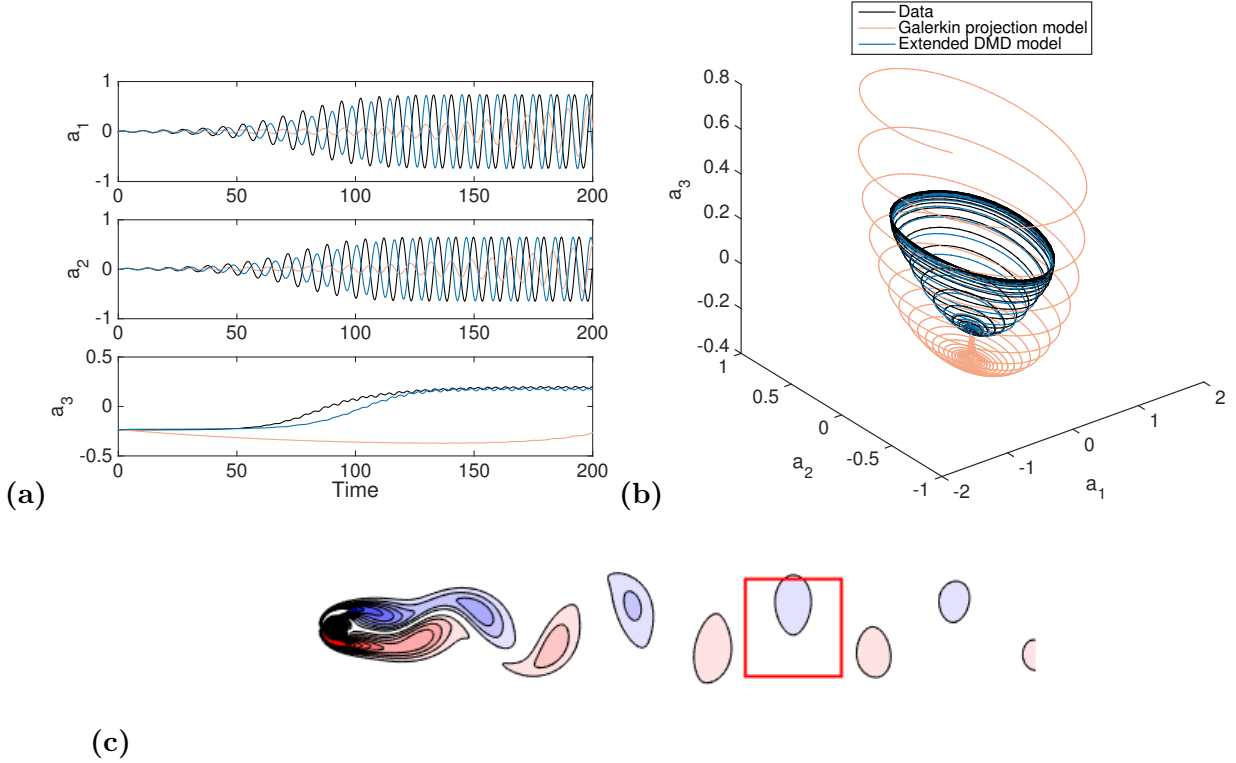


Figure 5.9: Performance of 3rd order GP and EDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, showing (a) time evolution and (b) phase portrait plots. Only data contained from inside the square shown in (c) is used for model identification. The EDMD models were obtained with $\beta = 0.1$

5.3.5 Data from a restricted spatial domain

Another potential drawback of GP models is that they require spatially resolved data across a wide domain, ideally such that the boundary conditions are constant. In Figure 5.9, we attempt to obtain reduced-order models using a small portion of the domain. Unlike GP, EDMD is still able to produce an accurate model with such a restricted data set. Here, we are only using the data collected from this limited domain to compute POD, so the modes that are computed will differ somewhat from the modes shown in Figure 5.2. The fact that EDMD models in this case retain their accuracy for both noisy data and restricted spatial domains suggests that this modeling framework could be particularly useful for experimental data, which often possesses both limitations.

An advantage of EDMD is that it does not require spatially resolved data, which is required for the computation of spatial derivatives in GP. In Figure 5.10, we show that EDMD can produce an accurate ROM with only data from a small number of sparse, randomly chosen points. This gives cause for optimism that the method could additionally be used in experiments for which only point sensor measurements are available, rather than the spatially resolved flowfields, which can be more costly and time consuming to obtain. While the selection of the sensor locations in this case was random, we observed empirically that best

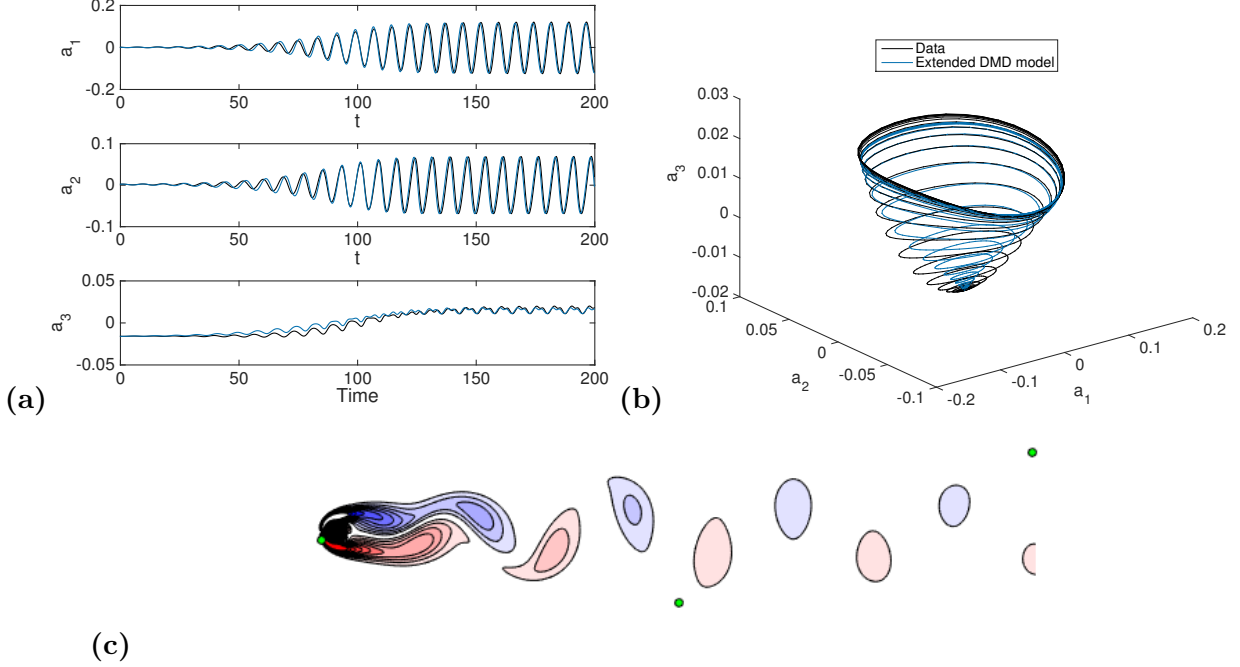


Figure 5.10: Performance of 3rd order GP and EDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, showing (a) time evolution and (b) phase portrait plots. Only data sampled at green circles shown in (c) are used for model identification

results were obtained when the sensor locations were spread across different spatial regions of the domain. When using only a small number of point sensors, we found that better models were obtained when we did not include a constant term in our set of observables. We additionally did not use any form of regularization.

5.3.6 Data from limited temporal sampling

In this section, we explore the ability of EDMD and GP models in predicting the evolution of a system using only a small amount of temporal data. As in Section 5.3.5, we compute POD modes (including the mean modes) from the subset of data that is used. Note in particular that this means that the POD basis that we use as our state space will be different to those identified from the full dataset. In Figure 5.11, we show the results of applying both EDMD and GP when using only the first 40%, 20%, and 10% of the time series of data used for the previous sections. We find that, even when the data available is significantly reduced and has clearly not reached the limit cycle, the EDMD models still accurately predict the presence (and to some extent, the location) of the limit cycle. This is in contrast to the GP model, where the limited data available renders the model qualitatively incorrect for all amounts of data used. Note that the results in this section did not use any form of regularization.

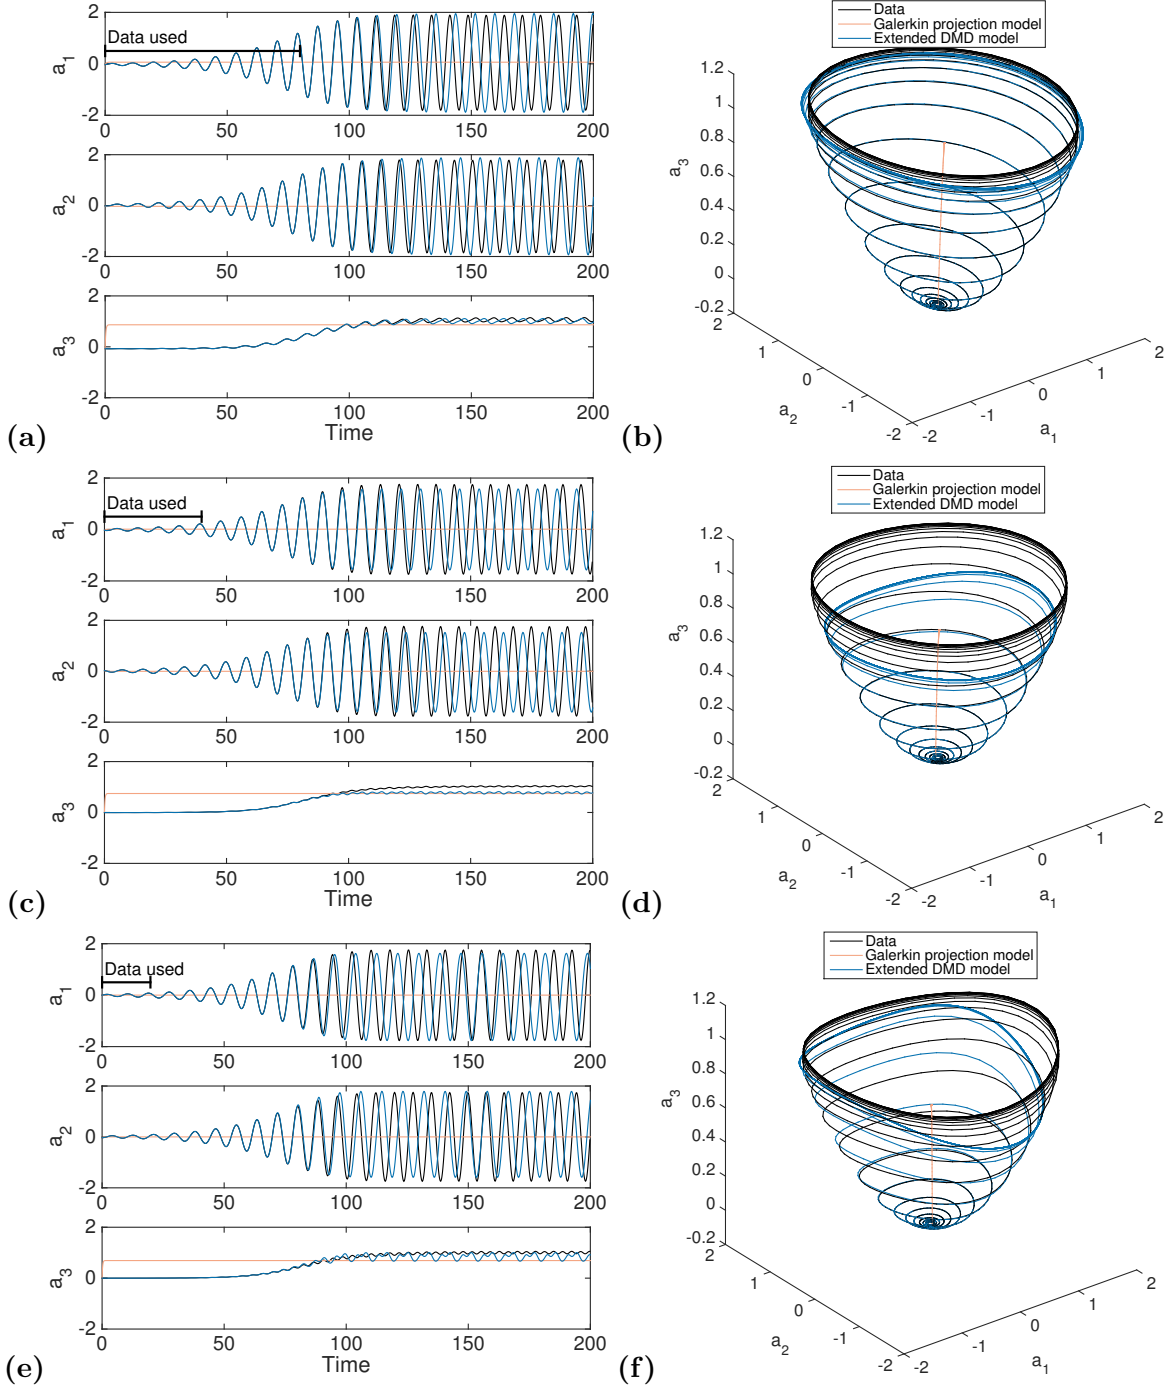


Figure 5.11: Performance of 3rd order GP and eDMD nonlinear models in predicting the evolution of POD coefficients for transitional flow past a cylinder, showing **(a,c,e)** time evolution and **(b,d,f)** phase portrait plots. Models are identified using only the first 400 (top), 200 (center) and 100 (bottom) snapshots of data, as shown

5.4 Discussion and conclusions

This chapter has introduced a method to extract low-order, nonlinear models from time-resolved data, by utilizing an extension of the DMD algorithm. Through the use of a variety of different types of data from a simple example of a nonlinear fluids system, we have attempted to evaluate the relative performance of this method in across a range of circumstances.

We believe that the advantages of the present approach are numerous. We have demonstrated that this data-driven approach is more robust to noisy data, and exhibits greater flexibility in terms of the spatial extent and resolution of data for which the method produces accurate models. We therefore hold strong hopes that this method will be useful in particular to model experimental data, which often contains such limitations.

Beyond this, the present method requires no explicit knowledge of the underlying governing equations, though here knowing at least the form of the nonlinearity was advantageous. In particular, note that the correct form of the observables is most likely important when the nonlinear model is used to extrapolate a small amount of data, and ultimately correctly predict the limit cycle in Section 5.3.6. In other words, if the correct form of the nonlinearity is known, then even a small sample of data can be sufficient to identify a model that is at least qualitatively accurate. This finding has important implications for more complicated systems, where sampling across all possible regimes and/or locations in phase space may be infeasible.

In Section 5.3.1, it was observed that EDMD models are particularly accurate when the dimension of the model is small, in comparison to both GP models of the same order, and EDMD models of higher order. The fact that GP models are less accurate for lower model orders is unsurprising, since (without any additional modification) they have no way of accounting for the effect of unmodeled modes. EDMD models, on the other hand, can at least account for the time-averaged effect of modes that are not explicitly included as variables. Note, however, that unmodeled modes with highly intermittent dynamics could be impossible to accurately account for. As the model order increases, GP models should be expected to become more accurate, as the effect of unmodeled modes reduces. Conversely, we find that EDMD models can become less accurate and robust, possibly due to the rapid increase in the number of parameters requiring to be fit. Note that a phenomena has been observed previously in polynomial identification [112].

The mathematical machinery behind DMD and it's variants involves a least squares (or minimum norm, in the underconstrained case) fit of data. In this sense, the method presented in this work shares strong similarities with a number of other methods that have been used to identify nonlinear systems in fluids, such as polynomial identification [112], modified quadratic stochastic estimation [104], Volterra system identification [12], and linear parameter varying system identification [70], which all employ a similar type of least-squares coefficient-fitting methodology in their respective algorithms. The recently developed SINDy [25] (sparse identification of nonlinear dynamics) also employs a similar basic framework, though with an additional sparsity-promoting algorithm to decide upon the appropriate observables from a larger library, which is particularly useful when the form of the nonlinearity is unknown. In this sense, we do not claim that the algorithm we use is exceptionally novel, but hope that this exposition, both in terms of the connections of DMD and Koopman

theory, and in highlighting a number of situations where such a modeling approach might be advantageous or necessary. On the former point, one could hope that EDMD, with the correct choice of observables, can allow for accurate prediction not just of Koopman eigenvalues and eigenmodes, but also Koopman eigenfunctions. There is indeed strong evidence to suggest that identified Koopman eigenfunctions can indeed give an accurate characterization of the dynamics of the cylinder wake near the limit cycle [8]. We do not make any claims that we have achieved such convergence here over the full transient regime, but rather show that the results can still be useful from a practical, model-reduction standpoint.

As well as pursuing this direction, further work could seek to utilize such models for purposes of flow control. Preliminary investigations also suggest that such EDMD-based models can be effective at identifying stability properties, such as eigenvalues of the linearized system at equilibrium points.

Chapter 6

Characterizing and correcting for the effect of sensor noise in the dynamic mode decomposition

Dynamic mode decomposition (DMD) provides a practical means of extracting insightful dynamical information from fluids datasets. Like any data processing technique, DMD’s usefulness is limited by its ability to extract real and accurate dynamical features from noise-corrupted data. In this chapter, we show analytically that DMD is biased to sensor noise, and quantify how this bias depends on the size and noise level of the data. We present three modifications to DMD that can be used to remove this bias: (i) a direct correction of the identified bias using known noise properties, (ii) combining the results of performing DMD forwards and backwards in time, and (iii) a total least-squares-inspired algorithm. We discuss the relative merits of each algorithm, and demonstrate the performance of these modifications on a range of synthetic, numerical, and experimental datasets. We further compare our modified DMD algorithms with other variants proposed in recent literature.

6.1 Introduction

One of the major advantages of DMD over techniques such as global stability analysis, for example, is that it can be applied directly to data, without the need for the knowledge or construction of the system matrix, which is typically not available for experiments [134]. For this reason, analysis of the sensitivity of DMD to the type of noise typically found in experimental results is of particular importance. The effects of noise on the accuracy of the DMD procedure was systematically investigated in the empirical study by Duke et al. [48], for the case of a synthetic waveform inspired by canonical periodic shear flow instabilities. More recently, Pan et al. [110] have extended this type of analysis to more complex data with multiple frequencies, as might be found in typical fluids systems.

This chapter builds upon these previous studies by analytically deriving an expression that explicitly shows how DMD should be affected by noise, for the case where the noise is assumed to be sensor noise that is uncorrelated with the dynamics of the system. Our analysis complements the “noise-robust” DMD formulation in Hemati et al. [69] by explicitly

quantifying the influence of noise on DMD. Further, while our analysis is consistent with the total least-squares formulation in Hemati et al. [69], we use the insights gained from our analysis to develop alternative techniques to total least-squares DMD that may be preferable in certain applications. Ultimately, the availability of multiple “noise-aware” DMD algorithms allows the user to approach dynamical analysis of noisy data from multiple angles, thus garnering more confidence in the computations. We note that the case of process noise, where noise can interact with the dynamics of the system, is also the subject of recent work [9].

The analysis will use a recent characterization of DMD [161], described in Section 2.1.3, which, along with other advantages, highlights the connection of DMD to related techniques that are used in other communities for the extraction of dynamical information from data. Many linear system identification techniques are closely related in that they are based around singular value decomposition of a data matrix; aside from DMD there is the eigensystem realization algorithm [79] (see Section 2.1.4) and balanced proper orthogonal decomposition [122], for example. Indeed, the origin of such an approach seems to date back to the work of Ho and Kalman [72].

In this chapter, we first show that the dominant effect of noise on DMD is often deterministic. This not only allows us to accurately predict its effect, but also allows for a correction to be implemented to recover the noise-free dynamics. As well as directly correcting for the noise, we present two other modifications of DMD, that both are able to remove this bias without needing to know the noise characteristics. Section 6.2 develops the theory that characterizes the effect of noise on DMD, which subsequently motivates the formulation of our modified algorithms, which we term *noise-corrected* DMD (ncDMD), *forward-backward* DMD (fbDMD) and *total least-squares* DMD (tlsDMD). In Section 6.3, we analyze the performance of these algorithms on a number of synthetic data sets, which are corrupted by Gaussian white noise. We additionally investigate how the algorithms perform on data with both sensor and process noise. In Section 6.4, we use numerical and experimental data from flow past a cylinder undergoing periodic vortex shedding, to demonstrate the utility of the proposed modifications of DMD for real fluids data.

6.2 Characterizing noise in dynamic mode decomposition

This section details the methodology that is used to analyze the effect of noise in DMD. The exposition is designed to follow on from the description of DMD given in Section 2.1.3. The effect of sensor noise in the data on the results of DMD is studied in Section 6.2.1, which in particular shows that DMD is biased to sensor noise. Sections 6.2.2–6.2.4 formulate three different modifications of the DMD algorithm that are designed to remove this bias.

6.2.1 Sensor noise in DMD

In this work we use the term sensor noise to describe additive noise that affects only our measurements of a given system, and does not interact with the true dynamics. If we have

a discrete-time dynamical system

$$\mathbf{x}(t + \Delta t) = F[\mathbf{x}(t)],$$

then we assume that our measurements take the form

$$\mathbf{x}_m(t) = \mathbf{x}(t) + \mathbf{n}(t),$$

where $\mathbf{n}(t)$ is a random noise vector. For the purposes of this paper, we will take each component of $\mathbf{n}(t)$ to be independent and normally distributed with zero mean and a given variance. With \mathbf{Y} and $\mathbf{Y}^\#$ as described in Section 2.1.3, suppose that we measure $\mathbf{Y}_m = \mathbf{Y} + \mathbf{N}$ and $\mathbf{Y}_m^\# = \mathbf{Y}^\# + \mathbf{N}^\#$, where \mathbf{N} and $\mathbf{N}^\#$ are random matrices of sensor noise. Note that some (or most) columns of random data in \mathbf{N} might also be in $\mathbf{N}^\#$, but shifted to a different column. We will assume that the noise is independent of the true data, and is independent in both space and time, so that each element of a given noise matrix is a random variable taken from a fixed zero-mean normal distribution. From Equation (2.1.7), the measured DMD matrix $\tilde{\mathbf{A}}_m$ can be computed from

$$\begin{aligned} \tilde{\mathbf{A}}_m &= \tilde{\mathbf{Y}}_m^\# \tilde{\mathbf{Y}}_m^+ = (\tilde{\mathbf{Y}}^\# + \tilde{\mathbf{N}}^\#)(\tilde{\mathbf{Y}} + \tilde{\mathbf{N}})^+ \\ &= (\tilde{\mathbf{Y}}^\# + \tilde{\mathbf{N}}^\#)(\tilde{\mathbf{Y}} + \tilde{\mathbf{N}})^* [(\tilde{\mathbf{Y}} + \tilde{\mathbf{N}})(\tilde{\mathbf{Y}} + \tilde{\mathbf{N}})^*]^{-1} \\ &= (\tilde{\mathbf{Y}}^\# \tilde{\mathbf{Y}}^* + \tilde{\mathbf{N}}^\# \tilde{\mathbf{Y}}^* + \tilde{\mathbf{Y}}^\# \tilde{\mathbf{N}}^* + \tilde{\mathbf{N}}^\# \tilde{\mathbf{N}}^*) \\ &\quad \times \left[\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^* + \tilde{\mathbf{N}} \tilde{\mathbf{Y}}^* + \tilde{\mathbf{Y}} \tilde{\mathbf{N}}^* + \tilde{\mathbf{N}} \tilde{\mathbf{N}}^* \right]^{-1}, \end{aligned} \quad (6.2.1)$$

where we have used the identity $\mathbf{M}^+ = \mathbf{M}^*(\mathbf{M}\mathbf{M}^*)^{-1}$. Note again that here the \sim notation means that the data is expressed in the POD basis obtained from the noisy data. We perform our analysis in this POD space rather than with the original data to allow for truncation of low energy modes, and because the computation of the pseudoinverse \mathbf{Y}^+ can be prohibitive for large datasets. We expect that the presence of noise should result in some error in the computation of $\tilde{\mathbf{A}}_m$ (in comparison to the noise free matrix $\tilde{\mathbf{A}}$) and thus some amount of error in the computed DMD eigenvalues and modes. Since elements of $\tilde{\mathbf{A}}_m$ are statistical quantities dependent on the noise, it will make sense to compute statistical properties of the matrix. We begin by computing $\mathbb{E}[\tilde{\mathbf{A}}_m]$, the expected value of the computed DMD matrix. Provided that we have truncated any POD modes with zero energy, $\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*$ should be invertible. If the noise terms are sufficiently small, then we can make use of the matrix perturbed inverse expansion $(\mathbf{M} + \mathbf{P})^{-1} = \mathbf{M}^{-1} - \mathbf{M}^{-1} \mathbf{P} \mathbf{M}^{-1} + \dots$, where higher order terms will be small for $\mathbf{M} \gg \mathbf{P}$. Equation 6.2.1 then becomes

$$\begin{aligned} \tilde{\mathbf{A}}_m &= (\tilde{\mathbf{Y}}^\# \tilde{\mathbf{Y}}^* + \tilde{\mathbf{N}}^\# \tilde{\mathbf{Y}}^* + \tilde{\mathbf{Y}}^\# \tilde{\mathbf{N}}^* + \tilde{\mathbf{N}}^\# \tilde{\mathbf{N}}^*)(\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1} \\ &\quad \times \left[\mathbf{I} - (\tilde{\mathbf{N}} \tilde{\mathbf{Y}}^* + \tilde{\mathbf{Y}} \tilde{\mathbf{N}}^* + \tilde{\mathbf{N}} \tilde{\mathbf{N}}^*)(\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1} + \dots \right]. \end{aligned} \quad (6.2.2)$$

Taking the expected value of Equation (6.2.2), we may classify the terms into three categories: a deterministic terms that does not involve $\tilde{\mathbf{N}}$ or $\tilde{\mathbf{N}}^\#$ (which ends up being $\tilde{\mathbf{A}}$), terms involving one or three noise matrices, which will have expected values of 0 (e.g.,

$\tilde{\mathbf{N}}^\# \tilde{\mathbf{Y}}^* (\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1}$), and terms which involve two or four noise matrices. It is terms this latter category that may have non-zero expected values, and thus bias the result of applying DMD to noisy data. Discarding terms containing a single noise matrix, and additionally discarding higher order terms from the expansion, we have

$$\begin{aligned} \mathbb{E}(\tilde{\mathbf{A}}_m) &= \tilde{\mathbf{A}}(\mathbf{I} - \mathbb{E}(\tilde{\mathbf{N}} \tilde{\mathbf{N}}^*) (\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1}) + \mathbb{E}(\tilde{\mathbf{N}}^\# \tilde{\mathbf{Y}}^+ \tilde{\mathbf{N}}) \tilde{\mathbf{Y}}^+ \\ &\quad + \mathbb{E}(\tilde{\mathbf{N}}^\# \tilde{\mathbf{Y}}^+ \tilde{\mathbf{Y}} \tilde{\mathbf{N}}^*) (\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1} + \tilde{\mathbf{Y}}^\# \mathbb{E}(\tilde{\mathbf{N}}^* (\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1} \tilde{\mathbf{N}}) \tilde{\mathbf{Y}}^+ \\ &\quad + \tilde{\mathbf{Y}}^\# \mathbb{E}(\tilde{\mathbf{N}}^* (\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1} \tilde{\mathbf{Y}} \tilde{\mathbf{N}}^* (\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1}) \\ &\quad + \mathbb{E} \left[\tilde{\mathbf{N}}^\# \tilde{\mathbf{N}}^* (\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1} (\mathbf{I} - \tilde{\mathbf{N}} \tilde{\mathbf{N}}^* (\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1}) \right], \end{aligned} \quad (6.2.3)$$

where we have noted that $\tilde{\mathbf{Y}}^\# \tilde{\mathbf{Y}}^* (\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1} = \tilde{\mathbf{Y}}^\# \tilde{\mathbf{Y}}^+ = \tilde{\mathbf{A}}$. Assuming that the noise is sufficiently small compared with the true data, we can further neglect the term involving four noise matrices. The largest of the remaining terms will be that which contains the product $\tilde{\mathbf{N}} \tilde{\mathbf{N}}^*$. The remaining terms do not necessarily have zero mean, but for the purposes of this investigation will be neglected. Our results will demonstrate that this simplification is justifiable. This reduces Equation (6.2.3) to the following expression, relating the identified and true DMD matrices:

$$\mathbb{E}(\tilde{\mathbf{A}}_m) = \tilde{\mathbf{A}}(\mathbf{I} - \mathbb{E}(\tilde{\mathbf{N}} \tilde{\mathbf{N}}^*) (\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^*)^{-1}). \quad (6.2.4)$$

It might seem surprising that Equation (6.2.4) contains \mathbf{N} , but not $\mathbf{N}^\#$. The reason for this will become apparent in Section 6.2.4, where casting DMD in an optimization framework shows that the standard algorithm is optimal only when assuming that all of the noise is in $\mathbf{Y}^\#$, but not \mathbf{Y} . From a mathematical point of view, it is because the expression $\tilde{\mathbf{A}} = \tilde{\mathbf{Y}}^\# \tilde{\mathbf{Y}}^+$ is linear in $\tilde{\mathbf{Y}}^\#$ but not in $\tilde{\mathbf{Y}}$, which is why perturbations to $\tilde{\mathbf{Y}}$ do not have to propagate through the equation in an unbiased manner. Note that the same analysis can be performed without transforming into POD space (i.e., without the $\tilde{\cdot}$ notation), with the analogous expression to Equation (6.2.4) being

$$\mathbb{E}(\mathbf{A}_m) = \mathbf{A}(\mathbf{I} - \mathbb{E}(\mathbf{N} \mathbf{N}^*) (\mathbf{Y} \mathbf{Y}^*)^{-1}), \quad (6.2.5)$$

subject to $\mathbf{Y} \mathbf{Y}^*$ being invertible. For systems where the size of each snapshot is larger than the number of snapshots (i.e, $n > m$, which is typical for fluids systems), $\mathbf{Y} \mathbf{Y}^*$ will not be invertible, thus motivating our choice to work in POD space. Moreover, one might want the option to truncate all but a certain number of POD modes, in order to obtain a low-dimensional model for the dominant system dynamics. Up until this point, we have not made a distinction between the POD modes of the clean data, \mathbf{Y} , and the noisy measured data, \mathbf{Y}_m , with the latter typically being all that we have access to. This issue will be explicitly addressed in Section 6.2.2.

Equation (6.2.4) shows that DMD is biased to sensor noise. In practice, the importance of this finding will depend on how the magnitude of this bias compares to the random component of error, that will fluctuate with different samples of random noise. Figure 6.1 shows an illustration of how bias and random components of error contribute to the total




Figure 6.1: Illustrative diagram showing how the error in estimation of a given quantity can be decomposed into bias error (being the difference between the true and expected value of the identified quantity), and random error (representing the fluctuation in the estimated quantity between different noise realizations)

error in the estimation of some quantity from noisy data. The appendix to this chapter (Section 6.6) provides scaling arguments that suggest that the bias will be the dominant component of error in DMD whenever $m^{1/2}SNR > n^{1/2}$, where SNR is the signal-to-noise ratio. When this is the case, it would be particularly advantageous if one had access to a bias-free alternative to DMD. The remainder of this section will present a number of such alternatives.

6.2.2 Direct correction of sensor noise bias in DMD

Referring back to Equation (6.2.4), we can form a bias-free estimate of the true DMD matrix $\tilde{\mathbf{A}}$ via

$$\tilde{\mathbf{A}} \approx \tilde{\mathbf{A}}_m (\mathbf{I} - \mathbb{E}(\tilde{\mathbf{N}}\tilde{\mathbf{N}}^*)(\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^*)^{-1})^{-1}. \quad (6.2.6)$$

Making this modification in practice requires an accurate estimate of both the noise covariance, $\mathbb{E}(\tilde{\mathbf{N}}\tilde{\mathbf{N}}^*)$, and the true data covariance, $\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^*$, in POD space. For noise that is sufficiently small, we can utilize the approximation

$$\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^* = \mathbf{U}^* \mathbf{Y} \mathbf{Y}^* \mathbf{U} \approx \mathbf{U}_m^* \mathbf{Y}_m \mathbf{Y}_m^* \mathbf{U}_m = \Sigma_m^2, \quad (6.2.7)$$

where $\mathbf{U}_m \Sigma_m \mathbf{V}_m^*$ is the singular value decomposition of the noisy data, \mathbf{Y}_m . This allows for us to express the bias of DMD in terms of quantities that are measurable from noisy data. The assumption that $\mathbf{Y} \mathbf{Y}^* = (\mathbf{Y}_m - \mathbf{N})(\mathbf{Y}_m - \mathbf{N})^* \approx \mathbf{Y}_m \mathbf{Y}_m^*$ can be further refined by retaining the $\mathbf{N} \mathbf{N}^*$ term, but for small noise this higher order term will typically be small enough to neglect after being inserted into Equation (6.2.6). The assumption that $\mathbf{U} \approx \mathbf{U}_m$ will largely be justified by means of results that show the utility of this analysis. Analyzing the precise relationship between \mathbf{U} and \mathbf{U}_m in more detail is beyond the scope of this work, and is indeed an active area of research. We direct the interested reader to relevant results in perturbed SVD's [147, 148, 10], random inner product matrices [29, 151], and POD-type operations on noisy data [52, 142, 181].

Assuming that the noise is uniform as well as spatially and temporally independent, then $\mathbb{E}(\tilde{\mathbf{N}}\tilde{\mathbf{N}}^*) = \mathbb{E}(\mathbf{U}^*\mathbf{N}\mathbf{N}^*\mathbf{U}) = \mathbf{U}^*m\sigma_N^2\mathbf{U} = m\sigma_N^2\mathbf{I}$, where σ_N^2 is the variance of each independent component of the noise matrix. With this assumption, and the approximation given in Equation (6.2.7), Equation (6.2.6) becomes

$$\tilde{\mathbf{A}} \approx \tilde{\mathbf{A}}_m(\mathbf{I} - m\sigma_N^2\boldsymbol{\Sigma}_m^{-2})^{-1}. \quad (6.2.8)$$

If the noise is sufficiently small, then a perturbed inverse approximation gives

$$\tilde{\mathbf{A}} \approx \tilde{\mathbf{A}}_m(\mathbf{I} + m\sigma_N^2\boldsymbol{\Sigma}_m^{-2}). \quad (6.2.9)$$

We thus have derived a correction to the bias that is present in the original DMD matrix \mathbf{A}_m due to the effect of sensor noise. We note that this approximation relies on an accurate knowledge of the noise covariance matrix. There are numerous means to estimate noise properties from data, see Pyatykh et al. [118], for example. The approximations used in deriving this expression also rely on the magnitude of the noise being smaller than that of the true data within each non-truncated POD mode. We now state explicitly the algorithm by which we can correct for the effect of sensor noise in the DMD algorithm, which we refer to as noise-corrected DMD, or ncDMD:

Algorithm 3 (Noise-corrected DMD (ncDMD)).

1. Compute $\tilde{\mathbf{A}}_m$ from the measured data as per steps 1–3 of Algorithm 1
2. Compute the approximation of $\tilde{\mathbf{A}}$ from Equation (6.2.8)
3. Compute the DMD eigenvalues and modes via steps 4–5 of Algorithm 1, using the bias-free estimate of $\tilde{\mathbf{A}}$.

As was also noted in Section 6.2.1, we could have performed all of the above analysis without first projecting onto the space of POD coefficients, which gives us the following as analogous to Equations 6.2.8 and (6.2.9) respectively, subject to the appropriate inverses existing:

$$\mathbf{A} \approx \mathbf{A}_m(\mathbf{I} - m\sigma_N^2(\mathbf{Y}_m\mathbf{Y}_m^*)^{-1})^{-1} \approx \mathbf{A}_m(\mathbf{I} + m\sigma_N^2(\mathbf{Y}_m\mathbf{Y}_m^*)^{-1}). \quad (6.2.10)$$

While this approach might be computationally prohibitive for many applications of DMD (since it requires inverting large $n \times n$ matrices), it could in theory be more accurate, since it does not rely on any assumption that the POD modes for the measured and true data are sufficiently close to each other. Note again that $\mathbf{Y}\mathbf{Y}^*$ can only be invertible if $m > n$, as otherwise it cannot be full rank.

6.2.3 Forward-backward DMD

If we were to swap the data in \mathbf{Y} and $\mathbf{Y}^\#$, then (for suitably well behaved data) we should expect to identify the inverse dynamical system, with state propagation matrix \mathbf{B}_m (or $\tilde{\mathbf{B}}_m$ in POD space), which approximates the true dynamics \mathbf{B} (and $\tilde{\mathbf{B}}$). Note that it is not guaranteed that the dynamics of the original system are invertible, but this assumption should not be too restrictive for the majority of physical systems under consideration (particularly

after projection onto an appropriate POD subspace). It is argued in Section 6.6 that sensor noise has the effect of shifting the computed DMD eigenvalues to appear to be more stable than they actually are (i.e., moving them further inside the unit circle). Since our analysis was independent of the nature of the data, we should expect the same effect to be present for the computation of the inverse system. However, if $\tilde{\mathbf{B}}$ is invertible, then we should have $\tilde{\mathbf{B}} = \tilde{\mathbf{A}}^{-1}$, meaning that we should be able to compute an estimate of the forward-time propagation matrix using backward-time DMD, via $\tilde{\mathbf{A}}_m^{back} = \tilde{\mathbf{B}}_m^{-1}$. However, given that the eigenvalues of $\tilde{\mathbf{B}}_m$ should have their growth rates underestimated, those of the eigenvalues of $\tilde{\mathbf{A}}_m^{back}$ will be overestimated. Specifically, from consideration of Equation (6.2.4), we have

$$\mathbb{E}(\tilde{\mathbf{B}}_m) \approx \tilde{\mathbf{B}} \left(\mathbf{I} - \mathbb{E}(\tilde{\mathbf{N}}\tilde{\mathbf{N}}^*)(\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^*)^{-1} \right),$$

and so

$$\tilde{\mathbf{A}}_m^{back} \approx \left(\mathbf{I} - \mathbb{E}(\tilde{\mathbf{N}}\tilde{\mathbf{N}}^*)(\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^*)^{-1} \right)^{-1} \tilde{\mathbf{A}}, \quad (6.2.11)$$

where we are using the fact that the noise and POD energy components are the same for forward- and backward-DMD. We can then combine estimates of the dynamics from forward- and backward-time DMD to obtain

$$\tilde{\mathbf{A}}_m \tilde{\mathbf{A}}_m^{back} = \tilde{\mathbf{A}} \left(\mathbf{I} - \mathbb{E}(\tilde{\mathbf{N}}\tilde{\mathbf{N}}^*)(\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^*)^{-1} \right) \left(\mathbf{I} - \mathbb{E}(\tilde{\mathbf{N}}\tilde{\mathbf{N}}^*)(\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^*)^{-1} \right)^{-1} \tilde{\mathbf{A}} = \tilde{\mathbf{A}}^2. \quad (6.2.12)$$

We thus have the estimate

$$\tilde{\mathbf{A}} \approx (\tilde{\mathbf{A}}_m \tilde{\mathbf{A}}_m^{back})^{1/2}. \quad (6.2.13)$$

Note that this square root will in general be non-unique, and thus determining which root is the relevant solution could be nontrivial. One reasonable method, if there is any ambiguity, is to take the square root which is closest to $\tilde{\mathbf{A}}_m$ (or $\tilde{\mathbf{A}}_m^{back}$). See Golub and Van Loan [59] for a more detailed discussion of the computation of matrix square roots. As an aside, note that if we assume that we know the equivalent continuous time matrices $\tilde{\mathbf{A}}_m^c = \log(\tilde{\mathbf{A}}_m)/\Delta t$ and $\tilde{\mathbf{A}}_m^{c,back} = \log(\tilde{\mathbf{A}}_m^{back})/\Delta t$, then the equivalent of Equation (6.2.13) is

$$\tilde{\mathbf{A}}^c \approx \frac{1}{2}(\tilde{\mathbf{A}}_m^c + \tilde{\mathbf{A}}_m^{c,back}).$$

We are now in a position to formalize this algorithm, which we refer to as *forward-backward* DMD or fbDMD.

Algorithm 4 (forward-backward DMD (fbDMD)).

1. Compute $\tilde{\mathbf{A}}_m$ from the measured data as per steps 1–3 of Algorithm 1
2. Compute $\tilde{\mathbf{B}}_m$ from the measured data as per steps 1–3 of Algorithm 1, where \mathbf{Y} and $\mathbf{Y}^\#$ are interchanged
3. Compute the approximation of $\tilde{\mathbf{A}}$ from Equation (6.2.13)
4. Compute the DMD eigenvalues and modes via steps 4–5 of Algorithm 1, using the improved estimate of $\tilde{\mathbf{A}}$ from step 4.

Note that in the case where most data snapshots are in both \mathbf{Y} and $\mathbf{Y}^\#$ (e.g., for a sequential time series of data) we can reduce the computational cost of steps 1–2 in Algorithm 4 by first taking the SVD of the entire data set, and then working in the space of the resulting POD modes.

6.2.4 Total least-squares DMD

For the case where the number of snapshots, m , is greater than the size of each snapshot, n , the DMD matrix \mathbf{A} can be interpreted as the least-squares solution to the overdetermined system $\mathbf{A}\mathbf{Y} = \mathbf{Y}^\#$. When $n > m$, then the solution for the now underdetermined system is the minimum Frobenius norm solution to $\mathbf{A}\mathbf{Y} = \mathbf{Y}^\#$. In both cases, this solution is $\mathbf{A} = \mathbf{Y}^\# \mathbf{Y}^+$. Note that it is possible to turn an under-determined system into an over-determined system by truncating the number of POD modes used to less than m (truncating to precisely m results in a unique solution when the data is full column rank, with no loss of data). A least-squares solution of this form minimizes the error in $\mathbf{Y}^\#$, but implicitly assumes that there is no error in \mathbf{Y} . This can explain why the bias in DMD (Equation (6.2.4)) is dependent on $\tilde{\mathbf{N}}$, but not $\tilde{\mathbf{N}}^\#$. That is, in the least-squares case DMD can be viewed as finding

$$\mathbf{A} : \mathbf{Y}^\# + \mathbf{E}^\# = \mathbf{A}\mathbf{Y}, \quad \text{minimizing } \|\mathbf{E}^\#\|_F,$$

where $\|\cdot\|_F$ denotes the Frobenius norm of a matrix. When doing backwards-time DMD in Section 6.2.3, we conversely assume that $\mathbf{Y}^\#$ is known exactly and minimize the error in \mathbf{Y} . That is, assuming the identified dynamics are invertible, we find

$$\mathbf{A} : \mathbf{Y}^\# = \mathbf{A}(\mathbf{Y} + \mathbf{E}), \quad \text{minimizing } \|\mathbf{E}\|_F.$$

For this reason, combining forward- and backward-time DMD takes into account the error in both \mathbf{Y} and $\mathbf{Y}^\#$. A more direct means of doing this is to use a single algorithm that finds a least-squares solution for the error in both \mathbf{Y} and $\mathbf{Y}^\#$. It is possible to adapt standard TLS algorithms [59] to a DMD setting, which we perform here. We seek

$$\mathbf{A} : (\mathbf{Y}^\# + \mathbf{E}^\#) = \mathbf{A}(\mathbf{Y} + \mathbf{E}), \quad \text{minimizing } \|\mathbf{E}\|_F, \quad \text{where } \mathbf{E} = \begin{bmatrix} \mathbf{E} \\ \mathbf{E}^\# \end{bmatrix}.$$

The expressions $\mathbf{Y}^\# + \mathbf{E}^\#$ and $\mathbf{Y} + \mathbf{E}$ can be interpreted as $\mathbf{Y}_m^\# - \mathbf{N}^\#$ and $\mathbf{Y}_m - \mathbf{N}$. To solve for this, we can rearrange the equation to obtain

$$[\mathbf{A} \quad -\mathbf{I}] \begin{bmatrix} \mathbf{Y} + \mathbf{E} \\ \mathbf{Y}^\# + \mathbf{E}^\# \end{bmatrix} = 0. \quad (6.2.14)$$

We would now like to assume that $2n < m$. This might not be the case, particularly for high-dimensional fluids data. To get around this, and improve computational tractability, we may project Equation (6.2.14) onto a POD subspace of dimension $r < m/2$, to obtain

$$[\tilde{\mathbf{A}} \quad -\mathbf{I}] \begin{bmatrix} \tilde{\mathbf{Y}} + \tilde{\mathbf{E}} \\ \tilde{\mathbf{Y}}^\# + \tilde{\mathbf{E}}^\# \end{bmatrix} = 0. \quad (6.2.15)$$

This POD projection step is in contrast to the TLS DMD formulation in Hemati et al. [69], where a projection is performed onto a basis determined from an augmented snapshot matrix $Z = \begin{bmatrix} \mathbf{Y} \\ \mathbf{Y}^\# \end{bmatrix}$. We find that the present formulation yields more accurate eigenvalues in a number of examples. Note that the nullspace of $\begin{bmatrix} \tilde{\mathbf{A}} & -\mathbf{I} \end{bmatrix}$ is r -dimensional, meaning that the $2r$ by m matrix $\begin{bmatrix} \tilde{\mathbf{Y}} + \tilde{\mathbf{E}} \\ \tilde{\mathbf{Y}}^\# + \tilde{\mathbf{E}}^\# \end{bmatrix}$ can have rank at most r .

Let the full SVD of $\begin{bmatrix} \tilde{\mathbf{Y}} \\ \tilde{\mathbf{Y}}^\# \end{bmatrix}$ be given by $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$. If the data is noisy, we should expect that all $2r$ diagonal entries of $\mathbf{\Sigma}$ are nonzero. By the Eckart-Young theorem [50], the nearest (in the sense of Frobenius norm) rank r matrix will be given by

$$\begin{bmatrix} \tilde{\mathbf{Y}} + \tilde{\mathbf{E}} \\ \tilde{\mathbf{Y}}^\# + \tilde{\mathbf{E}}^\# \end{bmatrix} = \mathbf{U}\mathbf{\Sigma}_{1:r}\mathbf{V}^*,$$

where $\mathbf{\Sigma}_{1:r}$ contains the leading r singular values of $\mathbf{\Sigma}$, with the rest replaced by zeros. We then have that

$$\begin{bmatrix} \tilde{\mathbf{Y}} + \tilde{\mathbf{E}} \\ \tilde{\mathbf{Y}}^\# + \tilde{\mathbf{E}}^\# \end{bmatrix} = \mathbf{U}\mathbf{\Sigma}_{1:r}\mathbf{V}^* = \begin{bmatrix} \mathbf{U}_{11} & \mathbf{U}_{12} \\ \mathbf{U}_{21} & \mathbf{U}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{\Sigma}_1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{V}_1^* \\ \mathbf{V}_2^* \end{bmatrix} = \begin{bmatrix} \mathbf{U}_{11}\mathbf{\Sigma}_1\mathbf{V}_1^* \\ \mathbf{U}_{21}\mathbf{\Sigma}_1\mathbf{V}_1^* \end{bmatrix},$$

where \mathbf{U}_{ij} are r by r matrices, and \mathbf{V}_1 is the first r columns of \mathbf{V} . Rearranging this equation, we obtain the total least-squares estimate for $\tilde{\mathbf{A}}$:

$$\tilde{\mathbf{A}} = \mathbf{U}_{21}\mathbf{U}_{11}^{-1}. \quad (6.2.16)$$

Note that this derivation requires that \mathbf{U}_{11} be invertible. While the derivation includes the full SVD of the augmented data, Equation (6.2.16) indicates that we only need the first r columns of \mathbf{U} , meaning that only a reduced SVD is required. Algorithm 5 summarizes this total least-squares approach to DMD, which we refer to as *total least-squares* DMD, or *tlsDMD*.

Algorithm 5 (total least-squares DMD (tlsDMD)).

1. Collect data \mathbf{Y} and $\mathbf{Y}^\#$, and project onto $r < m/2$ POD modes to obtain $\tilde{\mathbf{Y}}$ and $\tilde{\mathbf{Y}}^\#$.
2. Take the SVD of $\begin{bmatrix} \tilde{\mathbf{Y}} \\ \tilde{\mathbf{Y}}^\# \end{bmatrix}$, letting $\begin{bmatrix} \tilde{\mathbf{Y}} \\ \tilde{\mathbf{Y}}^\# \end{bmatrix} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$.
3. Partition the $2r$ by $2r$ matrix \mathbf{U} into r by r sub-matrices, letting $\mathbf{U} = \begin{bmatrix} \mathbf{U}_{11} & \mathbf{U}_{12} \\ \mathbf{U}_{21} & \mathbf{U}_{22} \end{bmatrix}$ (note that only the first r columns need to be computed).
4. Compute the total least-squares DMD matrix $\tilde{\mathbf{A}}$, using Equation (6.2.16).
5. Compute the DMD eigenvalues and modes using steps 4–5 of Algorithm 1.

An alternative and more focused exposition of `tlsDMD` is given in Hemati et al. [69]. We note that Algorithm 5 is not identical to that presented in this work (due to the lack of pre-truncation of POD modes), however we find that Algorithm 5 gives marginally better results in terms of the accuracy of identified eigenvalues.

6.3 Results with synthetic data

In this section we will test our proposed modifications to DMD on a number of examples. Using known dynamics with the addition of random noise will allow us to examine the performance of these proposed modifications (Algorithms 3–5) in comparison to regular DMD (Algorithm 1). We begin by considering a simple 2-dimensional linear system in Section 6.3.1. In Section 6.3.2, we consider the same system with an expanded set of observables, which tests the important case of high-dimensional data that is described by low-dimensional dynamics. Section 6.3.3 compares the performance of the proposed modifications of DMD to other DMD variants in recent literature, while Section 6.3.4 considers the problem of identifying dynamics that are quickly decaying and obscured by dominant modes and noise, a case where DMD-like algorithms could be of most use. Finally, in Section 6.3.5 we analyze how the proposed DMD modifications treat process noise.

6.3.1 Example: A periodic linear system

We consider first a simple two-dimensional linear system, with dynamics given by

$$\dot{\mathbf{x}} = \begin{bmatrix} 1 & -2 \\ 1 & -1 \end{bmatrix} \mathbf{x}. \quad (6.3.1)$$

This system has (continuous-time) eigenvalues $\lambda_{c1,2} = \pm i$, so gives purely periodic dynamics, with no growth or decay. We discretize with a timestep $\Delta t = 0.1$, so the discrete-time eigenvalues are then $\lambda_{1,2} = e^{\pm \Delta t i}$. We use 100 timesteps of data (i.e., $m = 99$), corrupted with Gaussian white noise of variance $\sigma_N^2 = 0.01$. The identified continuous-time eigenvalues from both regular DMD (Algorithm 1), and the direct noise-correction (Algorithm 3) are shown in Figure 6.2(a), for 1000 different trials from the initial condition $\mathbf{x}_0 = [1 \ 0.1]^T$. We assume that the correction term is given by $m\sigma_N^2 I_n$, and observe that this corrects almost perfectly for the bias in the DMD algorithm in terms of identifying eigenvalues. Also shown in Figure 6.2(a) are ellipses representing the 95% confidence region, with the major and minor axes of the ellipse aligned with the principal component directions of the eigenvalue data. For clarity, in the presentation of subsequent results, we will omit individual data points and show only such ellipses. In Figure 6.2(b) we show the mean and 95% confidence ellipses for Algorithms 4 and 5. As with `ncDMD`, both `fbDMD` (Algorithm 4) and `tlsDMD` (Algorithm 5) accurately correct for the bias in the mean of the identified eigenvalue. Further to this, `fbDMD` and `tlsDMD` also both reduce the area of the 95% confidence ellipse, which indicates that they are more likely to attain a closer approximation to the correct eigenvalue on any given trial.

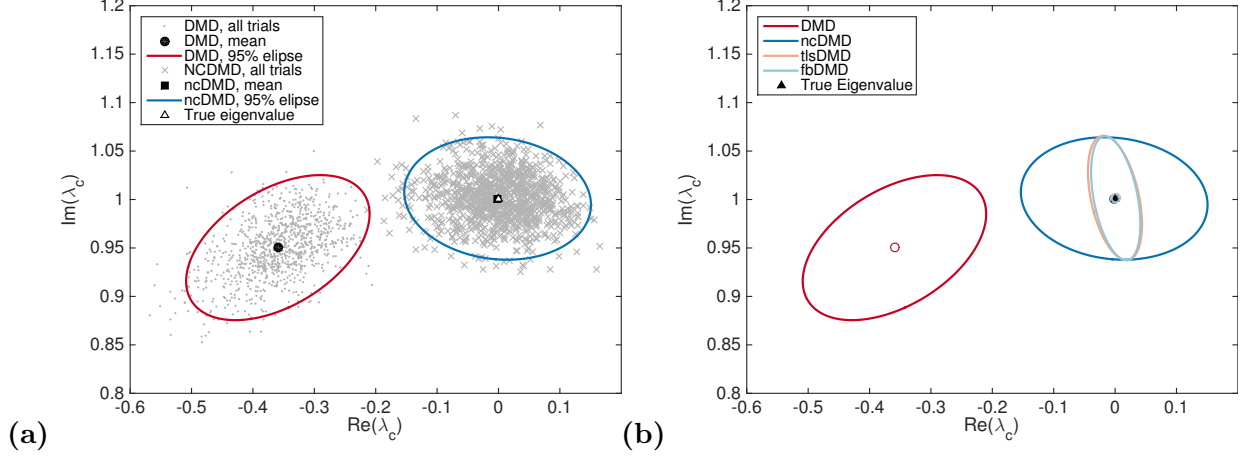


Figure 6.2: **(a)** Eigenvalues (in continuous-time) identified from regular DMD (Algorithm 1, dots) and noise-corrected DMD (Algorithm 3, crosses) from 100 snapshots of data from Equation (6.3.1), with $\Delta t = 0.1$ and $\sigma_N^2 = 0.01$. Only one of the complex conjugate pair of eigenvalues is shown. The mean and 95% confidence ellipse of 1000 trials are given for each data set. **(b)** shows the mean and 95% confidence ellipse for the same data set for each of the algorithms

Focusing back on comparing Algorithms 1 and 3, we show results for a variety of values of m and σ_N^2 in Figure 6.3. In Figure 6.3(a), rather than looking at the error in the eigenvalues, we instead consider the Frobenius norm of the difference between the true and identified propagation matrices, $\|\mathbf{A}_{\text{true}} - \mathbf{A}_{\text{pred}}\|_F$. For very small noise, the correction makes little difference, since the random error is larger than the bias error. For larger values of noise, we observe that the error saturates when using standard DMD, which is due to the presence of the bias term identified in Section 6.2.1, which has a size independent of the number of samples, m . We note that the magnitude of this bias term is proportional to σ_N^2 , as predicted by Equation (6.2.8). Evidence of this error saturation phenomena can also be seen in past studies of the effect of sensor noise on DMD [48, 176, 110]. After this bias term is corrected for, we see that the error decays proportional to $m^{-1/2}$ for all values of noise, as predicted from the analysis in the appendix, Section 6.6. The more rapid decay in error with m for small numbers of samples seems to arise from the fact that the data has not yet completed one full period of oscillation. Figure 6.3 shows the corrections to DMD made using both the sampled ($\mathbf{N}\mathbf{N}^*$) and theoretical ($m\sigma_N^2\mathbf{I}$) covariance matrices. Normally the sample noise covariance would not be known, and so we demonstrate here that the theoretical covariance achieves almost the same decrease in error. Figure 6.3(b) shows that the ncDMD error curves collapse when the error is normalized by the standard deviation of noise, σ_N (note that we could also multiply the error by the SNR to get the same scaling).

Figure 6.4 shows the performance of Algorithms 4 and 5 on the same data as Figure 6.3. Again, we find that both of these algorithms can prevent the error saturation present in standard DMD, and indeed can perform noticeably better than Algorithm 3 for larger noise

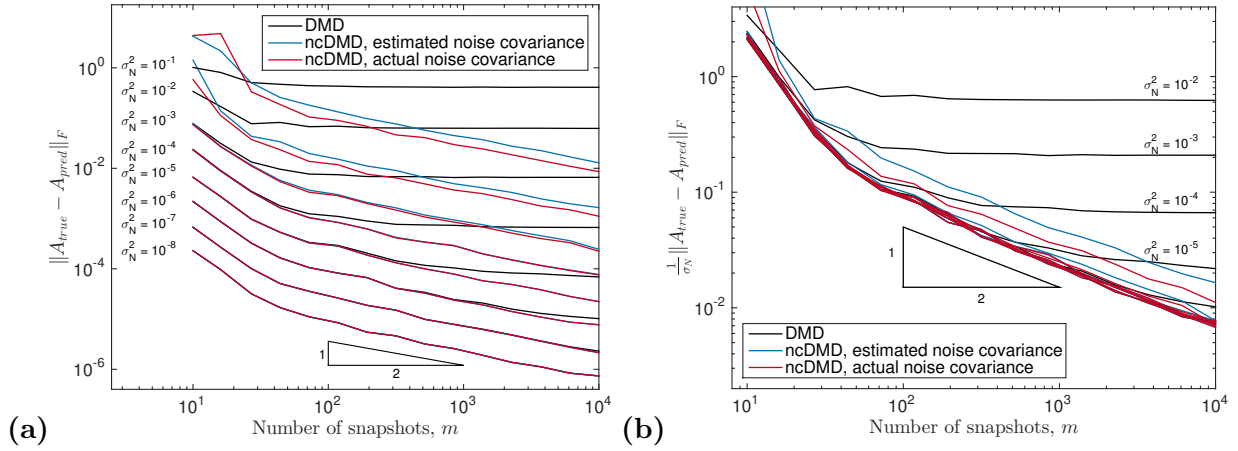


Figure 6.3: Error in the estimated propagation matrix $\tilde{\mathbf{A}}$ identified from performing DMD and ncDMD on noise-corrupted data generated from Equation (6.3.1), for different values of m and σ_N^2 . In (a) the error is given as $\|\mathbf{A}_{\text{true}} - \mathbf{A}_{\text{pred}}\|_F$, while in (b) this quantity is normalized by the standard deviation of the noise, σ_N . In both cases, the error is averaged over 100 trials for each m and σ_N^2 . Note that for clarity, (b) excludes the two largest noise levels shown in (a)




Figure 6.4: Error in the estimated propagation matrix $\tilde{\mathbf{A}}$ identified from performing DMD, ncDMD, fbDMD, and tIsDMD on noise-corrupted data generated from Equation (6.3.1), for different values of m and σ_N^2 . The error is given as $\|\mathbf{A}_{\text{true}} - \mathbf{A}_{\text{pred}}\|_F$, and is averaged over 100 trials for each m and σ_N^2

levels. Algorithms 3–5 all appear to exhibit the same asymptotic behavior as the number of snapshots, m , increases, with the error decreasing proportional to $m^{-1/2}$.

A common means to mitigate the effect of noisy data is to collect multiple time-series of data, and process this in such a way to improve the results over just using one data set. One can ask the question if it is better to concatenate the snapshots of data from each time series and apply DMD to this collection, or to apply DMD to phase-averaged data. Our results suggest that the latter option is preferable if using standard DMD, since adding additional pairs of snapshots will not decrease the error beyond a certain level, due to this bias saturation at large m . If we are using ncDMD, fbDMD, or tlsDMD, however, then we should get the same scalings regardless of which option is chosen, since in both cases the error should be proportional to $p^{-1/2}$, where p is the number of trials of data collected.

6.3.2 A periodic linear system with a high-dimensional state of observables

This example considered in Section 6.3.1 has $m \gg n$, which is atypical of many fluids systems for which DMD is used. To consider the case where the size of the state n is larger than the number of snapshots m , we expand the state of our system to include time-shifts of the data. In this sense, we have new observables given by

$$\mathbf{z}_k = \begin{bmatrix} \mathbf{x}_k \\ \mathbf{x}_{k-1} \\ \vdots \\ \mathbf{x}_{k-q} \end{bmatrix}, \quad (6.3.2)$$

with the size of the state $n = 2(q+1)$. This periodic system can equivalently be viewed as a traveling wave, which is now observed over a larger spatial domain. Similar data (but with a non-zero growth rate) was considered in Duke et al. [48] and Wynn et al. [176]. Since the dynamics are still only two-dimensional despite the higher dimensional state, we use only the first two POD modes of the data to identify a 2×2 propagation matrix $\tilde{\mathbf{A}}$. The next section will examine alternative means of performing this dimensionality reduction.

Figure 6.5 shows the statistical results (in terms of DMD eigenvalues) of performing variants of DMD on such data, using $m = 50$ and a range of snapshot sizes, n . We find that a bias exists for regular DMD, but the magnitude of this bias decreases as the size of each snapshot increases (note that the scale between subplots changes, though the aspect ratio remains the same). We find that Algorithms 3–5 all outperform regular DMD in terms of giving mean (expected) eigenvalues that are closer to the true value. For small state sizes, Algorithms 4 and 5 also give a smaller confidence ellipse, though this is not observed for larger state sizes. As the size of the state increases, the bias component of the error of DMD (evidenced by the difference between the true and mean identified eigenvalue) becomes smaller relative to the random component of the error (indicated by the size of the confidence ellipse). This means that the modifications to DMD presented in Algorithms 3–5 give the largest improvement when the size of the state is small, due to the fact that in this regime the bias component of error is larger than the random component. Note that these conclusions

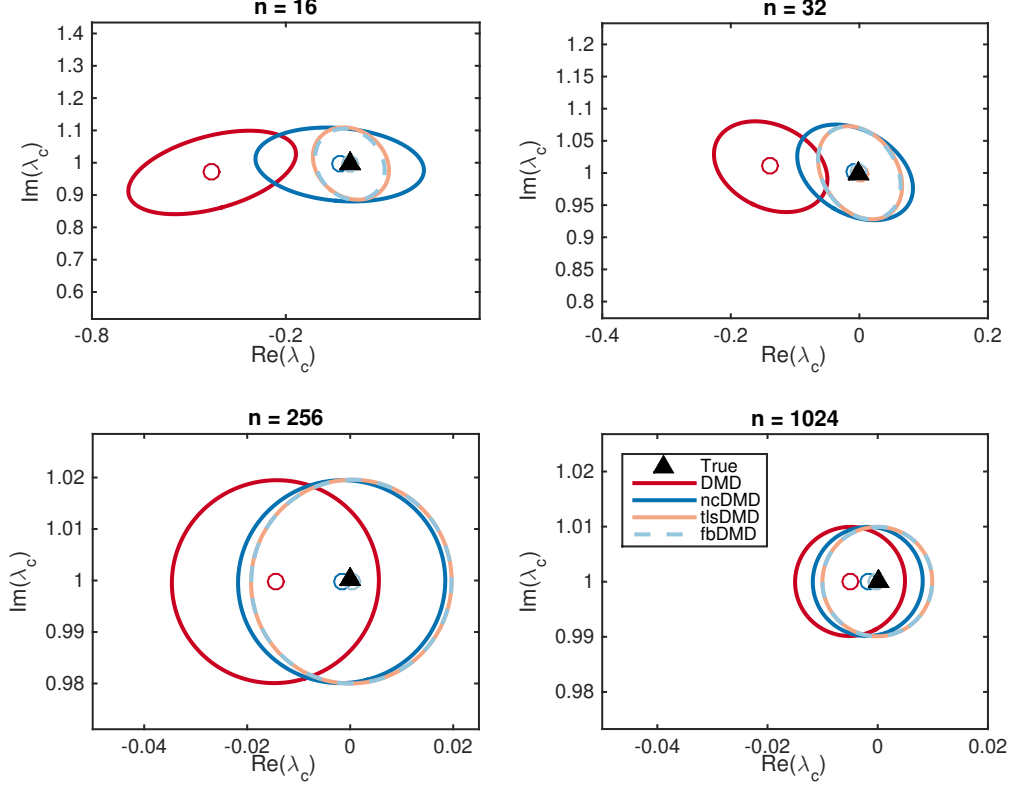


Figure 6.5: Mean and 95% confidence ellipses of continuous-time eigenvalues identified by applying regular DMD (Algorithm 1), noise-corrected DMD (Algorithm 3), forward-backward DMD (Algorithm 4) and total least-squares DMD (Algorithm 5) on 1000 trials of noisy data generated by Equation (6.3.1), and observed as in Equation (6.3.2). Here the number of snapshots m is fixed to be 50, $\Delta t = 0.1$, and $\sigma_N^2 = 0.1$. Only one of the complex conjugate pair of eigenvalues is shown

may be predicted from the scaling laws given in Equations (6.6.2) and (6.6.5). Moreover, one can verify that as the size of the state (n) increases, the size of the ellipses decrease proportional to $n^{-1/2}$.

6.3.3 Comparison to other modified DMD algorithms

Without any modification, applying DMD on noisy data will give $\min(m, n)$ eigenvalue-mode pairs, many of which may be mostly or entirely due to noise, particularly if the underlying dynamics are low-dimensional. For this reason, a number of modifications of DMD that aim to identify a small number of dynamically important modes have been developed. The most simple means of reducing the dimension of the data is to simply project onto a reduced number of POD modes, which is explicitly mentioned as an optional step in Algorithm 1. This projection step was also used within Algorithms 3–5 in Section 6.3.2. A number of alternative means to obtain a small number of dynamic modes from DMD-type algorithms have been proposed, as briefly mentioned in Section 1.2. These variants all start with the observation

that standard DMD can be formulated within an optimization framework, in the sense that it identifies a least-squares or minimum-norm propagation matrix for a given data set. Chen et al. [28] proposes a modification termed *optimized* DMD that seeks to find optimal low-rank dynamics that best matches a sequential time-series of data. While the fact that this method optimizes over the entire time-sequence of data rather than just pairwise snapshots should increase its robustness to noise, the non-convexity of the optimization potentially limits its utility. Optimal Mode Decomposition (OMD, [62, 176]) finds an optimal low-dimensional subspace on which the identified dynamics reside, rather than assuming that this subspace is simply the most energetic POD modes. This approach was shown to give an improvement on the DMD eigenvalues obtained for noisy data in Wynn et al. [176]. Sparsity-promoting DMD (spDMD, [78]) adds an l_1 regularization term that penalizes the number of DMD modes with non-zero coefficients in the approximation of the time-series of data.

This section will compare OMD and spDMD with the algorithms presented in the present work. Of the algorithms presented here, we will focus on fbDMD (Algorithm 4), which was found to perform equally well as tlsDMD, and better than ncDMD, in Sections 6.3.1 and 6.3.2. Figure 6.6 shows identified eigenvalue statistics (mean and confidence ellipses) for each of these algorithms, using the same data as that for Figure 6.5. We observe that OMD gives a more accurate mean eigenvalue than DMD, and a confidence ellipse of approximately the same size. spDMD gives a mean identified eigenvalue that is closer again to the mean, although the variance in the eigenvalues identified for each trial is larger. We note that spDMD occasionally produced erroneous results, which were excluded as outliers from the statistical analysis. This highlights an important advantage to the modifications to DMD presented here: the algorithms are given in closed form, and do not rely on an appropriate selection of parameters and tolerances that are most likely required for an optimization procedure. In all of the cases, fbDMD (and tlsDMD, which is not shown but barely distinguishable from fbDMD) gives the best estimate of the true eigenvalue.

While these results suggest that fbDMD/tlsDMD is more accurate than OMD and spDMD, we must remember that the results from one data set do not show the global superiority of any given algorithm. Indeed, one could most likely find data sets for which any given algorithm is superior (by some chosen metric) to others. We conclude this section by noting that it should be possible to combine the optimization procedures presented in Chen et al. [28], Wynn et al. [176], and Jovanović et al. [78] with the modifications to DMD presented here. Indeed, a simple means to do this might be to modify Algorithm 4 so that the results of applying a given algorithm forwards and backwards in time are geometrically averaged, as in Equation (6.2.13).

6.3.4 Identifying hidden dynamics

The systems considered in Sections 6.3.1 and 6.3.2 could be considered “easy” in the sense that the dominant dynamics are simple, and of consistently larger magnitude than the noise. Indeed, it is not difficult to qualitatively identify such dynamics by eye from simply looking at some visualization of the data. A more difficult case occurs when some of the dynamics are of low magnitude and/or are quickly decaying, and thus might quickly be lost among the noise in the measurements. A major benefit of data processing techniques such as DMD

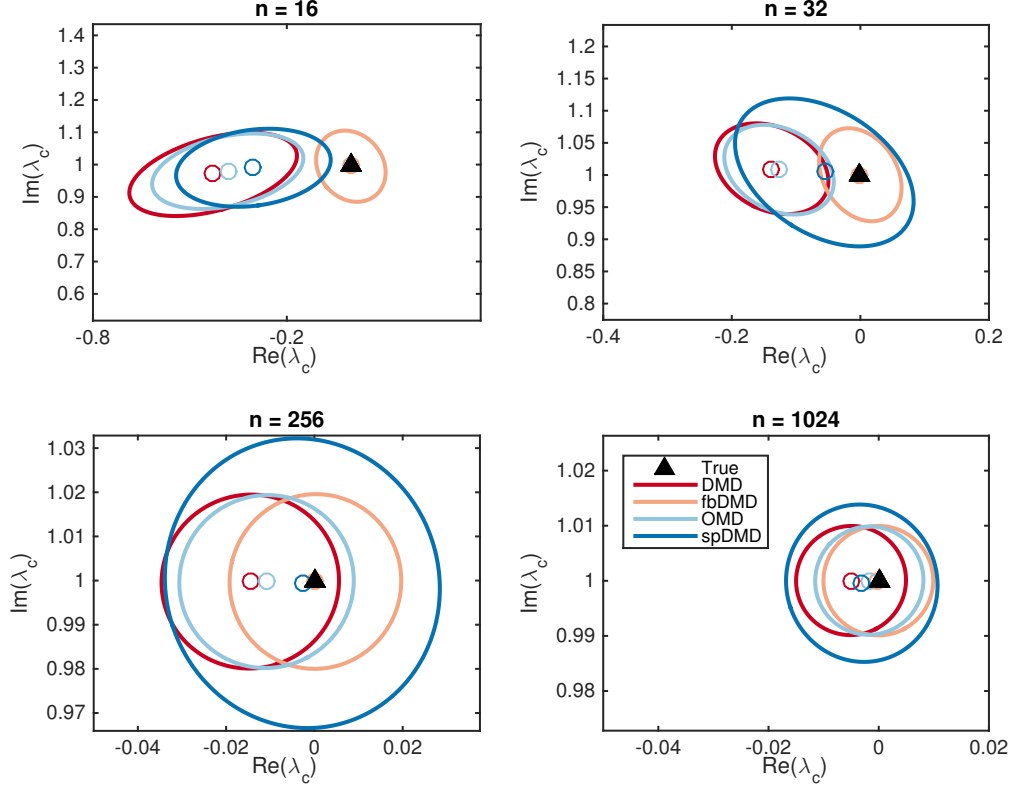


Figure 6.6: Mean and 95% confidence ellipses of continuous-time eigenvalues identified by applying regular DMD (Algorithm 1), forward-backward DMD (Algorithm 5), OMD and spDMD for noisy data generated from 1000 trials of data generated by Equation (6.3.1), and observed as in Equation (6.3.2). Here the number of snapshots m is fixed to be 50, $\Delta t = 0.1$, and $\sigma_N^2 = 0.1$. Only one of the complex conjugate pair of eigenvalues is shown




Figure 6.7: Visualization of data generated by Equation (6.3.3), with $k_1 = 1$, $\omega_1 = 1$, $\gamma_1 = 1$, $k_2 = 0.4$, $\omega_2 = 3.7$, $\gamma_2 = -0.2$, and $\sigma = 0.5$

is the ability to identify dynamics that might otherwise remain hidden. With this in mind, we now consider a superposition of two sinusoidal signals that are traveling across a spatial domain in time, with the amplitude of one mode growing, and the other decaying:

$$f(x, t) = \sin(k_1x - \omega_1t)e^{\gamma_1t} + \sin(k_2x - \omega_2t)e^{\gamma_2t} + n_\sigma(x, t), \quad (6.3.3)$$

where we set $k_1 = 1$, $\omega_1 = 1$, $\gamma_1 = 1$, $k_2 = 0.4$, $\omega_2 = 3.7$ and $\gamma_2 = -0.2$. We thus have the superposition of a growing, traveling wave, and a decaying signal that is quickly hidden by the unstable dynamics. The four continuous-time eigenvalues of this data are $\gamma_1 \pm \omega_1$ and $\gamma_2 \pm \omega_2$. This data is again similar to that considered in Wynn et al. [176] and Duke et al. [48], if we neglect the decaying dynamics. Figure 6.7 shows the data with white noise of standard deviation $\sigma = 0.5$. Figure 6.8 shows the performance of various DMD-type algorithms in identifying one of the dominant eigenvalues ($1 + i$) and one of the “hidden” eigenvalues ($-0.2 + 3.7i$). Mean eigenvalues and error ellipses are computed from 1000 different noise samples. Unsurprisingly, all methods are quite accurate at identifying the dominant eigenvalue, though the variants proposed in the present work show improvements in both the mean and scatter over the 1000 trials. In terms of the hidden eigenvalue, we observe that DMD (as well as OMD) estimates a decay rate that is almost double the true value. In contrast, all of ncDMD, fbDMD, and tlsDMD predict the eigenvalue accurately, with a reduction in the error of the mean eigenvalue between DMD and fbDMD (for example) of 88%. In addition, we note that the scatter in the identified hidden eigenvalue across the trials is smaller for fbDMD and tlsDMD (as indicated by smaller confidence ellipses).




Figure 6.8: Mean and 95% confidence ellipses of continuous-time eigenvalues identified by applying regular DMD (Algorithm 1), OMD, noise-corrected DMD (Algorithm 3), forward-backward DMD (Algorithm 4) and total least-squares DMD (Algorithm 5) to 1000 trials of noisy data generated by Equation (6.3.3), with $k_1 = 1$, $\omega_1 = 1$, $\gamma_1 = 1$, $k_2 = 0.4$, $\omega_2 = 3.7$, $\gamma_2 = -0.2$, and $\sigma = 0.5$

6.3.5 Differentiating between process and sensor noise

This section will primarily address the issue of comparing and distinguishing between the effects of process and sensor noise. We consider the Stuart-Landau equation, which has been used as a model for the transient and periodic dynamics of flow past a cylinder in the vortex shedding regime [106, 8]. In discrete time, we can express this system in polar coordinates by

$$\begin{aligned} r_{k+1} &= r_k + dt(\mu r_k - r_k^3 + n_r), \\ \theta_{k+1} &= \theta_k + dt(\gamma - \beta r_k^2 + \frac{n_\theta}{r_k}), \end{aligned} \quad (6.3.4)$$

where we have included process noise terms n_r and n_θ , which are assumed to be independent in time, and sampled from separate zero-mean Gaussian distributions with variance σ_P^2 . We take as our data snapshots of the form

$$\mathbf{x}_k = [e^{-Ji\theta_k} \ e^{(-J+1)i\theta_k} \ \dots \ e^{Ji\theta_k}]^T, \quad (6.3.5)$$

for some integer J . We may add sensor noise to this data as in previous sections. For $\mu > 0$, Equation (6.3.4) contains a stable limit cycle at $r = \sqrt{\mu}$, with period $2\pi/(\gamma - \beta\mu)$. Starting on the limit cycle, we consider data with process noise, sensor noise, neither, and both. Without any noise, the eigenvalues identified from this data will lie upon the imaginary axis, at locations given by $\lambda_c = ij(\gamma - \mu\beta)$. Process noise acts to perturb the system from its limit cycle, which ultimately leads to phase diffusion, and a “bending” of the eigenvalues such that they instead lie on a parabola. The behavior of this system with process noise is

described more extensively in Bagheri [9]. Figure 6.9 shows the results of applying variants of DMD on data generated by Equation (6.3.4) with $\mu = 1$, $\gamma = 1$, $\beta = 0$, and $dt = 0.01$, with data collected using Equation (6.3.5) with $J = 10$. Applying DMD on noise-free data gives eigenvalues along the imaginary axis, while data from the system with process noise gives a parabola of eigenvalues, as expected. For data collected using Equation (6.3.5), each data channel will be orthogonal in time, and will contain the same energy. As a result, sensor noise will act to shift all identified eigenvalues into the left half plane by the same amount, as observed in Figure 6.9(a). Figure 6.9(b) shows that applying ncDMD accurately corrects for this shift, for the system with and without process noise. This shows that it is possible to distinguish between the effects of these two forms of noise, given only an estimate of the magnitude of the sensor noise. That is, we are able to eliminate the effects of the noise that is due to imperfections in our observations, while retaining the effects of actual disturbances to the system. Figure 6.9(c) shows that tlsDMD corrects for the effects of both process and sensor noise, which is desirable if one wishes to recover the dynamics of the noise-free system. The results for fbDMD are not shown, but were very similar to those for tlsDMD. The ability of tlsDMD and fbDMD on process noise is not surprising, since they treat \mathbf{Y} and $\mathbf{Y}^\#$ in a symmetric manner, and thus consider phase diffusion both forwards and backwards in time.

6.4 Results with numerical and experimental data

Having analyzed the performance of the various proposed modifications of DMD on synthetic data sets, we now turn our attention to data obtained from fluids simulations and experiments. We will focus on the canonical case of the unsteady wake of a circular cylinder exhibiting periodic vortex shedding. In Section 6.4.1 we present results from data obtained from a two-dimensional direct numerical simulation, while Section 6.4.2 considers data obtained from PIV measurements in a water channel.

6.4.1 Cylinder wake: simulation data

We use an immersed boundary projection method [149, 36], with a domain consisting of a series of nested grids, with the finest grid enclosing the body, and each successive grid twice as large as the previous. The finest grid consists of uniformly spaced points with grid spacing equal to $0.02D$ (where D is the cylinder diameter), extending $2D$ upstream and $4D$ downstream of the center of the cylinder, and spanning $4D$ in the direction normal to the flow. Each successively larger grid contains the same number of grid points, with twice the grid spacing as the previous grid. More details about the numerical method were given in Section 4.2. The Reynolds number $Re = \frac{U_\infty D}{\nu}$ was set to be 100, where ν is the kinematic velocity. This Reynolds number is above that for which the wake is stable (47 [117]), and below that for which three-dimensional instabilities emerge (approximately 194 [175]). At this Reynolds Number, the wake is hence unstable, and approaches a single periodic limit cycle characterized by a von Kàrmàn vortex street in the wake. The data to be analyzed was taken from 234 snapshots of the vorticity field, spaced $0.1 \frac{D}{U_\infty}$ time units apart. This

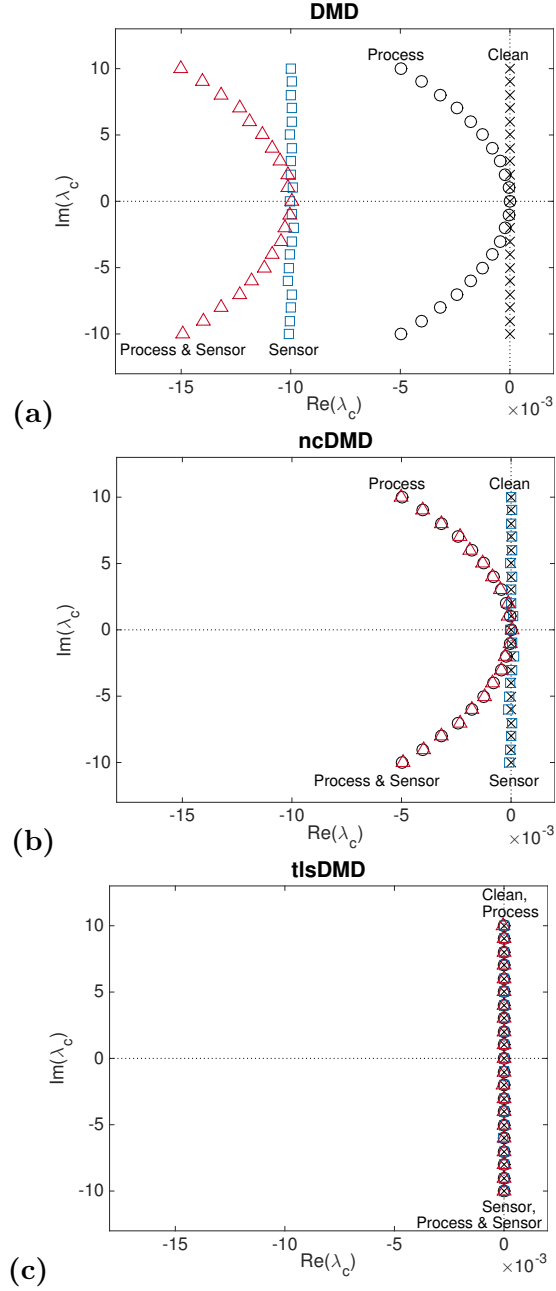


Figure 6.9: Eigenvalues identified using (a) DMD, (b) ncDMD, and (c) tIsDMD for the Stuart-Landau equation (Equation (6.3.4)), with 100,000 snapshots of data from Equation (6.3.5), with $r_0 = 1$, $\mu = 1$, $\gamma = 1$, $\beta = 0$, and $dt = 0.01$. Data with sensor noise, process noise, neither and both are considered, with noise levels for process and sensor noise being $\sigma_P^2 = 0.01$ and $\sigma_N^2 = 10^{-4}$ respectively. Note that in the absence of sensor noise, DMD and ncDMD are identical

corresponds to approximately 4 complete periods of vortex shedding. We truncate the data to only consider first 15 POD modes. These first 15 POD modes contain 99.99% of the total energy of the clean data, and 92.96% of the total energy of the data after the addition of Gaussian white noise with standard deviation $\sigma = 0.2 \frac{U_\infty}{D}$. Thus it is almost entirely noise that is truncated for the noisy data.

Figure 6.10 shows results from applying various variants of DMD to such data. Though not shown, the results of applying `tlsDMD` were visually indistinguishable as using `fbDMD`. Since we are artificially adding noise, we can compare the results using noisy data to those generated from the noise-free data. When applying regular DMD to noisy data, we observe significant errors in the growth rate associated with the highest-frequency eigenvalues (Figure 6.10(a)). For an oscillatory system such as this, the DMD eigenmodes are very similar to the POD modes, with a DMD mode corresponding to $\lambda_c \approx 0$ that is almost the mean flow, and the modes associated with conjugate pairs of DMD eigenvalues corresponding to pairs of POD modes with equal energy, see Chen et al. [28] for further discussion of this phenomenon. This means that the observed measured eigenvalues are in line with the analysis given in Sections 6.2.1 and 6.6, since the lower energy POD modes oscillate the most. We can see the effect of this error in Figure 6.10(b), which shows the prediction of a number of POD coefficients as evolved by the identified system, starting from the true initial condition. The dominant, low frequency POD modes are accurately predicted, but the higher “harmonics” are erroneously predicted to decay when using regular DMD. `ncDMD` improves the performance marginally, while `fbDMD` and `tlsDMD` both almost completely remove the erroneous decay of the high-frequency modes.

As well as considering eigenvalues, we also validate in Figure 6.11 that the modifications of DMD do not adversely affect the identified DMD modes. This is shown both visually in Figure 6.11(a), and quantitatively in Figure 6.11(b), where we give the inner product $\langle \phi_{i,\text{noisy}}, \phi_{i,\text{clean}} \rangle$ of the i^{th} modes identified from clean and noisy data, where we have pre-scaled the modes to be of unit norm. We enumerate the modes by the imaginary component of the associated eigenvalue, with mode 0 corresponding to the eigenvalue on the real axis. For modes that come in complex conjugate pairs, we arbitrarily consider those with positive imaginary component. We see that both `fbDMD` and `tlsDMD` marginally outperform regular DMD, in terms of identifying modes that are at least as close to those identified from noise-free data. The decrease in the inner product as the mode number increases is indicative of noise being more significant in higher-frequency modes, which contain less energy.

6.4.2 Cylinder wake: experimental data

We now turn our attention to data acquired from water channel experiment. An anodized aluminum cylinder of diameter $D = 9.5$ mm and length $L = 260$ mm was immersed in a recirculating, free-surface water channel with freestream velocity $U_\infty = 4.35$ cm/s. This gives a Reynolds number $Re = \frac{DU_\infty}{\nu} = 413$. Further details of the experimental setup and methodology are provided in Tu et al. [160]. We apply variants of DMD to 500 snapshots from a vorticity field of size 135×80 . Figure 6.12 shows the identified eigenvalues and the predicted POD coefficients from the models identified from DMD and `tlsDMD`. As in

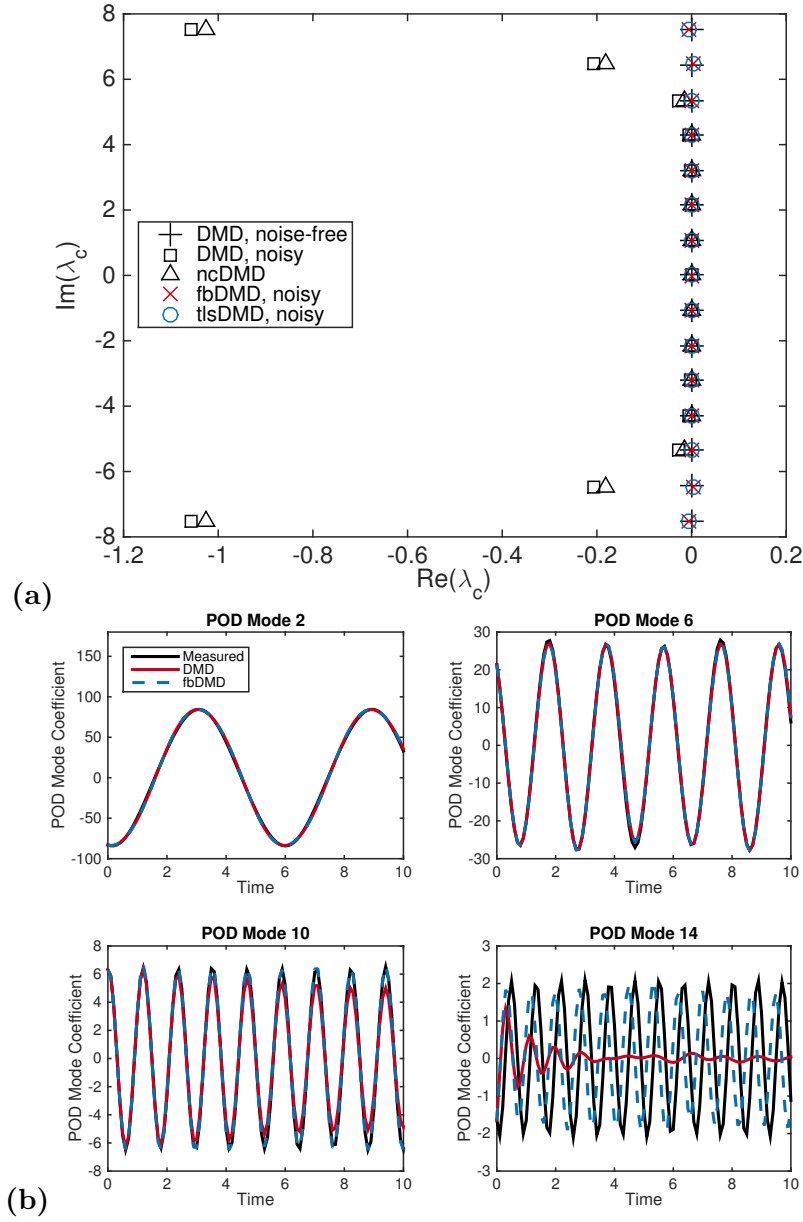


Figure 6.10: (a) Eigenvalues and (b) POD coefficients identified from applying DMD, ncDMD, fbDMD, and tIsDMD to DNS vorticity data from a cylinder wake at $Re = 100$. Noisy data was corrupted with Gaussian white noise with $\sigma = 0.2 \frac{U_\infty}{D}$.

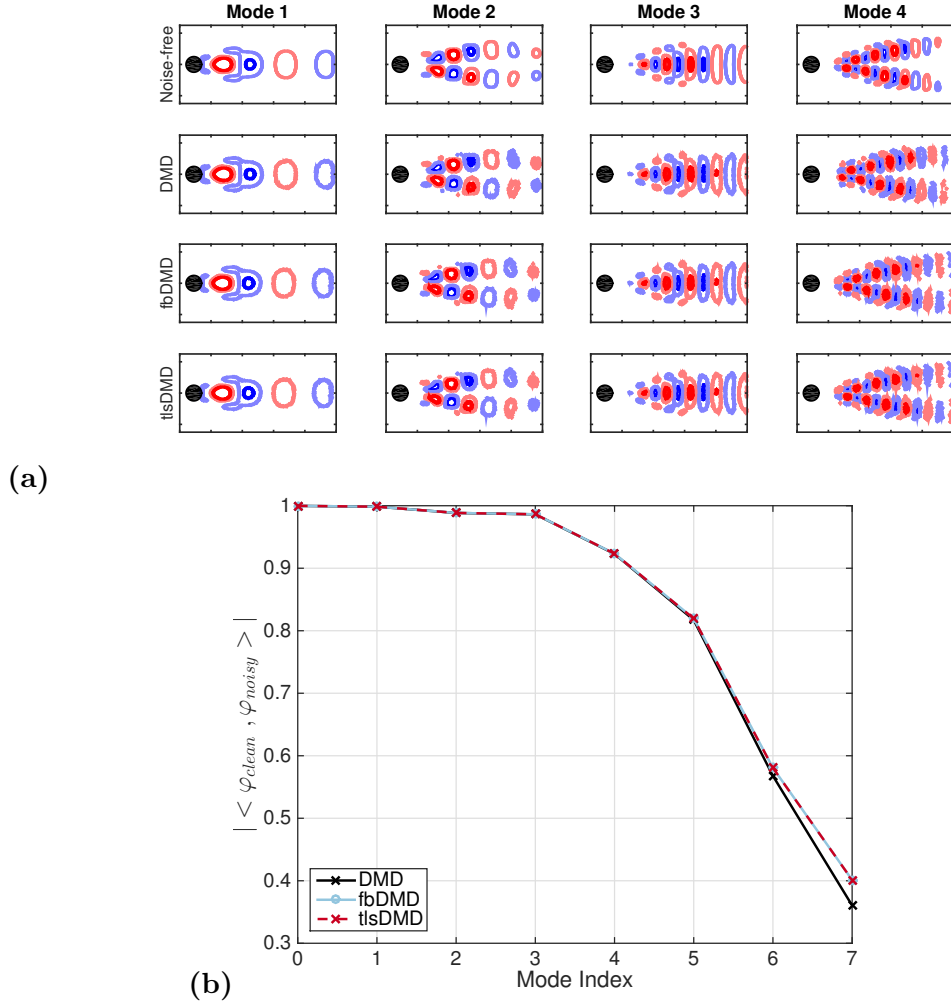


Figure 6.11: (a) A subset of the DMD modes (real components) computed from applying various variants of DMD to DNS data of flow around a cylinder. (b) Inner product between (normalized) clean modes, and modes obtained from noisy data (with $\sigma_N = 0.2 \frac{U_\infty}{D}$)

Section 6.4.1, we first project onto the 15 most energetic POD modes. Note that some eigenvalues (typically quickly decaying) are outside the range of the plot. As was the case with DNS data, we observe that DMD gives eigenvalues that are further into the left half plane than those identified using *tlsDMD*. This manifests in the erroneous prediction of decaying POD coefficients (Figure 6.12(b)), particularly for modes that are less energetic, and more rapidly oscillating. We thus conclude that more accurate low-dimensional models for the experimental results can be achieved by using *tlsDMD*. We note that this improvement can be attained without explicit knowledge of the process and sensor noise characteristics.

6.5 Discussion and conclusions

It was shown in Section 6.2 that simple linear algebraic considerations can allow us to derive an estimate for the bias that exists in all standard formulations of DMD. This subsequently led to the formulation of the three modified algorithms that we suggest can be used to eliminate this bias. Section 6.3 showed that this predicted bias is indeed present in the results of DMD. Directly correcting for this bias term (Algorithm 3, *ncDMD*) was shown to almost completely eliminate this bias. While this modification demonstrates that our characterization of the dominant effects of noise was accurate, its usefulness is limited by the fact that it requires an accurate estimate of the noise covariance. Additionally, the presence of a Σ^{-2} term in correction factor used in *ncDMD* makes this computation unsuitable for cases with small singular values that are not truncated. On the other hand, the correction factor in Algorithm 3 may be applied to existing DMD results with minimal computational effort. Algorithms 4 (*fbDMD*) and 5 (*tlsDMD*), which do not require knowledge of the noise characteristics, were also found to correct for the bias, and also were able to reduce the random error across many noise realizations (as seen by smaller associated confidence ellipses in Figure 6.2, for example). Furthermore, *fbDMD* and *tlsDMD* were found in Section 6.3.5 to also compensate for the effect of process noise. This feature could be desirable or undesirable, depending on the purpose for which DMD is being applied. Note that this is also consistent for the findings in Section 6.4.2, where for a notionally periodic system, *tlsDMD* was found to give eigenvalues very close to the imaginary axis, despite (presumably) the presence of both sensor and process noise.

In practice, the examples examined in Sections 6.3.4, 6.4.1 and 6.4.2 suggest an overarching principle: while regular DMD can be accurate for identifying dominant dynamics that have much larger amplitudes than the noise in the data, accurate identification of the eigenvalues associated with lower amplitude modes (and in particular, their real components) can be significantly improved when using the modified DMD algorithms presented here. Conversely, if one is primarily concerned with the identification of modes and their frequencies of oscillation, and less concerned with accurate identification of growth/decay rates, then the effect of sensor noise is comparatively minimal, and subsequently the choice of DMD algorithm is less important.

Fundamentally, the bias in DMD arises because the algorithm is essentially a least-squares algorithm, which is designed for cases where the “independent” variable (which in DMD takes the form of the data \mathbf{Y}) is known accurately, and the “dependent” variable $\mathbf{Y}^\#$ contains

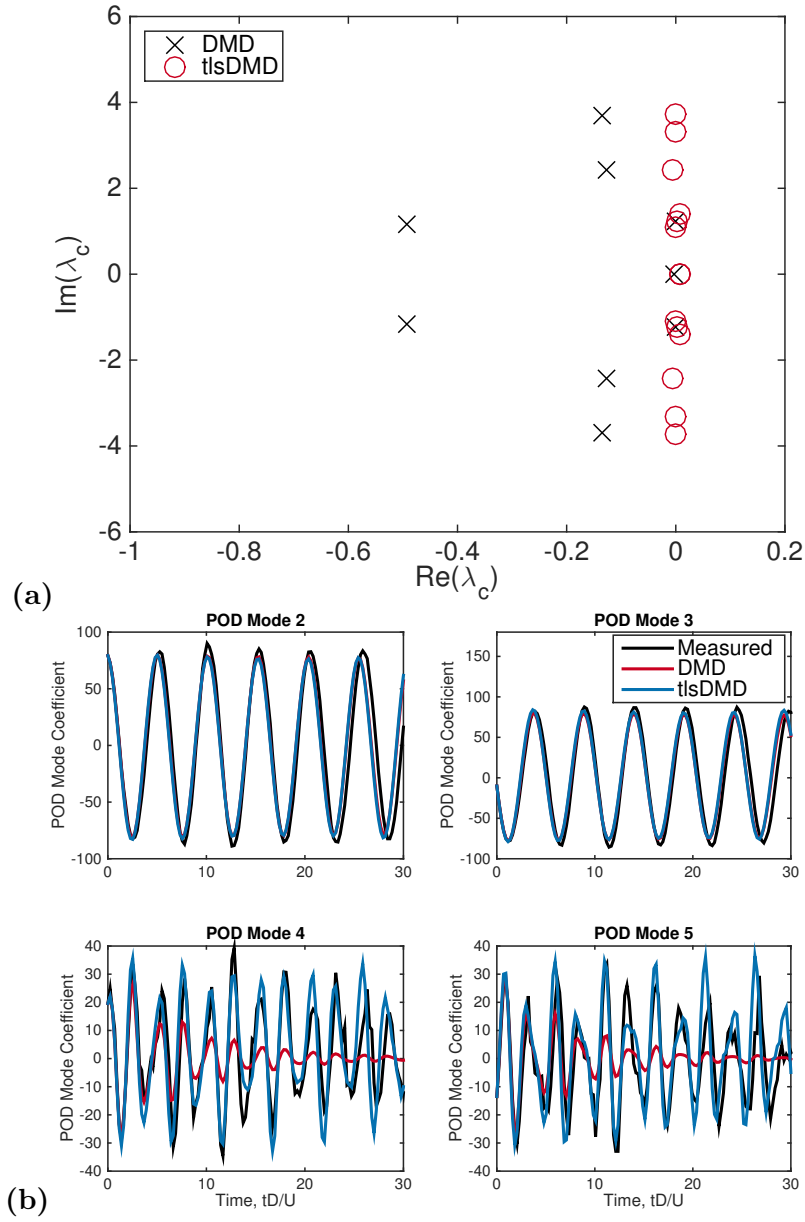


Figure 6.12: (a) Eigenvalues and (b) POD coefficients identified from applying DMD and tlsDMD to experimental vorticity data

the noise/error. In reality, since \mathbf{Y} and $\mathbf{Y}^\#$ should both be affected by noise, minimizing the error in both the \mathbf{Y} and $\mathbf{Y}^\#$ “directions” can allow for a more accurate answer to be obtained. One drawback of tlsDMD is that it requires taking the SVD of a larger matrix. Note that for cases where $n > m$ (i.e., the size of each snapshot is larger than the number of snapshots) and there is no truncation of POD modes corresponding to small singular values, DMD gives the minimum Frobenius norm (of \mathbf{A}) solution to $\mathbf{A}\mathbf{Y} = \mathbf{Y}^\#$. In this case, in principle neither fbDMD or tlsDMD should yield any improvements. In reality, however, if there is noise in the data, then we do not necessarily want an exact fit to the data, but rather an unbiased estimate of the noise-free dynamics. We may obtain this by truncating POD modes that are deemed to be mostly noise, and use some variant of DMD to identify the remaining dynamics. tlsDMD and fbDMD give very similar results, which suggests that fbDMD can be viewed as a computationally cheaper alternative to approximating the results of tlsDMD . Note that while fbDMD is often computationally cheaper, it relies on being able to invert the matrix $\tilde{\mathbf{B}}_m$, which might be an ill-conditioned operation for some data.

In Section 6.3.3, we compared the variants presented here with two recent optimization algorithms that have been proposed. The results show our algorithms outperforming both sparsity promoting DMD and OMD. Note that since these algorithms are not in closed form, but instead contain optimization procedures, the results depend somewhat on the specification of the relevant optimization parameters. In this comparison, our use of Algorithms 3–5 relied upon the projection onto a low dimensional subspace before applying DMD-type algorithms. We particularly note that the tlsDMD algorithm proposed here is slightly different from that given in Hemati et al. [69] due to this POD projection, which we found empirically to give improved results. We suggest that this is because the initial truncation of low-energy POD modes has a filtering effect that better isolates the true dynamics, at least for the datasets considered here. One could imagine, however, that in certain cases this projection could lead to significant degradation of results. For example, where the dynamically important modes are highly dissimilar to the dominant POD modes, the flexibility for the projection basis to be modified could be particularly advantageous. In such cases, sparsity promoting DMD or OMD could give more favorable results. In general, it is relatively common in system identification to use a subspace that is larger than the dimension of the underlying dynamics, and then later truncate to obtain a reduced-order model of an appropriate size/rank. This can be particularly important when the dealing with specific system inputs and outputs [122]. Juang and Pappa [80] discuss a number of ways in which true dynamic modes can be distinguished for noisy modes, in the context of the eigensystem realization algorithm. Tu et al. [161] further discusses how DMD modes can be scaled, from which appropriate modes can be chosen. The spDMD algorithm in Jovanović et al. [78] essentially automates this procedure, and comes with the additional potential advantage of not requiring a-priori knowledge of the dimension of the reduced-order dynamics to be identified. Note that it is also possible to combine the modifications to DMD proposed here with the OMD and spDMD optimization procedures, which could result in further improvements in some circumstances.

Though we used a large number of trials when testing our results on synthetic data in order to obtain statistically meaningful findings, in reality one would most likely not have this luxury with real data. In this case, it is important to understand for the size and quality

of the data to be analyzed, both the best algorithm to use, and the amount of confidence that should be had in the results of the chosen algorithm.

While this work has been motivated by and has largely focused on sensor noise (that is, noise which only affects measurements, and not the system dynamics), the characterization and removal of process noise (i.e., disturbances to the system states) is entirely another matter. Interestingly, the effect of process noise was identified analytically in [9] to be a parabolic decay in the growth rate of identified eigenvalues with increasing frequency. It turns out that a similar effect is observed here for the case of measurement noise. Isolating sensor noise from process noise (especially with limited prior information about the statistics of either) is an important and challenging task, particularly when dealing with more complex, turbulent flows, where the true dynamics exist on a wide range of spatial and temporal scales. The fact that DMD, ncDMD and tlsDMD/fbDMD each perform differently on these different forms of noise could itself be an important tool to this end.

Particularly in experimental data, users might typically preprocess data in a number of ways before considering applying DMD-type algorithms. It could be advantageous to investigate precisely how various averaging and smoothing operations affect the subsequent analysis of dynamics, and subsequently whether such post-processing and analysis can be achieved through a single algorithm.

Ultimately, having a larger selection of possible algorithms should be of benefit to researchers who desire the dynamical information that DMD-type algorithms can provide, who can choose based on the size of the data, amount of noise present, required accuracy of the results, and amount of computational resources available. One of the major advantages of DMD (and related algorithms) advocated in Schmid [134] is the fact that it requires only direct data measurements, without needing knowledge of any underlying system matrix, thus making it well suited to use on experimentally acquired data. Inevitably, however, data from experiments is always affected to some extent by noise. It is thus important to properly understand and quantify how noise can influence the results of DMD. Conversely, the quest for high quality data can often require large investments of both time and money. Formulating algorithms that are more robust to noisy data can be a cheaper alternative to obtain results of sufficient accuracy. As it becomes easier to generate and store increasingly large datasets, it is also important to recognize that simply feeding larger quantities of data (e.g., more snapshots) into a given algorithm does not guarantee desired improvements in the accuracy of their outputs, as illustrated in Figures 6.3 and 6.4.

The problem that fluid dynamicists face in extracting tractable information from large datasets is not unique to fluids, and rather transcends a wide variety of fields of study (although other fields are often not afforded the luxury of knowing the underlying differential equations). It thus should be fruitful to continue to investigate how previous and current developments across a wide range of other fields may be utilized in the study of fluids systems. We likewise hope that other areas can benefit from the work that is motivated by the desire to understand how fluids flow.

6.6 Appendix: Quantifying the size of the bias in DMD

We seek to quantify the magnitude of this bias present in DMD that was derived in Section 6.2.1, subject to certain simplifying assumptions on the nature of the data and noise. If the noise is uniform, and spatially and temporally independent, then

$$\mathbb{E}(\tilde{\mathbf{N}}\tilde{\mathbf{N}}^*) = \mathbb{E}(\mathbf{U}^*\mathbf{N}\mathbf{N}^*\mathbf{U}) = \mathbf{U}^*m\sigma_N^2\mathbf{U} = m\sigma_N^2\mathbf{I},$$

where σ_N^2 is the variance of each independent component of the noise matrix. Furthermore, if we assure that we are projecting onto the POD modes of the clean data, then $(\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^*) = \Sigma^2$, where $\mathbf{U}\Sigma\mathbf{V}^*$ is the singular value decomposition of \mathbf{Y} . Thus with these assumptions, Equation (6.2.4) can be simplified to give

$$\mathbb{E}(\tilde{\mathbf{A}}_m) = \tilde{\mathbf{A}}(\mathbf{I} - m\sigma_N^2\Sigma^{-2}). \quad (6.6.1)$$

The (diagonal) entries Σ_i^2 of Σ^2 have the interpretation of being the energy content of the i^{th} POD mode. We then should expect that $\Sigma_i^2 \sim mnq_i\sigma_Y^2$, where σ_Y^2 is the RMS value of the elements in the data matrix \mathbf{Y} , and $q_i = \frac{\Sigma_i^2}{\text{Trace}(\Sigma^2)}$ is the proportion of the total energy of the system contained in the i^{th} POD mode. For this scaling, we make the assumption that adding/removing rows and columns of data (i.e., varying m and n) does not affect either σ_Y or q_i . The bias term $m\sigma_N^2\Sigma^{-2}$ is a diagonal matrix whose i^{th} entry has a size $(e_b)_i$ proportional to

$$(e_b)_i \sim \frac{1}{nq_iSNR^2}, \quad (6.6.2)$$

where SNR is the signal-to-noise ratio. Thus sensor noise has the effect of reducing the diagonal entries of the computed $\tilde{\mathbf{A}}_m$ matrix by a multiplicative factor of $1 - \frac{\sigma_N^2}{nq_i\sigma_Y^2}$, which means that POD coefficients are predicted to decay more rapidly than they actually do. This effect is most pronounced for lower energy modes, for which the q_i is smaller. We thus expect to identify with DMD (continuous-time) eigenvalues that are further into the left half plane than they should be (or would be if we applied DMD to noise-free data). Duke et al. [48] argues in the case of periodic data that the growth rate of the eigenvalues should typically be the most challenging to identify, since there are a range of pre-existing methods that can identify frequencies. Here we have argued that it is precisely this growth rate that is most affected by noise. Importantly, the amount of bias is independent of m , which suggests that the bias component of the error will be particularly dominant when we have a large number of low-dimensional snapshots. Importantly, this suggests that one cannot always effectively reduce the effect of noise by simply using more snapshots of data, since the bias error will eventually become the dominant error.

While we can now quantify the magnitude of the bias in DMD, we do not as yet know how it compares to the random component of the error that would arise from a given realization of noise. To do this, we will estimate the typical size of the variance of individual entries of $\tilde{\mathbf{A}}$, using the standard definition

$$var \left[\tilde{\mathbf{A}}_{ij} \right] = \mathbb{E} \left\{ \left((\tilde{\mathbf{Y}}_m^\# \tilde{\mathbf{Y}}_m^+)_{ij} - \mathbb{E} \left[(\tilde{\mathbf{Y}}_m^\# \tilde{\mathbf{Y}}_m^+)_{ij} \right] \right) \left((\tilde{\mathbf{Y}}_m^\# \tilde{\mathbf{Y}}_m^+)_{ij} - \mathbb{E} \left[(\tilde{\mathbf{Y}}_m^\# \tilde{\mathbf{Y}}_m^+)_{ij} \right] \right) \right\}. \quad (6.6.3)$$

Referring back to Equation (6.2.1), if we exclude terms that are quadratic or higher in noise, and assume that the noise covariance matrix is sufficiently close to its expected value, we find that

$$\begin{aligned}
(\tilde{\mathbf{Y}}_m^\# \tilde{\mathbf{Y}}_m^+) - \mathbb{E} \left[(\tilde{\mathbf{Y}}_m^\# \tilde{\mathbf{Y}}_m^+) \right] \\
&= (\tilde{\mathbf{Y}}^\# + \tilde{\mathbf{N}}^\#)(\tilde{\mathbf{Y}} + \tilde{\mathbf{N}})(\tilde{\mathbf{Y}}\tilde{\mathbf{Y}}^* + \tilde{\mathbf{Y}}\tilde{\mathbf{N}}^* + \tilde{\mathbf{N}}\tilde{\mathbf{Y}}^* + \tilde{\mathbf{N}}\tilde{\mathbf{N}}^*)^{-1} \\
&\quad - \tilde{\mathbf{Y}}^\# \tilde{\mathbf{Y}}^+ - \mathbb{E}(\tilde{\mathbf{N}}\tilde{\mathbf{N}}^*)\Sigma^{-2} \\
&= \left[\tilde{\mathbf{Y}}^\# \tilde{\mathbf{Y}}^+ (\tilde{\mathbf{Y}}\tilde{\mathbf{N}}^* + \tilde{\mathbf{N}}\tilde{\mathbf{Y}}^*) + \tilde{\mathbf{N}}^\# \tilde{\mathbf{Y}}^* + \tilde{\mathbf{Y}}^\# \tilde{\mathbf{N}}^* \right] \Sigma^{-2}.
\end{aligned}$$

Elements of the terms $\tilde{\mathbf{Y}}\tilde{\mathbf{N}}^*$, $\tilde{\mathbf{N}}(\tilde{\mathbf{Y}}^\#)^*$, $\tilde{\mathbf{N}}^\#\tilde{\mathbf{Y}}^*$, and $\tilde{\mathbf{Y}}^\#\tilde{\mathbf{N}}^*$ are uncorrelated sums over m random terms, with each term in the sum having variance $nq_i\sigma_Y^2\sigma_N^2$ where as before q_i is the energy fraction in the i^{th} POD mode. This means that the sum will have variance $mnq_i\sigma_Y^2\sigma_N^2$. Assuming that $\tilde{\mathbf{Y}}^\#\tilde{\mathbf{Y}}^+ (= \tilde{\mathbf{A}})$ does not greatly change the magnitude of quantities that it multiplies, and assuming that q_i remains constant when varying m and n , this means that we find that

$$var \left[\tilde{\mathbf{A}}_{ij} \right] \sim \frac{\sigma_N^2}{mn\sigma_Y^2}. \quad (6.6.4)$$

Thus the expected size of the random error in applying DMD to noisy data is

$$e_r \sim \frac{1}{m^{1/2}n^{1/2}SNR}. \quad (6.6.5)$$

Comparing Equation (6.6.5) with Equation (6.6.2), we propose that the bias in DMD will be the dominant source of error whenever

$$m^{1/2}SNR > n^{1/2}.$$

Chapter 7

Conclusions and future work

This thesis has considered a number of topics related to modeling the dynamics of pitching airfoils, and the methods that can be used to model these and other systems encountered in fluid mechanics and aerodynamics.

In Chapter 3, a framework was proposed to construct models for the dynamics of a rapidly pitching airfoil across a range of angles of attack, encompassing flow that is fully, partially, and nowhere attached to the suction surface. The success of this model in predicting pressures and forces without all regimes shows its utility in modeling systems that exhibit a broad range of features across different regimes. Moreover, we demonstrated that the approach is sufficiently robust to be of use for experimental data.

Chapter 4 focused on investigating a particular feature of pitching airfoils, namely the effect of sinusoidal pitching at close to the “preferred” wake frequency, which can substantially change the structure of the wake, and in turn the forces on the airfoil. This systematic investigation is necessary to understand the phenomena, and how it affects quantities of interest, such as the lift coefficient. By using DMD, we were able to isolate the structures corresponding to the natural and forcing frequencies. A thorough understanding such effects is important for designing modeling techniques that will be capable of predicting the influence of such phenomena.

An alternative approach to generate nonlinear models from fluids data was presented in Chapter 5, where it was demonstrated that an extension to the DMD algorithm could be utilized to construct nonlinear models. Furthermore, it was demonstrated that this method was particularly advantageous when the quality or quantity of the data available was limited. In spite of the apparent robustness in this case, the effectiveness of data-driven modeling procedures is unavoidably limited to some extent by the quality of data available. Because of this, it is valuable to be able to quantify this effect, either theoretically or empirically. This issue is explicitly considered for the DMD algorithm in Chapter 6. As well as providing a theoretical explanation for why DMD can be sensitive to noise, this understanding was leveraged to formulate alternative algorithms that give improved accuracy when dealing with noise-corrupted datasets.

Successful modeling of fluids systems requires a combination of well chosen algorithms and methods of analysis with correct physical insight and interpretations. Ideally, this process can be symbiotic. On one hand, knowledge of a physical system can suggest the type of data that should be collected, and the nature of the analysis and algorithms that should be used

on the data. Conversely, an understanding for what the desired outcomes of applying an algorithm can lead to refinements and extensions of that algorithm to meet such demands. It is hoped that the contents of this thesis help researchers in applying and advancing each of these aspects.

The work of this thesis was motivated, in part, by preliminary results that used simple linear models to control lift. Implicit in this motivation is the fact that the identification accurate reduced-order models is not necessarily the final goal: ultimately, being able to use such models for effective manipulation and control of a fluids system, for example, gives such models more tangible value. The tools developed in this thesis will hopefully allow for these steps to be taken for increasingly complex systems, in a manner which is robust to the type of data being used to drive the procedure. More explicitly, a number of possible directions for future work that would extend the research presented in this thesis include:

- **Real-time estimation with limited sensors:** Chapter 3 showed that, with a limited number of measurements, one could equip the switched model with a Kalman filter to improve the estimates of the measurements. However, this procedure was implemented with phase-averaged data, and thus was could not be applied in real-time. It would be valuable to investigate if effective real-time state estimation could be achieved using only raw, noisy measurements.
- **Modeling complex systems:** The fluid/aerodynamic systems considered in this work, particularly those studied with DNS, are relatively simple, both in terms of the geometry (two-dimensional flat plate or circular cylinder), and the laminar, low Reynolds number conditions. Applying the nonlinear system identification algorithm discussed in Chapter 5 to more complex systems would be a logical next step.
- **Assimilating physics with data:** In Chapter 5, knowledge of the governing equations was utilized to guide the choice of observable functions. It could be possible to use additional properties of the governing equations to further improve the accuracy of identified models. For example, the energy-preserving property of the nonlinear terms in the Navier-Stokes equations could be used to add constraints to the least-squares optimization that underlies data-driven system identification.
- **Further investigation and mitigation of noise effects on modeling algorithms:** The analysis in Chapter 6 was limited only to DMD. While the same analysis extends trivially to related algorithms such as ERA, the effect of noisy data on EDMD, for example, is complicated by the fact that the data typically undergoes nonlinear transformations before the DMD algorithm is applied. A proper investigation of this could be useful if trying to apply EDMD to noisy experimental data. Beyond this, one can also analyze the effects of inaccuracies in on the tasks that such models might be used for, such as state estimation or control. Interesting recent work related to this idea is given in Dovetta et al. [46].
- **Inferring other pertinent properties of fluids systems from modeling algorithms:** As well as being useful for prediction and control purposes, it would be interesting to investigate what additional fundamental understanding of fluids systems

can be extracted from models identified using the procedures discussed in this thesis. For example, one might hope that the nonlinear models obtained in Chapter 5 could be utilized to calculate quantities such as the dominant eigenmodes of the linear system about equilibria. On a somewhat related note, the DMD modes that emerge in Chapter 4 at the frequency of the marginally stable wake mode look very similar to what linear stability eigenmodes look like for such systems. Further investigation of this could be fruitful for understanding the interacting between the dynamics of the pitching and natural frequencies of instability in the wake.

Bibliography

- [1] M. Acharya and M. H. Metwally. Unsteady pressure field and vorticity production over a pitching airfoil. *AIAA Journal*, 30(2):403–411, 1992.
- [2] S. Ahuja and C. W. Rowley. Feedback control of unstable steady states of flow past a flat plate using reduced-order estimators. *Journal of Fluid Mechanics*, 645:447–478, 2010.
- [3] E. Åkervik, L. Brandt, D. S. Henningson, J. Höpfner, O. Marxen, and P. Schlatter. Steady solutions of the Navier-Stokes equations by selective frequency damping. *Physics of Fluids*, 18(6):068102, 2006.
- [4] A. Alla and J. N. Kutz. Nonlinear model order reduction via dynamic mode decomposition. *arXiv preprint arXiv:1602.05080*, 2016.
- [5] J. D. Anderson Jr. *A history of aerodynamics: and its impact on flying machines*, volume 8. Cambridge University Press, 1999.
- [6] J. Annoni and P. Seiler. A method to construct reduced-order parameter-varying models. *International Journal of Robust and Nonlinear Control*, pages n/a–n/a, 2016. ISSN 1099-1239. doi: 10.1002/rnc.3586. URL <http://dx.doi.org/10.1002/rnc.3586>. rnc.3586.
- [7] N. Aubry, P. Holmes, J. L. Lumley, and E. Stone. The dynamics of coherent structures in the wall region of a turbulent boundary layer. *Journal of Fluid Mechanics*, 192:115–173, 1988.
- [8] S. Bagheri. Koopman-mode decomposition of the cylinder wake. *Journal of Fluid Mechanics*, 726:596–623, 2013.
- [9] S. Bagheri. Effects of weak noise on oscillating flows: Linking quality factor, floquet modes, and koopman spectrum. *Physics of Fluids*, 26(9):094104, 2014.
- [10] Z. Bai and J. W. Silverstein. *Spectral analysis of large dimensional random matrices*. Springer, 2009.
- [11] Y. S. Baik, L. P. Bernal, K. Granlund, and M. V. Ol. Unsteady force generation and vortex dynamics of pitching and plunging aerofoils. *Journal of Fluid Mechanics*, 709: 37–68, 2012.

- [12] M. Balajewicz and E. Dowell. Reduced-order modeling of flutter and limit-cycle oscillations using the sparse Volterra series. *Journal of Aircraft*, 49(6):1803–1812, 2012.
- [13] M. J. Balajewicz, E. H. Dowell, and B. R. Noack. Low-dimensional modelling of high-reynolds-number shear flows incorporating constraints from the Navier–Stokes equation. *Journal of Fluid Mechanics*, 729:285–308, 2013.
- [14] M. J. Balajewicz, I. Tezaur, and E. H. Dowell. Minimal subspace rotation on the Stiefel manifold for stabilization and enhancement of projection-based reduced order models for the compressible Navier-Stokes equations. *arXiv preprint arXiv:1504.06661*, 2015.
- [15] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera. An ‘empirical interpolation’ method: application to efficient reduced-basis discretization of partial differential equations. *Comptes Rendus Mathématique*, 339(9):667–672, 2004.
- [16] B. A. Belson, O. Semeraro, C. W. Rowley, and D. S. Henningson. Feedback control of instabilities in the two-dimensional blasius boundary layer: The role of sensors and actuators. *Physics of Fluids*, 25(5):054106, 2013.
- [17] B. A. Belson, J. H. Tu, and C. W. Rowley. A parallelized model reduction library. *ACM T. Math. Software*, 2013.
- [18] G. Berkooz, P. Holmes, and J. L. Lumley. The proper orthogonal decomposition in the analysis of turbulent flows. *Annual review of fluid mechanics*, 25(1):539–575, 1993.
- [19] J. Birch and M. Dickinson. Spanwise flow and the attachment of the leading-edge vortex on insect wings. *Nature*, 412:729–733, 2001.
- [20] J. Borggaard, A. Duggeby, A. Hay, T. Iliescu, and Z. Wang. Reduced-order modeling of turbulent flows. In *Proceedings of MTNS*, 2008.
- [21] P. L. Boyland. Bifurcations of circle maps: Arnol’d tongues, bistability and rotation intervals. *Communications in Mathematical Physics*, 106(3):353–381, 1986.
- [22] S. L. Brunton and B. R. Noack. Closed-loop turbulence control: progress and challenges. *Applied Mechanics Reviews*, 67:050801, Sept. 2015.
- [23] S. L. Brunton, C. W. Rowley, and D. R. Williams. Reduced-order unsteady aerodynamic models at low Reynolds numbers. *Journal of Fluid Mechanics*, 724:203–233, 2013.
- [24] S. L. Brunton, S. T. M. Dawson, and C. W. Rowley. State-space model identification and feedback control of unsteady aerodynamic forces. *Journal of Fluids and Structures*, 50:253–270, 2014.
- [25] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.

- [26] R. H. Cabell, M. A. Kegerise, D. E. Cox, and G. P. Gibbs. Experimental feedback control of flow-induced cavity tones. *AIAA journal*, 44(8):1807–1816, 2006.
- [27] S. Chaturantabut and D. C. Sorensen. Nonlinear model reduction via discrete empirical interpolation. *SIAM Journal on Scientific Computing*, 32(5):2737–2764, 2010.
- [28] K. K. Chen, J. H. Tu, and C. W. Rowley. Variants of dynamic mode decomposition: boundary condition, Koopman, and Fourier analyses. *Journal of Nonlinear Science*, 22(6):887–915, 2011.
- [29] X. Cheng and A. Singer. The spectrum of random inner-product kernel matrices. *Random Matrices: Theory and Applications*, 2(04), 2013.
- [30] J. Choi, T. Colonius, and D. R. Williams. Surging and plunging oscillations of an airfoil at low Reynolds number. *Journal of Fluid Mechanics*, 763:237–253, 2015.
- [31] A. J. Chorin. Numerical solution of the navier-stokes equations. *Mathematics of computation*, 22(104):745–762, 1968.
- [32] P. G. Choudhuri and D. Knight. Effects of compressibility, pitch rate, and reynolds number on unsteady incipient leading-edge boundary layer separation over a pitching airfoil. *Journal of Fluid Mechanics*, 308:195–217, 1996.
- [33] D. Cleaver, Z. Wang, and I. Gursul. Investigation of high-lift mechanisms for a flat-plate airfoil undergoing small-amplitude plunging oscillations. *AIAA journal*, 51(4):968–980, 2013.
- [34] D. J. Cleaver, Z. Wang, I. Gursul, and M. Visbal. Lift enhancement by means of small-amplitude airfoil oscillations at low Reynolds numbers. *AIAA journal*, 49(9):2018–2033, 2011.
- [35] D. J. Cleaver, Z. Wang, and I. Gursul. Bifurcating flows of plunging aerofoils at high Strouhal numbers. *Journal of Fluid Mechanics*, 708:349–376, 2012.
- [36] T. Colonius and K. Taira. A fast immersed boundary method using a nullspace approach and multi-domain far-field boundary conditions. *Computer Methods in Applied Mechanics and Engineering*, 197:2131–2146, 2008.
- [37] L. Cordier, E. Majd, B. Abou, and J. Favier. Calibration of POD reduced-order models using Tikhonov regularization. *International Journal for Numerical Methods in Fluids*, 63(2):269–296, 2010.
- [38] L. Cordier, B. R. Noack, G. Tissot, G. Lehnasch, J. Delville, M. Balajewicz, G. Daviller, and R. K. Niven. Identification strategies for model-based control. *Experiments in Fluids*, 54(8):1–21, 2013.
- [39] J. Dandois and P. Pamart. NARX modeling and extremum-seeking control of a separation. *AerospaceLab*, 6:1–13, 2013.

- [40] S. T. M. Dawson, S. L. Brunton, and C. W. Rowley. Nonlinear switched models for control of unsteady forces on a rapidly pitching airfoil. *66th Annual Meeting of the APS Division of Fluid Dynamics*, 2013.
- [41] S. T. M. Dawson, N. K. Schiavone, C. W. Rowley, and D. 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.
- [42] S. T. M. Dawson, D. C. Floryan, C. W. Rowley, and M. S. Hemati. Lift enhancement of high angle of attack airfoils using periodic pitching. In *54th AIAA Aerospace Sciences Meeting*, page 2069, 2016.
- [43] S. T. M. Dawson, M. S. Hemati, M. O. Williams, and C. W. Rowley. Characterizing and correcting for the effect of sensor noise in the dynamic mode decomposition. *Experiments in Fluids*, 57(42):1–19, 2016.
- [44] A. Deane, I. Kevrekidis, G. E. Karniadakis, and S. Orszag. Low-dimensional models for complex geometry flows: Application to grooved channels and circular cylinders. *Physics of Fluids A: Fluid Dynamics (1989-1993)*, 3(10):2337–2354, 1991.
- [45] P. A. Dewey, A. Carriou, and A. J. Smits. On the relationship between efficiency and wake structure of a batoid-inspired oscillating fin. *Journal of Fluid Mechanics*, 691:245–266, 2012.
- [46] N. Dovetta, P. J. Schmid, and D. Sipp. Uncertainty propagation in model extraction by system identification and its implication for control design. *Journal of Fluid Mechanics*, 791:214–236, 2016.
- [47] D. Duke, D. Honnery, and J. Soria. Experimental investigation of nonlinear instabilities in annular liquid sheets. *Journal of Fluid Mechanics*, 691:594–604, 2012.
- [48] D. Duke, J. Soria, and D. Honnery. An error analysis of the dynamic mode decomposition. *Experiments in Fluids*, 52(2):529–542, 2012.
- [49] R. Dunne and B. J. McKeon. Dynamic stall on a pitching and surging airfoil. *Experiments in Fluids*, 56(8):1–15, 2015. ISSN 1432-1114.
- [50] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.
- [51] J. D. Eldredge, C. Wang, and M. V. OL. A computational study of a canonical pitch-up, pitch-down wing maneuver. AIAA Paper 2009-3687, 39th Fluid Dynamics Conference, June 2009.
- [52] B. P. Epps and A. H. Techet. An error threshold criterion for singular value decomposition modes extracted from PIV data. *Experiments in Fluids*, 48(2):355–367, 2010.
- [53] N. B. Erichson and C. Donovan. Randomized low-rank dynamic mode decomposition for motion detection. *Computer Vision and Image Understanding*, 146:40–50, 2016.

- [54] T. L. Flinois and A. S. Morgans. Feedback control of unstable flows: a direct modelling approach using the eigensystem realisation algorithm. *Journal of Fluid Mechanics*, 793: 41–78, 2016.
- [55] M. Gad-el Hak. *Flow Control: Passive, Active, and Reactive Flow Management*. Cambridge University Press, 2007.
- [56] I. E. Garrick. On some reciprocal relations in the theory of nonstationary flows. Technical Report 629, NACA, 1938.
- [57] M. Gharib, E. Rambod, and K. Shariff. A universal time scale for vortex ring formation. *Journal of Fluid Mechanics*, 360:121–140, 1998.
- [58] B. Glaz, L. Liu, and P. P. Friedmann. Reduced-order nonlinear unsteady aerodynamic modeling using a surrogate-based recurrence framework. *AIAA Journal*, 48(10):2418–2429, 2010.
- [59] G. H. Golub and C. F. Van Loan. *Matrix computations*, volume 3. JHU Press, 2012.
- [60] M. Goman and A. Khrabrov. State-space representation of aerodynamic characteristics of an aircraft at high angles of attack. *Journal of Aircraft*, 31(5):1109–1115, 1994.
- [61] F. Gómez, H. M. Blackburn, M. Rudman, B. J. McKeon, M. Luhar, R. Moarref, and A. S. Sharma. On the origin of frequency sparsity in direct numerical simulations of turbulent pipe flow. *Physics of Fluids*, 26(10):101703, 2014.
- [62] P. J. Goulart, A. Wynn, and D. Pearson. Optimal mode decomposition for high dimensional systems. In *CDC*, pages 4965–4970, 2012.
- [63] K. O. Granlund, M. V. Ol, and L. P. Bernal. Unsteady pitching flat plates. *Journal of Fluid Mechanics*, 733:R5, 2013.
- [64] M. Grilli, P. J. Schmid, S. Hickel, and N. A. Adams. Analysis of unsteady behaviour in shockwave turbulent boundary layer interaction. *Journal of Fluid Mechanics*, 700: 16–28, 2012.
- [65] J. Grosek and J. N. Kutz. Dynamic mode decomposition for real-time background/-foreground separation in video. *arXiv preprint arXiv:1404.7592*, 2014.
- [66] M. S. Hemati, J. D. Eldredge, and J. L. Speyer. Improving vortex models via optimal control theory. *Journal of Fluids and Structures*, 49:91–111, 2014.
- [67] M. S. Hemati, M. O. Williams, and C. W. Rowley. Dynamic mode decomposition for large and streaming datasets. *Physics of Fluids*, 26(11):111701, 2014.
- [68] M. S. Hemati, S. T. M. Dawson, and C. W. Rowley. Unsteady aerodynamic response modeling: A parameter-varying approach. *AIAA Aerospace Sciences Meeting*, 2015.

- [69] M. S. Hemati, C. W. Rowley, E. A. Deem, and L. N. Cattafesta. De-biasing the dynamic mode decomposition for applied Koopman spectral analysis. *arXiv preprint arXiv:1502.03854*, 2015.
- [70] M. S. Hemati, S. T. M. Dawson, and C. W. Rowley. Parameter-varying models for unsteady aerodynamic response prediction. *to appear in AIAA Journal*, 2016.
- [71] A. Hervé, D. Sipp, P. J. Schmid, and M. Samuelides. A physics-based approach to flow control using system identification. *Journal of Fluid Mechanics*, 702:26–58, 2012.
- [72] B. L. Ho and R. E. Kalman. Effective construction of linear state-variable models from input/output data. In *Proceedings of the 3rd Annual Allerton Conference on Circuit and System Theory*, pages 449–459, 1965.
- [73] P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley. *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge University Press, 2012.
- [74] S. J. Illingworth. Model-based control of vortex shedding at low reynolds numbers. *Theoretical and Computational Fluid Dynamics*, pages 1–20, 2016.
- [75] S. J. Illingworth, A. S. Morgans, and C. W. Rowley. Feedback control of flow resonances using balanced reduced-order models. *Journal of Sound and Vibration*, 330(8):1567–1581, 2011.
- [76] S. J. Illingworth, A. S. Morgans, and C. W. Rowley. Feedback control of cavity flow oscillations using simple linear models. *Journal of Fluid Mechanics*, 709:223–248, 2012.
- [77] T. Jardin and Y. Bury. Lagrangian and spectral analysis of the forced flow past a circular cylinder using pulsed tangential jets. *Journal of Fluid Mechanics*, 696:285–300, 2012.
- [78] M. R. Jovanović, P. J. Schmid, and J. W. Nichols. Sparsity-promoting dynamic mode decomposition. *Physics of Fluids (1994-present)*, 26(2):024103, 2014.
- [79] J. N. Juang and R. S. Pappa. An eigensystem realization algorithm for modal parameter identification and model reduction. *Journal of Guidance, Control and Dynamics*, 8(5):620–627, 1985.
- [80] J. N. Juang and R. S. Pappa. Effects of noise on modal parameters identified by the eigensystem realization algorithm. *Journal of Guidance, Control, and Dynamics*, 9(3): 294–303, 1986.
- [81] J. N. Juang, J. Cooper, and J. Wright. An eigensystem realization algorithm using data correlations (era/dc) for modal parameter identification. *Control-Theory and Advanced Technology*, 4(1):5–14, 1988.
- [82] J. N. Juang, M. Phan, L. G. Horta, and R. W. Longman. Identification of observer/Kalman filter Markov parameters: Theory and experiments. Technical Memorandum 104069, NASA, 1991.

- [83] F. Juillet, B. McKeon, and P. J. Schmid. Experimental control of natural perturbations in channel flow. *Journal of Fluid Mechanics*, 752:296–309, 2014.
- [84] E. Kaiser, B. R. Noack, L. Cordier, A. Spohn, M. Segond, M. Abel, G. Daviller, J. Östh, S. Krajnović, and R. K. Niven. Cluster-based reduced-order modelling of a mixing layer. *Journal of Fluid Mechanics*, 754:365–414, 2014.
- [85] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Fluids Engineering*, 82(1):35–45, 1960.
- [86] S. M. Kaplan, A. Altman, and M. Ol. Wake vorticity measurements for low aspect ratio wings at low Reynolds number. *Journal of Aircraft*, 44(1):241–251, 2007.
- [87] D.-H. Kim and J.-W. Chang. Unsteady boundary layer for a pitching airfoil at low Reynolds numbers. *Journal of mechanical science and technology*, 24(1):429–440, 2010.
- [88] B. O. Koopman. Hamiltonian systems and transformation in Hilbert space. *Proceedings of the National Academy of Sciences of the United States of America*, 17(5):315, 1931.
- [89] S.-Y. Kung. A new identification and model reduction algorithm via singular value decomposition. In *Proceedings of the 12th Asilomar conference on circuits, systems and computers, Pacific Grove*, pages 705–714, 1978.
- [90] J. H. Lee, A. Seena, S. hyun Lee, and H. J. Sung. Turbulent boundary layers over rod- and cube-roughened walls. *Journal of Turbulence*, 13:N40, 2012.
- [91] J. G. Leishman. *Principles of Helicopter Aerodynamics*. Cambridge University Press, Cambridge, England, 2 edition, 2006.
- [92] D. Luchtenburg, B. Noack, and M. Schlegel. An introduction to the POD Galerkin method for fluid flows with analytical examples and matlab source codes. *Berlin Institute of Technology MB1, Muller-Breslau-Strabe*, 11, 2009.
- [93] D. M. Luchtenburg, C. W. Rowley, M. W. Lohry, L. Martinelli, and R. F. Stengel. Unsteady high-angle-of-attack aerodynamic models of a generic jet transport. *Journal of Aircraft*, 52(3):890–895, 2015.
- [94] D. J. Lucia, P. S. Beran, and W. A. Silva. Aeroelastic system development using proper orthogonal decomposition and Volterra theory. *Journal of aircraft*, 42(2):509–518, 2005.
- [95] J. L. Lumley. *Stochastic tools in turbulence*. Courier Dover Publications, 2007.
- [96] K. Y. Ma, P. Chirarattananon, S. B. Fuller, and R. J. Wood. Controlled flight of a biologically inspired, insect-scale robot. *Science*, 340(6132):603–607, 2013.
- [97] W. McCroskey. Unsteady airfoils. *Annual Review of Fluid Mechanics*, 14(1):285–311, 1982.

- [98] I. Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1-3):309–325, 2005.
- [99] I. Mezić. Analysis of fluid flows via spectral properties of the Koopman operator. *Annual Review of Fluid Mechanics*, 45:357–378, 2013.
- [100] M. Milano and M. Gharib. Uncovering the physics of flapping flat plates with artificial evolution. *Journal of Fluid Mechanics*, 534:403–409, 2005.
- [101] J. Moehlis, T. Smith, P. Holmes, and H. Faisst. Models for turbulent plane couette flow using the proper orthogonal decomposition. *Physics of Fluids*, 14(7):2493–2507, 2002.
- [102] A. T. Mohan, M. R. Visbal, and D. V. Gaitonde. Model reduction and analysis of deep dynamic stall on a plunging airfoil using dynamic mode decomposition. *AIAA Aerospace Sciences Meeting*, 2015.
- [103] T. W. Muld, G. Efraimsson, and D. S. Henningson. Flow structures around a high-speed train extracted using proper orthogonal decomposition and dynamic mode decomposition. *Computers & Fluids*, 57:87–97, 2012.
- [104] N. E. Murray and L. S. Ukeiley. Modified quadratic stochastic estimation of resonating subsonic cavity flow. *Journal of Turbulence*, 8:N53, 2007.
- [105] B. R. Noack and H. Eckelmann. A global stability analysis of the steady and periodic cylinder wake. *Journal of Fluid Mechanics*, 270:297–330, 1994.
- [106] B. R. Noack, K. Afanasiev, M. Morzynski, G. Tadmor, and F. Thiele. A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *Journal of Fluid Mechanics*, 497:335–363, 2003.
- [107] B. R. Noack, M. Morzynski, and G. Tadmor. *Reduced-order modelling for flow control*, volume 528. Springer Science & Business Media, 2011.
- [108] J. Östh, B. R. Noack, S. Krajnović, D. Barros, and J. Borée. On the need for a nonlinear subscale turbulence term in POD models as exemplified for a high-Reynolds-number flow over an ahmed body. *Journal of Fluid Mechanics*, 747:518–544, 2014.
- [109] C. Pan, D. Yu, and J. Wang. Dynamical mode decomposition of Gurney flap wake flow. *Theoretical and Applied Mechanics Letters*, 1(1):012002, 2011.
- [110] C. Pan, D. Xue, and J. Wang. On the accuracy of dynamic mode decomposition in estimating instability of wave packet. *Experiments in Fluids*, 56(8):1–15, 2015.
- [111] J. B. Perot. An analysis of the fractional step method. *Journal of Computational Physics*, 108(1):51–58, 1993.
- [112] L. Perret, E. Collin, and J. Delville. Polynomial identification of POD based low-order dynamical system. *Journal of Turbulence*, 7:N17, 2006.

- [113] C. S. Peskin. Flow patterns around heart valves: a numerical method. *Journal of computational physics*, 10(2):252–271, 1972.
- [114] B. Podvin. A proper-orthogonal-decomposition–based model for the wall layer of a turbulent channel flow. *Physics of Fluids*, 21(1):015111, 2009.
- [115] R. J. Prazenica, P. H. Reisenthel, A. J. Kurdila, and M. J. Brenner. Volterra kernel extrapolation for modeling nonlinear aeroelastic systems at novel flight conditions. *Journal of Aircraft*, 44(1):149–162, 2007.
- [116] J. L. Proctor, S. L. Brunton, and J. N. Kutz. Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems*, 15(1):142–161, 2016.
- [117] M. Provansal, C. Mathis, and L. Boyer. Bénard-von Kármán instability: transient and forced regimes. *Journal of Fluid Mechanics*, 182:1–22, 1987.
- [118] S. Pyatykh, J. Hesser, and L. Zheng. Image noise level estimation by principal component analysis. *Image Processing, IEEE Transactions on*, 22(2):687–699, 2013.
- [119] S. J. Qin. An overview of subspace identification. *Computers & chemical engineering*, 30(10):1502–1513, 2006.
- [120] M. Rajaei, S. K. Karlsson, and L. Sirovich. Low-dimensional description of free-shear-flow coherent structures and their dynamical behaviour. *Journal of Fluid Mechanics*, 258:1–29, 1994.
- [121] D. Rempfer and H. F. Fasel. Dynamics of three-dimensional coherent structures in a flat-plate boundary layer. *Journal of Fluid Mechanics*, 275:257–283, 1994.
- [122] C. W. Rowley. Model reduction for fluids using balanced proper orthogonal decomposition. *International Journal of Bifurcation and Chaos*, 15(3):997–1013, 2005.
- [123] C. W. Rowley and S. T. Dawson. Model reduction for flow analysis and control. *Annual Review of Fluid Mechanics*, 49(1), 2017.
- [124] C. W. Rowley and J. E. Marsden. Reconstruction equations and the Karhunen–Loève expansion for systems with symmetry. *Physica D: Nonlinear Phenomena*, 142(1):1–19, 2000.
- [125] C. W. Rowley and D. R. Williams. Dynamics and control of high-Reynolds-number flow over open cavities. *Annual Review of Fluid Mechanics*, 38:251–276, 2006.
- [126] C. W. Rowley, T. Colonius, and R. M. Murray. Dynamical models for control of cavity oscillations. *AIAA paper*, 2126(2001):2126–34, 2001.
- [127] C. W. Rowley, I. G. Kevrekidis, J. E. Marsden, and K. Lust. Reduction and reconstruction for self-similar dynamical systems. *Nonlinearity*, 16(4):1257, 2003.

- [128] C. W. Rowley, T. Colonius, and R. M. Murray. Model reduction for compressible flows using POD and Galerkin projection. *Physica D: Nonlinear Phenomena*, 189(1):115–129, 2004.
- [129] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson. Spectral analysis of nonlinear flows. *Journal of Fluid Mechanics*, 641(1):115–127, 2009.
- [130] S. P. Sane. The aerodynamics of insect flight. *The Journal of Experimental Biology*, 206(23):4191–4208, 2003.
- [131] T. Sarpkaya. Vortex-induced oscillations: a selective review. *Journal of Applied Mechanics*, 46(2):241–258, 1979.
- [132] T. Sayadi and P. J. Schmid. Parallel data-driven decomposition algorithm for large-scale datasets: with application to transitional boundary layers. *Theoretical and Computational Fluid Dynamics*, pages 1–14, 2016.
- [133] T. Sayadi, P. J. Schmid, J. W. Nichols, and P. Moin. Reduced-order representation of near-wall structures in the late transitional boundary layer. *Journal of Fluid Mechanics*, 748:278–301, 6 2014.
- [134] P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, 2010.
- [135] P. J. Schmid. Application of the dynamic mode decomposition to experimental data. *Experiments in Fluids*, 50(4):1123–1130, 2011.
- [136] P. J. Schmid and J. Sesterhenn. Dynamic mode decomposition of numerical and experimental data. In *61st Annual Meeting of the APS Division of Fluid Dynamics*. American Physical Society, 2008.
- [137] P. J. Schmid, L. Li, M. P. Juniper, and O. Pust. Applications of the dynamic mode decomposition. *Theoretical and Computational Fluid Dynamics*, 25(1-4):249–259, 2011.
- [138] A. Seena and H. J. Sung. Dynamic mode decomposition of turbulent cavity flows for self-sustained oscillations. *International Journal of Heat and Fluid Flow*, 32(6):1098–1110, 2011.
- [139] J. S. Shamma and M. Athans. Gain scheduling: potential hazards and possible remedies. *IEEE Control Systems Magazine*, 12(3):101–107, 1992.
- [140] C. Shih, L. Lourenco, L. Van Dommelen, and A. Krothapalli. Unsteady flow past an airfoil pitching at a constant rate. *AIAA journal*, 30(5):1153–1161, 1992.
- [141] W. A. Silva and R. E. Bartels. Development of reduced-order models for aeroelastic analysis and flutter prediction using the CFL3Dv6.0 code. *Journal of Fluids and Structures*, 19:729–745, 2004.
- [142] A. Singer and H.-T. Wu. Two-dimensional tomography from noisy projections taken at unknown random directions. *SIAM Journal on Imaging Sciences*, 6(1):136, 2013.

- [143] S. Skogestad and I. Postlethwaite. *Multivariable Feedback Control: Analysis and Design*. John Wiley & Sons, Inc., Hoboken, New Jersey, 2 edition, 2005.
- [144] T. R. Smith, J. Moehlis, and P. Holmes. Low-dimensional modelling of turbulence using the proper orthogonal decomposition: a tutorial. *Nonlinear Dynamics*, 41(1-3): 275–307, 2005.
- [145] K. Sreenivasan, P. Strykowski, and D. Olinger. Hopf bifurcation, Landau equation, and vortex shedding behind circular cylinders. In *Forum on unsteady flow separation*, volume 1, pages 1–13, 1987.
- [146] R. F. Stengel. *Optimal Control and Estimation*. Dover, 1994.
- [147] G. Stewart. Perturbation theory for the singular value decomposition. in *SVD and Signal Processing, 11: Algorithms, Analysis and Applications*, 1991.
- [148] M. Stewart. Perturbation of the svd in the presence of small singular values. *Linear Algebra and its Applications*, 419(1):53–77, 2006.
- [149] K. Taira and T. Colonius. The immersed boundary method: a projection approach. *Journal of Computational Physics*, 225(2):2118–2137, 2007.
- [150] K. Taira and T. Colonius. Three-dimensional flows around low-aspect-ratio flat-plate wings at low Reynolds numbers. *Journal of Fluid Mechanics*, 623:187–207, 2009.
- [151] T. Tao and V. Vu. Random covariance matrices: Universality of local statistics of eigenvalues. *The Annals of Probability*, 40(3):1285–1315, 2012.
- [152] M. Taya, E. Van Volkenburgh, M. Mizunami, and S.-h. Nomura. *Bioinspired Actuators and Sensors*. Cambridge University Press, 2016.
- [153] R. Temam. Sur l’approximation de la solution des équations de navier-stokes par la méthode des pas fractionnaires (ii). *Archive for Rational Mechanics and Analysis*, 33 (5):377–385, 1969.
- [154] T. Theodorsen. General theory of aerodynamic instability and the mechanism of flutter. Technical Report 496, NACA, 1935.
- [155] G. Tissot, L. Cordier, N. Benard, and B. R. Noack. Model reduction using dynamic mode decomposition. *Comptes Rendus Mécanique*, 342(6):410–416, 2014.
- [156] L. Tregidgo, Z. Wang, and I. Gursul. Frequency lock-in phenomenon for self-sustained roll oscillations of rectangular wings undergoing a forced periodic pitching motion. *Physics of Fluids (1994-present)*, 24(11):117101, 2012.
- [157] G. Triantafyllou, M. Triantafyllou, and M. Grosenbaugh. Optimal thrust development in oscillating foils with application to fish propulsion. *Journal of Fluids and Structures*, 7(2):205–224, 1993.

- [158] M. Triantafyllou, G. Triantafyllou, and R. Gopalkrishnan. Wake mechanics for thrust generation in oscillating foils. *Physics of Fluids A: Fluid Dynamics (1989-1993)*, 3(12):2835–2837, 1991.
- [159] J. H. Tu, C. W. Rowley, E. Aram, and R. Mittal. Koopman spectral analysis of separated flow over a finite-thickness flat plate with elliptical leading edge. In *49th AIAA Aerospace Sciences Meeting and Exhibit*, pages AIAA–Paper. AIAA, 2011.
- [160] J. H. Tu, C. W. Rowley, J. N. Kutz, and J. K. Shang. Spectral analysis of fluid flows using sub-Nyquist-rate PIV data. *Experiments in Fluids*, 55(9):1–13, 2014.
- [161] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz. On dynamic mode decomposition: Theory and applications. *Journal of Computational Dynamics*, 1(2):391–421, 2014.
- [162] L. Ukeiley, L. Cordier, R. Manceau, J. Delville, M. Glauser, and J. Bonnet. Examination of large-scale structures in a turbulent plane mixing layer. part 2. dynamical systems model. *Journal of Fluid Mechanics*, 441:67–108, 2001.
- [163] M. Verhaegen and P. Dewilde. Subspace model identification part 1. the output-error state-space model identification class of algorithms. *International Journal of Control*, 56(5):1187–1210, 1992.
- [164] M. Viberg. Subspace-based methods for the identification of linear time-invariant systems. *Automatica*, 31(12):1835–1851, 1995.
- [165] J. J. Videler, E. J. Samhuis, and G. D. E. Povel. Leading-edge vortex lifts swifts. *Science*, 306:1960–1962, 2004.
- [166] M. R. Visbal and J. Shang. Investigation of the flow structure around a rapidly pitching airfoil. *AIAA journal*, 27(8):1044–1051, 1989.
- [167] H. Wagner. Über die Entstehung des dynamischen Auftriebes von Tragflügeln. *Zeitschrift für Angewandte Mathematic und Mechanik*, 5(1):17–35, 1925.
- [168] Z. Wang, I. Akhtar, J. Borggaard, and T. Iliescu. Two-level discretizations of non-linear closure models for proper orthogonal decomposition. *Journal of Computational Physics*, 230(1):126–146, 2011.
- [169] Z. Wang, I. Akhtar, J. Borggaard, and T. Iliescu. Proper orthogonal decomposition closure models for turbulent flows: a numerical comparison. *Computer Methods in Applied Mechanics and Engineering*, 237:10–26, 2012.
- [170] Z. J. Wang. Dissecting insect flight. *Annual Review of Fluid Mechanics*, 37:183–210, 2005.
- [171] D. R. Williams, X. An, S. Iliev, R. King, and F. Reißner. Dynamic hysteresis control of lift on a pitching wing. *Experiments in Fluids*, 56(5):1–12, 2015.

- [172] M. O. Williams, I. G. Kevrekidis, and C. W. Rowley. A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, pages 1–40, 2015.
- [173] M. O. Williams, C. W. Rowley, and I. G. Kevrekidis. A kernel-based method for data-driven Koopman spectral analysis. *Journal of Computational Dynamics*, 2(2), 2015.
- [174] M. O. Williams, M. S. Hemati, S. T. M. Dawson, I. G. Kevrekidis, and C. W. Rowley. Extending data-driven Koopman analysis to actuated systems. In *IFAC Symposium on Nonlinear Control Systems*, 2016.
- [175] C. H. Williamson. Vortex dynamics in the cylinder wake. *Annual Review of Fluid Mechanics*, 28(1):477–539, 1996.
- [176] A. Wynn, D. S. Pearson, B. Ganapathisubramani, and P. J. Goulart. Optimal mode decomposition for unsteady flows. *Journal of Fluid Mechanics*, 733:473–503, 2013.
- [177] S. Yarusevych, P. E. Sullivan, and J. G. Kawall. Coherent structures in an airfoil boundary layer and wake at low Reynolds numbers. *Physics of Fluids*, 18(4):044101, 2006.
- [178] S. Yarusevych, P. E. Sullivan, and J. G. Kawall. On vortex shedding from an airfoil in low-Reynolds-number flows. *Journal of Fluid Mechanics*, 632:245–271, 2009.
- [179] J. Young and J. C. Lai. Vortex lock-in phenomenon in the wake of a plunging airfoil. *AIAA journal*, 45(2):485–490, 2007.
- [180] M. Y. Zakaria, H. E. Taha, M. R. Hajj, and A. A. Hussein. Experimental-based unified unsteady nonlinear aerodynamic modeling for two-dimensional airfoils. In *33rd AIAA Applied Aerodynamics Conference*, page 3167, 2015.
- [181] Z. Zhao and A. Singer. Fourier–Bessel rotational invariant eigenimages. *JOSA A*, 30(5):871–877, 2013.