

©Copyright 2022

Isabel Scherl

Optimization, Modeling, and Control of Cross-Flow Turbine Arrays

Isabel Scherl

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2022

Reading Committee:

Steven Brunton, Chair

Brian Polagye, Chair

Michelle DiBenedetto

Program Authorized to Offer Degree:
Mechanical Engineering

University of Washington

Abstract

Optimization, Modeling, and Control of
Cross-Flow Turbine Arrays

Isabel Scherl

Co-Chairs of the Supervisory Committee:

Professor Steven Brunton
Mechanical Engineering

Professor Brian Polagye
Mechanical Engineering

The ability to understand unsteady fluid flows is foundational to advancing technologies in energy, health, transportation, and defense. This work uses data-driven methods (i.e., machine learning) to interpret and control unsteady fluid flows through experiments. Specifically, these methods are used to control, optimize, and model cross-flow turbines. Cross-flow turbines (i.e. vertical axis turbines), are devices that can be used to convert the kinetic energy in wind to electricity. A key advantage of cross-flow turbines over axial-flow turbines is that they can efficiently operate in close-proximity in arrays. We demonstrate how data-driven methods can be used to efficiently explore, model, and interpret the high-dimensional space cross-flow turbine dynamics occupy through the following three projects.

First, robust principal component analysis (RPCA), a method borrowed from robust statistics, is used to improve flow-field data by leveraging global coherent structures to identify and replace spurious data points. We apply RPCA filtering to a range of fluid simulations and experiments of varying complexities and assess the accuracy of low-rank structure recovery. First, we analyze direct numerical simulations of flow past a circular cylinder at Reynolds number 100 with artificial outliers, alongside similar particle image velocimetry

(PIV) measurements at Reynolds number 413. Next, we apply RPCA filtering to a turbulent channel flow simulation from the Johns Hopkins Turbulence database, demonstrating that dominant coherent structures are preserved in the low-rank matrix. Finally, we investigate PIV measurements behind a two-bladed cross-flow turbine that exhibits both broadband and coherent phenomena. We demonstrate that more persistent dynamics can be identified when RPCA is utilized in lieu of traditional processing methods. In all cases, both simulated and experimental, we find that RPCA filtering extracts dominant coherent structures and identifies and fills in incorrect or missing measurements.

Second, the performance of a two-turbine array in a recirculating water channel was experimentally optimized across 64 unique array configurations using a hardware-in-the-loop approach. For each configuration, turbine performance was optimized using tip-speed ratio control, where the rotation rate for each turbine is optimized individually, and using coordinated control, where the turbines are optimized to operate at synchronous rotation rates but with a phase difference. For each configuration and control strategy, the consequences of co- and counter-rotation were also evaluated. Arrays with well-considered geometries and control strategies are found to outperform isolated turbines by up to 30%.

Third, the performance and wake of a two-turbine array in a fence configuration (side-by-side) are characterized. The turbines are operated under coordinated control. Measurements were made with turbines co-rotating, counter-rotating with the blades advancing upstream at the array midline, and counter-rotating with the blades retreating downstream at the array midline. From the performance and wake data, we found individual turbine and array efficiency to depend significantly on rotation direction and phase difference. Persistent dynamics that exist across all flow fields, as well as differences between cases are identified. Each of these projects demonstrate how data-driven methods can be used to explore, model, and interpret cross-flow turbine dynamics and other fluid systems.

TABLE OF CONTENTS

	Page
List of Figures	iii
Chapter 1: Introduction	1
1.1 Cross-Flow Turbines	1
1.2 Cross-Flow Turbine Arrays	2
1.3 Data-Driven Methods	5
1.4 Contributions	6
Chapter 2: Robust Principal Component Analysis for Modal Decomposition of Corrupt Fluid Flows	8
2.1 Experimental challenges	8
2.1.1 Contributions of this chapter	10
2.2 Modal analysis	11
2.2.1 Proper orthogonal decomposition (POD)	12
2.2.2 Dynamic mode decomposition	13
2.3 Robust extraction of fluid coherent structures	17
2.4 Model flows	20
2.4.1 Cylinder flow	20
2.4.2 Turbulent channel flow	21
2.4.3 Cross-flow turbine wake	22
2.5 Results	22
2.5.1 Cylinder flow	22
2.5.2 Turbulent channel flow	30
2.5.3 Cross-flow turbine wake	32
2.6 Conclusions and discussion	34

Chapter 3: Geometric and Control Optimization of a Two Turbine Array	38
3.1 Contributions of this chapter	38
3.2 Methods	39
3.2.1 Experimental design	41
3.2.2 Performance Metrics	41
3.2.3 Turbine Control	44
3.3 Results and Discussion	45
3.3.1 Co-rotation	45
3.3.2 Counter-rotation	47
3.3.3 Counter- vs. co-rotating arrays	49
3.3.4 Tip-speed ratio control vs. coordinated control	50
3.3.5 Inter-Array and Array-Channel Interactions	51
3.4 Conclusions	54
Chapter 4: Near Wake Dynamics of a Two Turbine Array	56
4.1 Methods	56
4.1.1 Cross-flow turbine performance and dynamics	57
4.1.2 Array Configurations, Control, & Symmetry	58
4.1.3 Experimental methods	59
4.1.4 PIV Methods	59
4.1.5 Array Performance	60
4.2 Results and Discussion	61
4.2.1 Array Performance Results	61
4.2.2 Time-average flow field evolution	63
4.3 Conclusions	68
Chapter 5: Conclusion	69
5.1 Future work	70
Bibliography	89
Appendix A: Appendix	90
A.1 Appendix: Optimized parameters	90
A.2 Appendix: Uncertainty analysis	90

LIST OF FIGURES

Figure Number	Page	
1.1	Cross-flow turbines operating side-by-side in the Alice C. Tyler Flume at the University of Washington.	3
2.1	Schematic of RPCA filtering applied to corrupt flow field data. Corrupted snapshots are arranged as column vectors in the matrix \mathbf{X} , which is decomposed into the sum of a low-rank matrix \mathbf{L} and a sparse matrix of outliers \mathbf{S} . (Videos: tinyurl.com/RPCA-PIV)	9
2.2	Flowchart showing how we apply RPCA filtering to a data matrix and analyze the results. Depending on the dataset in question, the data matrix may be artificially corrupted prior to RPCA filtering (Sec. 2.4). The results of principal component analysis and dynamic mode decomposition performed on the data matrix (\mathbf{X}) are referred to as PCA and DMD modes, respectively, whereas those same operations performed on the low rank matrix (\mathbf{L}) are referred to as RPCA and RDMD modes.	11
2.3	(left) Example flow field data. (right) Singular value spectrum for each data set. Mean flow travels from left to right in all cases.	20
2.4	RPCA filtering removes noise and outliers in the flow past a cylinder (black circle), from DNS (left) with 10% of velocity field measurements corrupted with salt and pepper noise, and PIV measurements (right). All frames show resultant vorticity fields. As the parameter λ is decreased, RPCA filtering is more aggressive, eventually incorrectly identifying coherent flow structures as <i>outliers</i>	23
2.5	RPCA filtering removes vorticity-biased corruption from simulated flow past a cylinder at Reynolds number 100. Unlike results shown in Fig. 2.4, corrupt entries are concentrated in regions of high vorticity instead of being uniformly distributed. In the flow on the left, $\eta = 1\%$ of the velocity field measurements are corrupted and on the right $\eta = 10\%$ of the velocity field measurements are corrupted. All frames show resultant vorticity fields.	25
2.6	Error ($\ \mathbf{X}_{uncorrupted} - \mathbf{L}\ _F / \ \mathbf{X}_{uncorrupted}\ _F$) and relative nuclear norm ($\ \mathbf{L}\ _* / \ \mathbf{X}\ _* = \mathbf{sum}(\boldsymbol{\sigma}_L) / \mathbf{sum}(\boldsymbol{\sigma}_{\mathbf{X},uncorrupted})$) of the low rank matrix \mathbf{L} compared with the uncorrupted data \mathbf{X} for varying percentages of corruption.	25

2.7	Odd PCA vorticity modes of the cylinder simulations from Fig. 2.5 with 1% of velocity measurements corrupted with a bias towards regions of high vorticity.	26
2.8	Odd PCA vorticity modes of the cylinder simulations from Fig. 2.5 with 10% of velocity measurements corrupted with a bias towards regions of high vorticity.	27
2.9	L_2 error between the true PCA modes of the clean cylinder simulation data \mathbf{X} and the RPCA and PCA modes for corrupted data with 1% (left) and 10% (right) vorticity-biased corruption.	28
2.10	Odd PCA vorticity modes of the experimental cylinder data for PCA and RPCA at $\lambda = 0.1, 1,$ and $10,$ along with their singular values.	28
2.11	Discrete-time DMD eigenvalues for the simulated cylinder data for small and large amounts of corruption and for increasing amounts of training data. In all cases, the RPCA-filtered DMD results dramatically outperform the standard DMD results.	29
2.12	Continuous-time DMD eigenvalues for the simulated cylinder data along with parabolic eigenvalue fits to estimate the error as in [13]. Here we use 5 vortex shedding periods with $\eta = 1\%$ corrupt values. The RDMD parabolic coefficient is approximately 2×10^4 times smaller than the DMD coefficient.	31
2.13	Continuous-time DMD eigenvalues for the PIV cylinder data along with parabolic eigenvalue fits to estimate the error as in [13]. The RDMD parabolic coefficient is approximately two times smaller than the DMD coefficient.	31
2.14	RPCA filtering for turbulent channel flow vorticity fields with various levels of added noise and tuning parameter λ . η represents the percentage of corrupted measurements in the velocity fields. The border colors match the color of the curve at the corresponding η in Fig. 2.15.	33
2.15	Turbulent kinetic energy (TKE) spectra for various levels of corruption and RPCA filtering. The TKE profiles provide a summary of the filtering that occurs at various scales. As corruption increases, the filtering remove more high-frequency information.	34
2.16	RPCA filtering of cross-flow turbine wake PIV data. The standard PIV processing pipeline (top row) includes several steps where RPCA filtering can be applied (bottom row). In the cross-correlated streamwise velocity field (top left), 23% of the velocity vectors are missing. Vector validation reduces the missing vectors to 20%. Finally, linear interpolation is used to fill in these missing vectors. In all cases, RPCA filtering captures the relevant phase-averaged coherent structures with fewer outliers and missing data, which appear as dark spots in the standard deviation plot.	35

2.17	Continuous-time DMD eigenvalues for the turbine wake PIV data, along with parabolic eigenvalue fits to estimate the error as in [13]. The RDMD parabolic coefficient is approximately six times smaller than the DMD coefficient. . . .	35
3.1	(a) Experimental set-up of the mobile, cantilevered turbine which consisted of a motor to enforce consistent tip-speed ratios and phase differences, a torque cell, an air bearing to absorb cantilever loads, and the rotor. (b) Experimental set-up of the fixed turbine, which consists of two load cells, a motor, the vacuum plate, and the rotor. (c) Photo of both turbines operating in the Bamfield Marine Science Centre flume.	39
3.2	(a) Experimental matrix showing the position where each control strategy was tested. (b) Each array configuration tested, where one turbine was fixed at $X/D = Y/D = 0$ and the other turbine was tested at each prescribed polar grid location. (c) Performance curve with the interquartile range for each turbine operating in isolation.	42
3.3	Optimization for one geometric configuration of co-rotating turbines under tip-speed ratio control contrasting initial case (light blue) with optimized case (red). From left to right: (a) the co-rotating turbines and direction of free-stream flow; (b) performance comparison of initial and optimized cases for each turbine as a function of the angular position of each turbine ($\theta = 0$ when the blade is pointing directly upstream) color encoded to denote interaction factor; (c) top shows the optimization path and bottom shows the performance evolution during optimization color encoded to denote interaction factor; (d) resulting performance of selected array geometry on a heat map corresponding to all test positions.	42
3.4	Interaction factor for co-rotating arrays under (a) tip-speed ratio control and (b) coordinated control. The fixed turbine is located at $Y/D = X/D = 0$ and the black dot denotes the location of the upstream turbine. Array performance greater than 1 (red) indicates that the array is outperforming the turbines in isolation. The color map contains a dead-band surrounding unity to highlight variations and interaction factor uncertainty of ± 0.039	46
3.5	Optimized parameters for co-rotating turbines under tip-speed ratio control. (a) Upstream turbine tip-speed ratio for a given array layout and (b) downstream turbine tip-speed ratio.	47
3.6	Optimized parameters for co-rotating turbines under coordinated control: (a) optimized tip-speed ratio for the turbine pair and (b) optimized phase difference.	48
3.7	Interaction factor for counter-rotating arrays under (a) tip-speed ratio control and (b) coordinated control.	48

3.8	Interaction factor comparison ($\Delta\kappa = \kappa_{\text{co-}} - \kappa_{\text{counter-}}$) of co-rotating versus counter rotating arrays for (a) coordinated control and (b) tip-speed ratio control. Green denotes a co-rotating array out-performing its counter-rotating counterpart.	49
3.9	(a) Mean wake velocity of a cross-flow turbine rotating counterclockwise with the skewed velocity deficit denoted with a red arrow [140], (b) deficit path and wake deflection direction for the co-rotating turbine array (c) deficit path and wake deflection direction for the counter-rotating turbine array.	50
3.10	Interaction factor comparison ($\Delta\kappa = \kappa_{\text{TSRC}} - \kappa_{\text{CC}}$) of tip-speed ratio control (TSRC) minus coordinated control (CC) for (a) counter-rotating and (b) co-rotating arrays.	51
3.11	Coefficient of performance (b) and coefficient of thrust (c) for the downstream turbine in three side-by-side array configurations (a) for coordinated control optimization. The point clouds correspond to variations in λ and ϕ during the control optimization. C_P , C_T , and λ are corrected according to the method of Barnsley and Wellicome [14].	51
4.1	Diagram of turbine showing the advancing and retreating directions within a rotation.	57
4.2	(a) Rotation schemes, co-, retreating, and advancing shown with phase difference calculations. Turbine 1 is located at a Y/D position of -0.55 and turbine 2 is located at a Y/D position of 0.55. For all cases, turbine 1 is at 0° and the phase difference shown is 45° , (b) turbine positions and PIV fields shown in the flume coordinates.	58
4.3	$C_P - \phi$ plot for the retreating case at $\lambda = 2.1 - 2.6$	61
4.4	(a) Array average performance for all rotation schemes (b) Mean-subtracted average performance for the retreating rotation scheme (c) Mean-subtracted performance for the advancing rotation scheme (d) Mean-subtracted performance for the co-rotation scheme.	63
4.5	(a-j) Average flow field magnitude for each rotation scheme and phase difference measured (k-n) Average flow field magnitude contours for a 90° phase difference for each rotation scheme (k), all co-rotating phase differences (l), all retreating counter-rotating phase differences (m), and all advancing counter-rotating phase differences (n); for each of the profiles, the contour represents the velocity magnitude present at $X/D = 1, 2, \text{ or } 3$. The curves are located such that the dotted lines are at $\frac{ U }{U_\infty} = 1$ and the spacing between the dotted lines (1 X/D) represents $\frac{ U }{U_\infty} = 1$	64

4.6	Selected phase averaged frames for the advancing and retreating counter-rotating cases at turbine 1 azimuthal positions of $\theta_1 \approx 0^\circ, 42^\circ, \text{ or } 85^\circ$	66
4.7	Second POD mode for the advancing and retreating counter-rotating cases at each phase difference.	67
A.1	Optimized parameters for counter-rotating turbines under tip-speed ratio control. (a) Upstream turbine tip-speed ratio for a given array layout and (b) downstream turbine tip-speed ratio for the same.	91
A.2	Optimized parameters for counter-rotating turbines under coordinated control: (a) optimized tip-speed ratio for the turbine pair and (b) optimized phase difference.	91
A.3	Uncertainty values ($\delta\kappa'$) calculated from Eq. 11 for both rotation directions and control strategies.	92

ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisors Brian Polagye and Steve Brunton. Throughout my PhD, Brian has demonstrated sincere trust in and commitment to his students and an unwavering pragmatism I hope to emulate. Steve has imparted persistent optimism and encouraged me to continually reframe my perspective both in research and my career on the whole. You have both guided my growth as a scientist and person and allowed me to pursue projects that pique my interest. For that, I am grateful.

In both the Marine Renewable Energy Lab and Steve Brunton's lab, I have had incredibly supportive labmates and friends. In particular I would like to mention Ama Hartman, Brittany Lydon, Emma Cotter, Gemma Calandra, Lindsey Jones, Kate van Ness, Trent Dillon, Trevor Harrison, Ariana Mendible, Benjamin Hermann, Krithika Manohar, Sam Rudy, Jared Callahan, Kardi Kaheman, Alan Kaptanoglu, Urban Faisel, and Michelle Hickner. I would *especially* like to thank the cross-flow turbine group, Aidan Hunt, Hannah Ross, Abigale Snortland, Carl Stringer, Benjamin Strom, and Greg Talpey, for graciously supporting one another in debugging or data interpretation. Learning and building community with all of you over the last 5 years has been an absolute joy.

Thank you to both the MREL discussion group and MESAR for bringing the imperative topics of equity and justice to the forefront and working to make our lab and department a more welcoming place for future students.

To my dear friends, Ariana for being a trusted confidant and agreeing to my impulsive or impromptu adventures, Ama and Julie for helping me work through my thoughts on long runs and beyond, and to Miranda, my faithful last minute proofreader, your support and friendship from afar have been incomparable, thank you. To the other friends I have made

along the way, the picnics, game nights, Lopez 4ths, backpacking trips, bike rides long and short, evening volleyball games, shared meals, and so many other adventures have filled me with such joy and gratitude.

To Ben, whose care and support has been steadfast over the last 5 years and for always being there if I want to talk through a problem, help me come up with a solution, and plan adventures when that is what we both need. Thank you to my family, both old and new, for supporting me through this journey including my parents, Sophie, Zeke, Stephanie, Jennifer, Eric, Suzanne, Mike, Laura, and Duna.

Thank you to my committee members, past and present, including Brian Polagye, Steve Brunton, Andy Stewart, Junlan Wang, Jim Riley, Michelle DiBenedetto, and Tom Daniel, for gracious service and perspective through snowmageddon, a pandemic, and whatever is to come on August 11th.

Chapter 1

INTRODUCTION

1.1 Cross-Flow Turbines

Cross-flow turbines (i.e., vertical-axis turbines), which use blades that rotate about an axis perpendicular to the incoming flow, are an innovative option for harnessing the renewable energy from marine currents and wind [52, 126]. Currently, wind energy comprises more than 8% of electricity generation in the U.S. and more than 20% in 10 states [157]. Though wind energy is predominantly produced by axial-flow turbines, cross-flow turbines are a promising alternative technology. Additionally, there has been significant interest in similarly converting the kinetic energy in marine currents and riverine flows, particularly in remote or isolated settings [100].

Advantages of cross-flow turbines over axial-flow turbines (i.e., horizontal-axis in wind) include bi- or omni-directional operation without active yaw control and less environmental disturbance due to lower maximum blade speeds for equivalent inflow conditions. Since large-scale energy harvesting relies on turbine arrays, the most significant advantage is that cross-flow turbines can interact beneficially, producing more power together than in isolation and outperform an equivalent axial-flow turbine array [155, 83]. In fact, larger arrays are estimated to have a power density 6-9 times greater than axial-flow turbine arrays [38]. The mechanisms responsible for these improvements are not yet well understood and are potentially augmented when turbines are operating in confined flow [121].

While the performance of a cross-flow turbine depends on its geometry (e.g., number of blades [32], chord-to-radius ratio [107], blade profile [31], and preset pitch angle [33]) and operating condition (e.g., channel blockage [120], Reynolds number [10]), all fixed-pitch cross-flow turbines experience similar unsteady fluid-structure interaction. Specifically, as a cross-

flow turbine blade rotates, the angles of attack of the blades are constantly changing, resulting in unsteady boundary separation and dynamic stall, an active area of research for which a universal model does not yet exist [37]. Following stall separation, periodic structures are advected into the wake [45, 39]. Cross-flow turbine wake dynamics are broadly characterized as unsteady separated flows with significant three dimensionality [124]. Consequently, cross-flow turbines exhibit high-dimensional, complex dynamics [107, 9, 11, 113] that are not present in traditional axial-flow turbines, but have commonalities with energy harvesters that utilize heaving and pitching foils [117]. Because of this, the flow fields around individual cross-flow turbines have received significant attention both computationally [17, 102, 130, 131, 134] and experimentally [6, 73, 119, 124, 142]. These studies show that an individual rotor wake has an asymmetric skew towards the direction where the blades are advancing upstream. The deficit, in conjunction with accelerated bypass flow around the rotor, produce a shear layer. The region of the wake with the maximum deficit is not immediately downstream of the rotor. Instead, its distance from the rotor varies with tip-speed ratio and the number of blades on the turbine.

While variable blade pitch mechanisms [125] and boundary-layer flow control have been demonstrated for cross-flow turbines, active control in most systems is limited to the rotor angular velocity. Performance is similar when turbines are operated at a constant angular velocity or a constant regulatory torque [109], though increased power production is possible when cross-flow turbines are forced to follow a non-constant rotation rate [139, 40].

1.2 Cross-Flow Turbine Arrays

Large-scale energy harvesting by cross-flow turbines typically requires arrays of multiple turbines. When expanding from a single turbine to a multi-turbine array, the system becomes increasingly high dimensional, turbulent, and nonlinear. A pair of closely-spaced cross-flow turbines in a laboratory setting is shown in Figure 1.1. The flow induced by each turbine depends heavily on the relative position of turbines in the array, rotation rates relative to the incoming flow, the direction of rotation for each turbine, and any phase difference




Figure 1.1: Cross-flow turbines operating side-by-side in the Alice C. Tyler Flume at the University of Washington.

between the turbines. These factors increase the dimensionality of the problem relative to an individual turbine due to the potential for multiple interaction mechanisms between turbines in the array. For example, mutual interactions can occur between turbines in relatively close proximity, either through mean flow alteration or coherent structure interaction, as well as downstream wake interaction at greater separation distances after coherent structures have dissipated.

Dense array interactions between nearby turbines have analogues to flow control in other fluid systems (e.g., fish schooling [56, 159, 151], flocking birds [110, 159], heaving and pitching foils for energy generation [99, 117]). Prior investigations of cross-flow turbine arrays in field experiments [38, 83, 82], laboratory experiments [20, 21, 5, 150, 43, 128], and simulation [47, 162, 19, 49, 35] have demonstrated that power output increases when the rotors are arranged in a side-by-side configuration. Further, as the rotor spacing decreases,

efficiency increases [162, 79].

Since large-scale energy harvesting typically involves multiple rows of turbines, the wake evolution is also of considerable interest, but wake studies for multiple cross-flow turbines are less common. Zanforlin et al., 2016 [162] simulated pairs of counter-rotating turbines with varied rotor spacing, non-dimensional rotation rate (tip-speed ratio), and wind direction. They note that turbines placed side-by-side induced beneficial mean flow alterations in the cross-stream direction increasing lift on the blade and turbine torque. In the wake, they found that the individual deficit regions contract downstream of the rotors. In Posa et al., 2019 [111], pairs of turbines were simulated co-rotating and counter-rotating. This work was then extended this work to include blade-tip effects and wake recovery, finding that three dimensional edge effects accelerate wake recovery [112]. In both of these computational studies, the turbines are operated in a synchronous manner without any phase offset (i.e., blades on both turbines at equal positions). As an exception, Jin et al., 2020 [79] includes a computational study on rotation scheme and rotor phase difference. They found that the main driver of wake dynamics is turbine rotation scheme and that effects from phase difference are secondary. Performance and wakes of turbine pairs offset in the streamwise direction have been explored [20]. These offset arrays were found to have increased efficiency relative to the individual turbines and the time-averaged flow fields showed varying dynamics dependent on rotation scheme. Lam et al., 2017 [86] experimentally measured asymmetric wake skew patterns in paired turbines that are consistent with those seen in individual rotor wakes.

Consequently, our understanding of rotor interactions and array dynamics, such as the the effect of turbine phase and rotor rotation schemes on wake dynamics and performance, remains limited. Additionally, there has been little exploration of phase-resolved wake dynamics. On the whole, we are limited by both the measurable quantities in our experiments and the inherent difficulty in disambiguating flow structures and their effects on turbine dynamics and control. Determining causality between performance and flow field alterations, not just correlation, remains a central challenge in understanding both individual turbines

and arrays.

1.3 Data-Driven Methods

The ability to understand, model, and control fluid flows is foundational to advancing technologies not only in energy, but also in health, transportation, and defense. Specifically, cross-flow turbine dynamics are not easily solved by first principles analysis without intense simplification. Instead, we rely on data from simulations and experiments [144, 46, 25]. However, challenges are introduced because system state estimation data is limited by the measurable and controllable parameters. Additionally, data quality is dependant on the precision and accuracy of our measurement or computational techniques.

Despite measurement limitations, we are collecting increasing amounts of data in these simulations and experiments. Analyzing vast amounts of data can be challenging, but recently, data-driven methods and machine learning have been shown to be useful in better interpreting fluid systems [16, 145]. These techniques have been used to model, optimize, and control fluid flows [48, 85, 23]. Modal decomposition techniques are designed to extract these meaningful patterns from high-dimensional fluids data [144, 145], resulting in a compact representation that can be used for accurate and efficient reduced-order models and control [23, 123].

The majority of modal decompositions are linear [144, 145], although emerging techniques in machine learning are providing improved nonlinear pattern extraction [25]. Linear regression and least-squares optimization are particularly widely used, as in the ubiquitous proper orthogonal decomposition (POD) (also commonly known as principal component analysis (PCA)), [93, 8, 16, 74, 146] and the dynamic mode decomposition (DMD) [132, 122, 148, 84]. POD provides a principled approach to decomposing high-dimensional fluid flow data into a hierarchy of orthogonal modes that are ordered in terms of their ability to capture the energy in the flow; because these modes are orthogonal, it is possible to obtain reduced-order models by Galerkin projection of the Navier-Stokes equations onto a truncated POD basis [103, 30, 90]. Robust principal component analysis (RPCA) is a robust variant of prin-

principal component analysis, that decomposes a data matrix into the sum of a low-rank matrix containing coherent structures and a sparse matrix of outliers and corrupt entries [29]. DMD is a technique related to POD to decompose a flow into spatiotemporal coherent structures that are each constrained to have coherent and linear dynamics in time.

In addition to methods that utilize data produced from simulations and experiments, there is a variety of techniques that can improve data acquisition. Data is commonly collected in uniform grids across parameters. Since the dimensionality of a parameter space increases exponentially with the parameters, traditional uniform sampling across parameters quickly becomes intractable and it becomes necessary to explore alternative acquisition methods. Optimization schemes [101] have been used to explore similarly large parameter spaces [128] and, while an excellent solution to find optimality, they cannot provide full information about how various parameters affect performance.

1.4 Contributions

Here we explore dynamics and apply data-driven methods to cross-flow turbines and other fluid systems. Contributions of this thesis are in three parts. First, robust principal component analysis (RPCA) is used to improve the quality of flow-field data by leveraging global coherent structures to identify and replace spurious data points. RPCA filtering is applied to a range of fluid simulations and experiments of varying complexities and the accuracy of low-rank structure recovery is assessed. This work has been published in *Physical Review Fluids* and was selected as the editors' selection [129]. Second, the performance of a two-turbine array in a recirculating water channel is experimentally optimized across 64 unique array configurations. For each configuration, turbine performance is optimized using tip-speed ratio control, where the rotation rate for each turbine is optimized individually, and using coordinated control, where the turbines are optimized to operate at synchronous rotation rates but with a phase difference. For each configuration and control strategy, the consequences of individual turbine rotation direction were also evaluated. This work has been published in the *Journal of Renewable and Sustainable Energy* [128]. Third, the per-

formance and wake of a two-turbine array in a fence configuration (side-by-side turbines) are characterized under coordinated control. Measurements were made with turbines co-rotating, counter-rotating with the blades advancing upstream at the array midline, and counter-rotating with the blades retreating downstream at the array midline. As with prior simulation results [79], we find that individual turbine and array efficiency depends primarily on rotation scheme and secondarily on phase difference. Similar variations are observed in the wake and, using these data, we hypothesize how rotation direction and phase influence interactions between adjacent turbines. This work is being prepared for publication.

Chapter 2

ROBUST PRINCIPAL COMPONENT ANALYSIS FOR MODAL DECOMPOSITION OF CORRUPT FLUID FLOWS

In this chapter, we explore the use of the robust principal component analysis (RPCA) [29] to process corrupt flow fields, leveraging global correlations in the data. We emphasize the impact of this approach on modal analysis, including POD/PCA and DMD.

2.1 *Experimental challenges*

Experimental techniques to measure fluid flows have evolved rapidly over the past century, with the ultimate goal of acquiring full flow fields with high spatial and temporal resolution. Laser-based imaging techniques have evolved from point measurements [160, 58] to 2D and 3D field measurements [156, 1, 2, 3, 153, 115]. Particle Image Velocimetry (PIV) has since become a cornerstone of experimental fluid mechanics, providing non-intrusive velocity field measurements across a range of applications. Improvements in PIV hardware, including more powerful lasers, higher resolution and frame rate cameras, advanced image processing technology, and the development of tomographic PIV are providing unprecedented views into real flows. Despite the undeniable success of PIV, there are several well-known challenges to acquiring clean and accurate data. Multiple factors in the PIV data acquisition and processing pipeline can contribute to velocity vector outliers that degrade the resulting velocity fields. These include inadequate illumination and irregularities in the light field, background speckle, seeding density and non-passivity of the particles, sharp gradients in flow properties, optical issues, such as alignment and aberration, limited resolution and shot noise in the image recording, and out of plane motion of the particles when measuring in 2D [76, 3]. Because of a fundamental tradeoff between the quantity and quality of PIV data,




Figure 2.1: Schematic of RPCA filtering applied to corrupt flow field data. Corrupted snapshots are arranged as column vectors in the matrix \mathbf{X} , which is decomposed into the sum of a low-rank matrix \mathbf{L} and a sparse matrix of outliers \mathbf{S} . (Videos: tinyurl.com/RPCA-PIV)

in both space and time, researchers continue to push the resolution limits of current systems. Thus, flow fields acquired with PIV often have missing and/or corrupt measurements. This has motivated processing techniques to improve PIV data quality of PIV data [71, 106, 61], including predictor-corrector schemes [133], spatial filtering to remove frequencies not possible for the measurement resolution [42], and POD-based background removal [95]. The identification of spurious vectors has been studied extensively [152, 44], and the normalized median filter is a robust and well-used method [154]. However, missing vectors often cluster in regions of high shear, presenting a challenge for standard vector validation and interpolation methods that rely on local flow information [154]. In this chapter, we leverage robust statistics and global spatiotemporal coherent structures across the entire dataset to fill in missing measurements and improve modal decomposition of fluid flow fields. We also utilize PIV to explore the wake dynamics of a two-turbine array in Chapter 4.

2.1.1 Contributions of this chapter

We investigate the use of RPCA [29], a robust variant of POD/PCA [108, 22]. RPCA uses a sparsity-promoting optimization to decompose a data matrix into the sum of a low-rank matrix containing coherent structures and a sparse matrix containing outliers. RPCA was originally popularized in the Netflix matrix completion algorithm for its recommender system [158] and has since been widely used for image and video processing [18], electrical capacitance tomography [88], and voice separation [77]. Here, we use RPCA filtering to process flow field data from several simulations and experiments. Figure 2.1 demonstrates the ability of RPCA to uncover and isolate the dominant low-rank coherent structures from sparse outliers in flow data from an idealized example. In addition to directly analyzing and processing flow field data, we also perform PCA and DMD modal analyses on the data before and after RPCA filtering to assess its performance.

Here we consider a range of simulated and experimentally acquired flow fields of varying complexity to isolate and analyze various aspects of the algorithm applied to data from fluid mechanics. First, we investigate high-fidelity flow fields from direct numerical simulations of a laminar flow past a cylinder and a turbulent channel flow, where it is possible to artificially add corrupt velocity field vectors to compare the RPCA filtered fields with a known ground truth. Next, we apply the method to two experimentally acquired datasets, including a companion laminar flow past a cylinder and measurements of a cross-flow turbine wake. Although there is not a ground truth model for these flows, it is possible to estimate the effect of RPCA filtering in reducing outliers and corruption by analyzing the DMD spectrum, which has well-stereotyped behavior for such periodic wake flows [13]. In all cases, we show that the RPCA filtered fields yield DMD spectra that are more consistent with a periodic wake in the absence of noise.

This chapter is organized as follows: First, we present the standard POD/PCA and DMD modal analysis techniques in Sec. 2.2, followed by the RPCA method in Sec. 2.3. Section 2.4 describes the four flow fields used in this analysis. Results of RPCA filtering on these flow

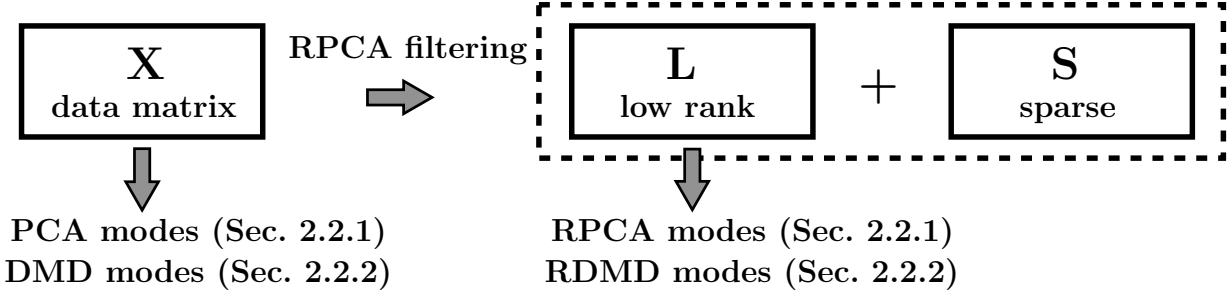


Figure 2.2: Flowchart showing how we apply RPCA filtering to a data matrix and analyze the results. Depending on the dataset in question, the data matrix may be artificially corrupted prior to RPCA filtering (Sec. 2.4). The results of principal component analysis and dynamic mode decomposition performed on the data matrix (\mathbf{X}) are referred to as PCA and DMD modes, respectively, whereas those same operations performed on the low rank matrix (\mathbf{L}) are referred to as RPCA and RDMD modes.

fields and its impact on downstream modal analysis are presented in Sec. 2.5.

2.2 Modal analysis

Extracting coherent structures from high-dimensional data has been a central challenge in fluid mechanics for decades. Here we review two leading modal decomposition techniques for data from fluid mechanics, the proper orthogonal decomposition (POD), also known as principal component analysis (PCA) (Sec. 2.2.1), and the dynamic mode decomposition (DMD) (Sec. 2.2.2). Both methods apply equally well to data from simulations or experiments. We use these two modal decomposition techniques to assess the effectiveness of RPCA filtering, in processing and correcting corrupt flow fields. Both techniques are based on the singular value decomposition (SVD) [66, 64, 28, 65] and there are several detailed discussions of these decompositions [84, 144, 145, 22]. The RPCA algorithm is explained in Sec. 2.3.

In this chapter, we follow the flowchart shown in Fig. 2.2. RPCA filtering is applied to a data matrix (\mathbf{X}) which is bisected into the low rank structure (\mathbf{L}) and sparse (\mathbf{S}) subspaces. From there, RPCA modes and DMD modes are calculated from the low rank data. We also calculate the POD/PCA and DMD modes on the data matrix.

2.2.1 Proper orthogonal decomposition (POD)

Proper orthogonal decomposition—referred to as principal component analysis (PCA) throughout results—is a widely used method to identify spatially correlated coherent structures from data, decomposing the flow field into a linear combination of orthogonal modes that are arranged hierarchically by energy content. There are several variants of POD [144, 93, 136, 16, 74, 146], and we will present a variant of the *snapshot POD* of Sirovich that relies on the numerically stable SVD [136, 22]. First, flow field data (e.g., a velocity or vorticity field) is measured or computed on a discrete spatial grid, and m snapshots of these flow fields are collected at various times t_1, t_2, \dots, t_m . The flow field data at time t_k may be reshaped into a column vector $\mathbf{x}_k = \mathbf{x}(t_k) \in \mathbb{R}^n$, where n denotes the number of flow variables times the number of spatial grid locations. Next, a data matrix is formed by arranging the column vectors \mathbf{x}_k in a matrix \mathbf{X} :

$$\mathbf{X} = \begin{bmatrix} | & | & & | \\ \mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_m \\ | & | & & | \end{bmatrix}. \quad (2.1)$$

Finally, POD modes are obtained by computing the singular value decomposition of $\mathbf{X} \in \mathbb{R}^{n \times m}$:

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad (2.2)$$

where T defines the matrix transpose, $\mathbf{U} \in \mathbb{R}^{n \times n}$, $\mathbf{\Sigma} \in \mathbb{R}^{n \times m}$ and $\mathbf{V} \in \mathbb{R}^{m \times m}$. The columns of \mathbf{U} are *POD modes* with the same dimension as a flow field \mathbf{x} . POD modes are orthonormal so that $\mathbf{U}^T\mathbf{U} = \mathbf{I}$; similarly $\mathbf{V}^T\mathbf{V} = \mathbf{I}$. Moreover, the columns of \mathbf{U} (resp. rows of \mathbf{V}^T) are arranged in order of their importance in describing the data. The importance of each mode (i.e., column of \mathbf{U}) is given by the corresponding entry of the non-negative, diagonal matrix of singular values $\mathbf{\Sigma} \in \mathbb{R}^{n \times m}$.

The matrix \mathbf{X} will exhibit *low-rank structure*, so that it is well approximated by the first

$r \ll m < n$ columns of \mathbf{U} and \mathbf{V} :

$$\mathbf{X} \approx \mathbf{U}_r \boldsymbol{\Sigma}_r \mathbf{V}_r^T, \quad (2.3)$$

where \mathbf{U}_r and \mathbf{V}_r denote the first r columns of each matrix and $\boldsymbol{\Sigma}_r$ denotes the first $r \times r$ sub-block of $\boldsymbol{\Sigma}$. In fact, the Eckart-Young theorem states that this is the *optimal* rank- r approximation of the matrix \mathbf{X} in a least-squares sense. More details about the SVD can be found in [22].

After truncating all but the first r dominant modes, a flow field snapshot \mathbf{x} may be approximated by a linear combination of these modes:

$$\mathbf{x} \approx \sum_{k=1}^r \mathbf{u}_k \alpha_k,$$

where α_k is the *POD mode coefficient*.

Because the POD modal basis is orthogonal, it is possible to obtain a reduced-order nonlinear dynamical system for the evolution of the coefficients $\alpha_k(t)$ in time via Galerkin projection of the Navier-Stokes equations onto the POD basis. In this way, the POD basis may be thought of a data-driven generalization of the Fourier basis that is tailored to a particular flow field. POD is also closely related to principal component analysis (PCA) [108], the Karhunen–Loève decomposition [81], empirical orthogonal functions [92], or the Hotelling transform [75].

2.2.2 *Dynamic mode decomposition*

DMD is a modal decomposition technique that simultaneously identifies spatially coherent modes that are constrained to have the same linear behavior in time, given by oscillations at a fixed frequency with growth or decay [132, 84]. Thus, the dynamic mode decomposition provides a dimensionality reduction into a set of spatial modes along with a linear model for how these modes evolve in time. This is in contrast to POD, which results in orthogonal

modes arranged in terms of energy content and without consideration of dynamics. However, in many formulations, DMD is closely related to POD, and may be thought of as a linear combination of POD modes that results in linear evolution in time. DMD also has deep connections to nonlinear dynamical systems via Koopman operator theory [96, 122, 148, 97, 84].

In the original formulation of DMD [132], the snapshots in the data matrix in Eq. (2.1) are spaced evenly in time, so that $t_k = k\Delta t$ with Δt sufficiently small to resolve the highest frequencies in the dynamics. Generalizations have since been formulated to allow for non-sequential time-series [148, 7] and for data that is under-resolved in space [24, 70] or time [147]; however, for simplicity, we will present the standard *exact DMD* formulation of Tu et al. [148] with evenly spaced and sequential snapshots. DMD seeks to identify the leading eigenvalues and eigenvectors of the best-fit linear operator \mathbf{A} that evolves snapshots forward in time:

$$\mathbf{x}_{k+1} \approx \mathbf{A}\mathbf{x}_k. \quad (2.4)$$

The eigenvectors ϕ of \mathbf{A} have the dimensions of a flow-field and correspond to spatiotemporal coherent structures whose dynamics in time evolve according to the associated eigenvalue γ .

In practice, this operator is identified by first splitting the data in Eq. (2.1) into two matrices:

$$\mathbf{X} = \begin{bmatrix} | & | & & | \\ \mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_{m-1} \\ | & | & & | \end{bmatrix} \quad \mathbf{X}' = \begin{bmatrix} | & | & & | \\ \mathbf{x}_2 & \mathbf{x}_3 & \cdots & \mathbf{x}_m \\ | & | & & | \end{bmatrix}. \quad (2.5a)$$

and then solving for the best fit operator that satisfies

$$\mathbf{X}' \approx \mathbf{A}\mathbf{X} \quad (2.6)$$

via the following *least-squares* optimization problem

$$\mathbf{A} = \arg \min_{\mathbf{A}} \|\mathbf{X}' - \mathbf{A}\mathbf{X}\|_F = \mathbf{X}'\mathbf{X}^\dagger \approx \mathbf{X}'\mathbf{V}_r\boldsymbol{\Sigma}_r^{-1}\mathbf{U}_r^T. \quad (2.7)$$

Here we are minimizing the Frobenius norm $\|\cdot\|_F$ via the pseudo-inverse $\mathbf{X}^\dagger = \mathbf{V}\boldsymbol{\Sigma}^{-1}\mathbf{U}^T \approx \mathbf{V}_r\boldsymbol{\Sigma}_r^{-1}\mathbf{U}_r^T$.

In practice, the matrix \mathbf{A} is far too large to analyze directly, and instead, we project \mathbf{A} onto an r -dimensional POD subspace, given by the columns of \mathbf{U}_r :

$$\tilde{\mathbf{A}} = \mathbf{U}_r^T \mathbf{A} \mathbf{U}_r = \mathbf{U}_r^T \mathbf{X}' \mathbf{V}_r \boldsymbol{\Sigma}_r^{-1}. \quad (2.8)$$

The eigenvalues of $\tilde{\mathbf{A}}$ are the same as the eigenvalues of \mathbf{A} , which are known as the *DMD eigenvalues*. They are computed via the eigendecomposition of the $r \times r$ matrix $\tilde{\mathbf{A}}$:

$$\tilde{\mathbf{A}}\mathbf{W} = \mathbf{W}\boldsymbol{\Gamma}. \quad (2.9)$$

Finally, the corresponding *DMD modes* are reconstructed using the full-dimensional data along with the reduced eigenvectors in \mathbf{W} :

$$\boldsymbol{\Phi} = \mathbf{X}'\mathbf{V}_r\boldsymbol{\Sigma}_r^{-1}\mathbf{W}. \quad (2.10)$$

This formula for the eigenvectors is from the exact DMD algorithm [148, 84]; the original formulation of Schmid [132] computes modes as $\boldsymbol{\Phi} = \mathbf{U}_r\mathbf{W}$.

With the DMD modes $\boldsymbol{\Phi}$ and eigenvalues $\boldsymbol{\Gamma}$ it is possible to reconstruct the state at time $k\Delta t$

$$\mathbf{x}_k = \sum_{j=1}^r \phi_j \gamma^{k-1} b_j = \boldsymbol{\Phi}\boldsymbol{\Gamma}^{k-1}\mathbf{b}, \quad (2.11)$$

where the vector \mathbf{b} of mode amplitudes is generally computed as

$$\mathbf{b} = \Phi^\dagger \mathbf{x}_1. \quad (2.12)$$

More principled approaches to select the few dominant modes have been considered based on sparsity-promoting optimization [80].

The spectral expansion above may also be written in continuous-time by introducing the continuous eigenvalues $\omega = \log(\gamma)/\Delta t$:

$$\mathbf{x}(t) = \sum_{j=1}^r \phi_j e^{\omega_j t} b_j = \Phi \exp(\Omega t) \mathbf{b}, \quad (2.13)$$

where Ω is a diagonal matrix containing the continuous-time eigenvalues ω_j .

DMD is known to be extremely sensitive to noisy data [13, 41, 72], and the eigenvalues specifically suffer from a bias that is not reduced with increasing data. There are several modifications to make DMD more robust to noise, including averaging forward-time and backward-time operators [41], total least squares [72], and variable projection [7]. For periodic wake data, as explored in three of the examples in this paper, the discrete-time eigenvalues should occur in complex conjugate pairs $\gamma, \bar{\gamma}$ exactly on the unit circle in the complex plane for clean data. Similarly, the continuous-time eigenvalues should be in complex conjugate pairs $\pm i\omega$ on the imaginary axis, indicating pure oscillations with no growth or decay [12, 13]. In [13], Bagheri characterized the perturbative effect of noise on these eigenvalues, deriving an asymptotic expression for how high frequency eigenvalues become increasingly affected by noise. If the true continuous-time eigenvalue should be $\pm i\omega$ in the absence of noise, Bagheri showed that in the presence of perturbative white noise with magnitude $\epsilon \ll 1$ the observed eigenvalue pair is

$$\pm i\omega - \epsilon C \omega^2 + \mathcal{O}(\epsilon^2), \quad (2.14)$$

where C is a sensitivity constant. Thus, low noise levels cause a spurious real-valued damping $-\epsilon C\omega^2$ that is quadratic in the frequency. We will make extensive use of this property to assess the quality of our RPCA filtered fields by computing the ratio of the best-fit factor ϵC before and after applying RPCA filtering. ϵC should decrease as a consequence of reduced noise and corruption.

2.3 Robust extraction of fluid coherent structures

Techniques based on least-squares regression, such as POD/PCA and DMD, are highly susceptible to outliers and corrupted data, making them *fragile* with respect to some experimental measurement errors. Outliers and corruption are defined as data points that differ significantly from the statistical distribution of the majority of the dataset [69], so that they cannot be considered as the original data plus a small-to-moderate amount of white noise. To mitigate this sensitivity, Candes et al. [29] have developed a robust principal component analysis (RPCA) that seeks to decompose a data matrix \mathbf{X} into a structured low-rank matrix \mathbf{L} that is characterized by dominant coherent structures and a sparse matrix \mathbf{S} containing outliers and corrupt data:

$$\mathbf{X} = \mathbf{L} + \mathbf{S}. \quad (2.15)$$

The principal components of \mathbf{L} are *robust* to outliers and corrupt data, which are isolated in \mathbf{S} . This decomposition, also referred to as a filter, has profound implications for many modern problems of interest, including video surveillance (where the background objects appear in \mathbf{L} and foreground objects appear in \mathbf{S}), facial recognition (eigenfaces are in \mathbf{L} and shadows, occlusions, etc. are in \mathbf{S}), natural language processing and latent semantic indexing, and ranking problems¹. Standard PCA/POD is effective at removing white noise that is smaller than the relevant singular values in the data [63]; however, it is not able to

¹The ranking problem may be thought of in terms of the Netflix prize for matrix completion. In the Netflix prize, a large matrix of preferences is constructed, with rows corresponding to users and columns corresponding to movies. This matrix is sparse, as most users only rate a handful of movies. The Netflix prize seeks to accurately fill in the missing entries of the matrix, revealing the likely user rating for movies the user has not seen.

remove outliers. Instead, RPCA is used to correct outliers that differ significantly from the distribution of the other observations.

Mathematically, the goal is to find \mathbf{L} and \mathbf{S} that satisfy the following:

$$\min_{\mathbf{L}, \mathbf{S}} \text{rank}(\mathbf{L}) + \|\mathbf{S}\|_0 \text{ subject to } \mathbf{L} + \mathbf{S} = \mathbf{X}. \quad (2.16)$$

$\|\mathbf{S}\|_0$ counts the number of nonzero elements in \mathbf{S} , quantifying how *sparse* it is. $\text{rank}(\mathbf{L})$ is the number of nonzero singular values in \mathbf{L} , quantifying how many linearly independent rows and columns describe the data. However, neither the $\text{rank}(\mathbf{L})$ nor the $\|\mathbf{S}\|_0$ terms are convex, making this optimization intractable. Similar to the compressed sensing problem, it is possible to solve for \mathbf{L} and \mathbf{S} with *high probability* using a convex relaxation of (2.16):

$$\min_{\mathbf{L}, \mathbf{S}} \|\mathbf{L}\|_* + \lambda_0 \|\mathbf{S}\|_1 \text{ subject to } \mathbf{L} + \mathbf{S} = \mathbf{X}, \quad (2.17)$$

where $\|\cdot\|_*$ is the nuclear norm, given by the sum of singular values which is a proxy for the rank of the matrix, and $\|\cdot\|_1$ is the 1-norm of the matrix viewed as a vector, given by the sum of the magnitudes of each entry in the matrix, which is a proxy for the $\|\cdot\|_0$ norm of a matrix; the hyperparameter λ_0 is given by $\lambda_0 = \lambda/\sqrt{\max(n, m)}$. The solution to (2.17) converges to the solution of (2.16) with high probability if $\lambda = 1$, where n and m are the dimensions of \mathbf{X} , given that \mathbf{L} is not sparse and \mathbf{S} is not low-rank. In the examples below, these assumptions may only be partially valid, so the optimal value of λ may vary slightly. The convex problem in (2.17) is known as *principal component pursuit* (PCP), and may be solved using the augmented Lagrange multiplier (ALM) algorithm. Specifically, an augmented Lagrangian may be constructed:

$$\mathcal{L}(\mathbf{L}, \mathbf{S}, \mathbf{Y}) = \|\mathbf{L}\|_* + \lambda_0 \|\mathbf{S}\|_1 + \langle \mathbf{Y}, \mathbf{X} - \mathbf{L} - \mathbf{S} \rangle + \frac{v}{2} \|\mathbf{X} - \mathbf{L} - \mathbf{S}\|_F^2. \quad (2.18)$$

Where \mathbf{Y} is the matrix of Lagrange multipliers and v is a hyperparameter. We then solve

for \mathbf{L}_k and \mathbf{S}_k to minimize \mathcal{L} , update the Lagrange multipliers

$$\mathbf{Y}_{k+1} = \mathbf{Y}_k + \nu(\mathbf{X} - \mathbf{L}_k - \mathbf{S}_k),$$

and iterate until convergence. Here, an inexact ALM implementation from [138] is used. The alternating directions method (ADM) [89, 161] provides another simple procedure.

After the low-rank matrix \mathbf{L} is obtained, it is possible to compute robust POD/PCA modes (Sec. 2.2.1) as in Eq. (2.2):

$$\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T. \quad (2.19)$$

Henceforth, we refer to the modes in \mathbf{U} from \mathbf{UL} as RPCA modes. We note that in many flow applications it is important to subtract the mean flow before computing POD, which allows the POD eigenvalues to be interpreted as the variance of fluctuations and the expansion to respect boundary conditions by construction. However, before computing RPCA, it may be difficult to obtain an accurate mean flow estimate. Instead, we advocate computing the RPCA, then subtracting the mean of \mathbf{L} from itself, and finally computing POD on the mean-subtracted low-rank matrix; this final POD step will remove small amounts of white noise.

Similarly, it is also possible to compute robust DMD (RDMD) modes and eigenvalues. We note that the RDMD should not be confused with the recursive DMD of Noack et al. [105], which uses the same acronym. We also note that this decomposition is similar in spirit to the coherent vortex simulation (CVS) approach [54], which separates turbulent flows into coherent and random parts based on a wavelet decomposition. However, RPCA does not perform this decomposition using a universal basis, such as wavelets, that rely on scale separation, but rather based on statistical correlations in the data.

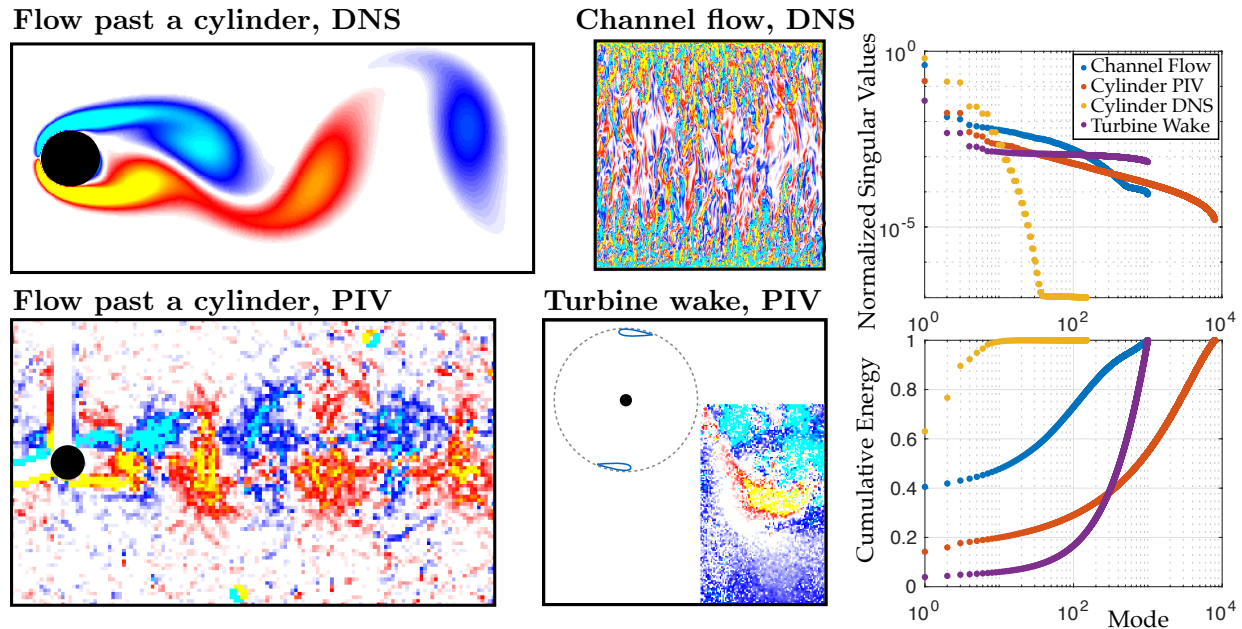


Figure 2.3: (left) Example flow field data. (right) Singular value spectrum for each data set. Mean flow travels from left to right in all cases.

2.4 Model flows

We demonstrate RPCA filtering on several example data sets of varying complexity, drawn from direct numerical simulations (DNS) and PIV data from experiments. Figure 2.3 provides an overview of the four example flow fields.

2.4.1 Cylinder flow

Flow past a cylinder is a canonical example in fluid mechanics. We consider data from DNS at a diameter-based Reynolds number of 100 and from PIV measurements at Reynolds number 413 [135].

The DNS data is generated by simulating the two-dimensional incompressible Navier-Stokes equations using the immersed boundary projection method [36, 143]. The computational domain comprises four nested grids: the finest grid covers a domain of 9×4 and the largest grid covers a domain of 72×32 , where lengths are non-dimensionalized by the

cylinder diameter. Each grid contains 449×199 points with a resolution of 50 points per cylinder diameter. The time-step is $\Delta t = 0.02$ and data is sampled at intervals of $10\Delta t$ (30 times the vortex shedding frequency) with $m = 150$ snapshots saved, covering 5 vortex shedding cycles. The DNS provides a benchmark, where the uncorrupted flow field is known, to quantitatively assess performance of RPCA filtering on data with artificial salt & pepper corruption. Corrupted sample points are chosen uniformly in space and time at a given rate, and both the u and v velocity components at each selected location are randomly assigned a value of ± 10 times the standard deviation of the streamwise velocity data. In addition, we consider a second case where corrupted sample points are chosen with a bias towards regions of high vorticity or shear magnitude, which is more physically realistic for PIV data. In this second case, we select measurements for corruption based on a probability density given by $\alpha + |\boldsymbol{\omega}|$, where α is a small positive constant and $|\boldsymbol{\omega}|$ is the absolute value of the vorticity; when the rate of corruption is sufficiently high, these corrupted fields begin to resemble the uniformly corrupted cases, but with more corruption in vortex cores. Because vorticity is calculated from velocity fields using a finite-difference derivative, there is a higher rate of corruption in the vorticity fields than in the velocity fields.

The PIV data has frame size of 135×80 grid points with a resolution of 8 points per cylinder diameter. Data is sampled at a rate of 20 Hz (125 times the shedding frequency) with $m = 8,000$ snapshots saved, which corresponds to 64 vortex shedding cycles.

2.4.2 Turbulent channel flow

For a more complex and multi-scale flow, we consider DNS data from a forced, fully developed turbulent channel flow data with a friction velocity Reynolds number of $Re_\tau = 1,000$, from the Johns Hopkins Turbulence Database (JHTDS) [68]. This example provides a test case to see how turbulent kinetic energy at various scales is filtered depending on the level of added noise. The addition of noise is similar to the cylinder DNS where randomly selected sample points of the streamwise and cross-stream velocity fields are assigned a value of ± 10 standard deviations of the streamwise velocity data. Due to the size of the full dataset,

we only consider two-dimensional fields on the mid-plane, with a 512×512 grid of three component velocity measurements spanning the channel width. Data is sampled at a rate of 966 times the mean flow-through time with $m = 1000$ snapshots.

2.4.3 Cross-flow turbine wake

Finally, we consider PIV wake data from a cross-flow turbine experiment conducted at the University of Washington. Cross-flow turbines can be used to extract power from wind and water currents for renewable energy generation. This flow exhibits both coherent and broad-band phenomena, and provides a challenging test-case RPCA filtering. The frame consists of 158×98 grid points, at a resolution of 99 points per rotor diameter. Data is sampled at a rate of 32 times the blade-pass frequency with $m = 1000$ snapshots. Vectors were calculated using a multi-grid, multi-pass algorithm with adaptive image deformation [127]. Resulting vector fields were then validated using a normalized median filter with potential replacement by secondary correlation peaks. The cross-correlation and validation steps result in missing data, particularly in regions of high vorticity and shear. To apply RPCA filtering, these missing values are randomly assigned a value of ± 10 standard deviations of the streamwise flow data, in contrast to the experimental cylinder wake where missing measurements were previously interpolated.

2.5 Results

We now explore the ability of RPCA filtering to isolate and remove noise and corruption from the example flow fields. We will begin with the simulated and experimental flow past a cylinder, followed by data from the Johns Hopkins turbulent channel flow simulation, and ending with the experimental wake of a cross-flow turbine.

2.5.1 Cylinder flow

Figure 2.4 shows the results of RPCA filtering for flow past a cylinder, providing a side-by-side comparison of PIV and corrupted DNS data. Although the Reynolds numbers differ by

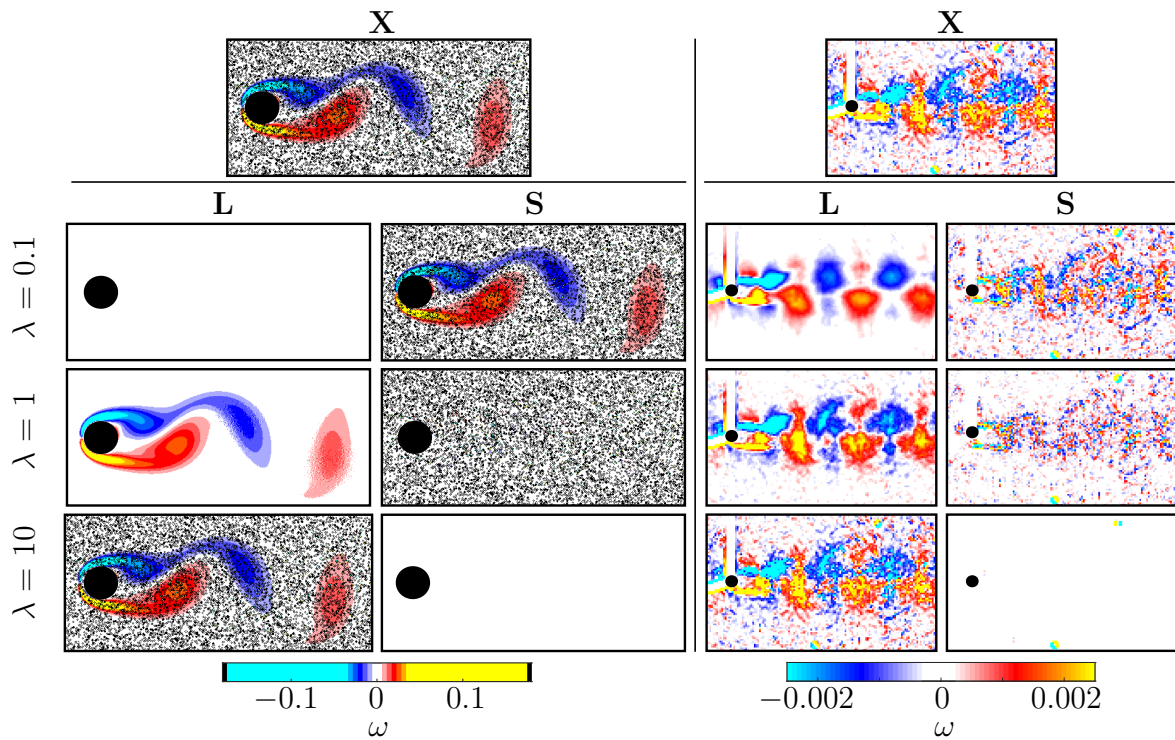


Figure 2.4: RPCA filtering removes noise and outliers in the flow past a cylinder (black circle), from DNS (left) with 10% of velocity field measurements corrupted with salt and pepper noise, and PIV measurements (right). All frames show resultant vorticity fields. As the parameter λ is decreased, RPCA filtering is more aggressive, eventually incorrectly identifying coherent flow structures as *outliers*.

a factor of four, the flow fields are qualitatively similar, characterized by periodic, laminar vortex shedding. For $\lambda = 1$, the data is correctly segmented with the coherent flow in \mathbf{L} and the sparse corruption in \mathbf{S} . When λ is too small, RPCA filtering is overly aggressive, incorrectly including relevant flow structures in \mathbf{S} , and when λ is too large, the corruption is not filtered.

For the experimental data in Fig. 2.4 (right), the optimal value of λ is less clear. For $\lambda = 0.1$, the low-rank field \mathbf{L} is visually smoother than the field at $\lambda = 1$, but the sparse matrix \mathbf{S} contains a significant portion of the wake structures, indicating over-filtering. This filtering becomes more pronounced in the movies, where it is clear that much of the high-frequency “noise” in the bypass flow is actually free-stream turbulence, which is consistent with the turbulence intensity of the experiments. Further, as subsequently discussed, when we compute the RPCA modes, it is clear that the $\lambda = 0.1$ case is heavily filtering out all but the first three modes. Thus, it appears that the theoretically optimal value $\lambda = 1$ has the best performance, although there may be a tradeoff between filtering ambient free-stream turbulence and coherent structures of interest in experiments.

Figure 2.5 shows the results of RPCA filtering on the simulated data for the second case of vorticity-biased corruption. Again, in all cases, the theoretically optimal value of $\lambda = 1$ yields the best segmentation of the corruption into the matrix \mathbf{S} . When the rate of corruption is increased from 1% (left) to 10% (right), the free-stream flow begins to become corrupted, resembling the uniform corruption case in Fig. 2.4. The mean error and relative nuclear norm of the low-rank matrix (\mathbf{L}) compared to the true, uncorrupted data (\mathbf{X}) are shown in Fig. 2.6 for varying percentages of corrupt entries. Statistically, results are similar for vorticity-biased and randomly-distributed corruption. In both cases, RPCA filtering is remarkably robust to corruption, even for corruption in excess of 50% of the measurements. This laminar vortex shedding example is an ideal application for RPCA filtering, as the true flow field is low rank and the corruption is sparse; it is unlikely that this will hold as well for data exhibiting broadband turbulence.

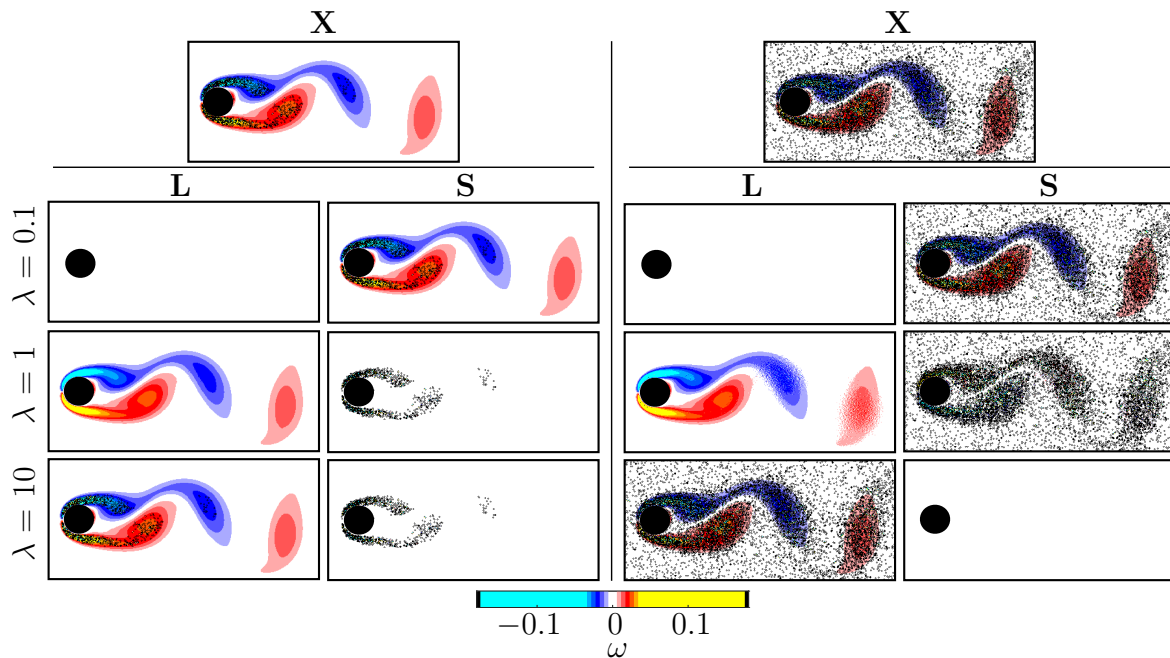


Figure 2.5: RPCA filtering removes vorticity-biased corruption from simulated flow past a cylinder at Reynolds number 100. Unlike results shown in Fig. 2.4, corrupt entries are concentrated in regions of high vorticity instead of being uniformly distributed. In the flow on the left, $\eta = 1\%$ of the velocity field measurements are corrupted and on the right $\eta = 10\%$ of the velocity field measurements are corrupted. All frames show resultant vorticity fields.




Figure 2.6: Error ($\|\mathbf{X}_{uncorrupted} - \mathbf{L}\|_F / \|\mathbf{X}_{uncorrupted}\|_F$) and relative nuclear norm ($\|\mathbf{L}\|_* / \|\mathbf{X}\|_* = \text{sum}(\sigma_L) / \text{sum}(\sigma_{\mathbf{X}, \text{uncorrupted}})$) of the low rank matrix \mathbf{L} compared with the uncorrupted data \mathbf{X} for varying percentages of corruption.

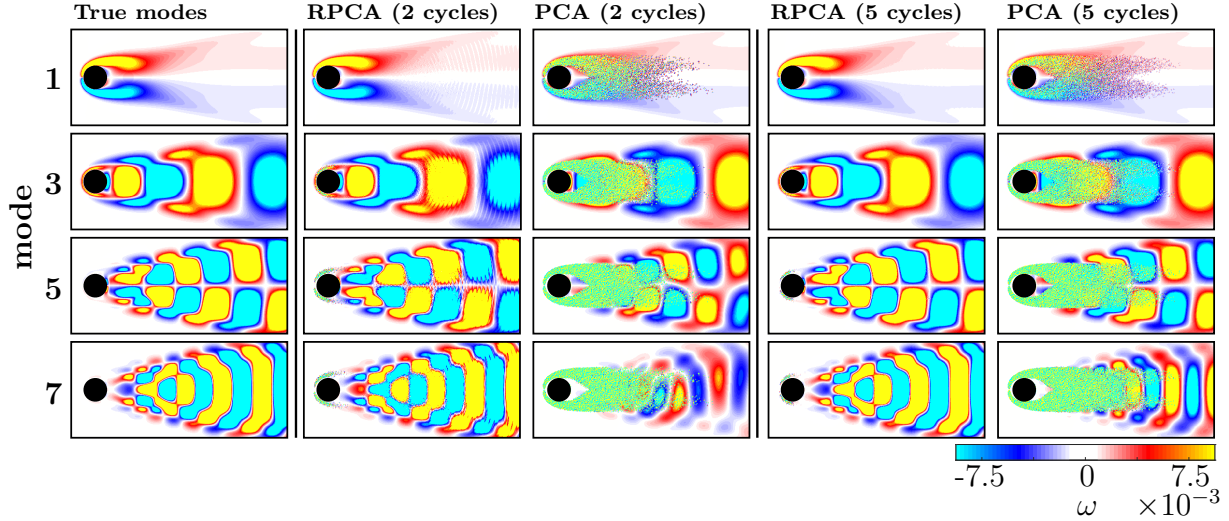


Figure 2.7: Odd PCA vorticity modes of the cylinder simulations from Fig. 2.5 with 1% of velocity measurements corrupted with a bias towards regions of high vorticity.

PCA analysis for cylinder wake flows

We now investigate the impact of RPCA filtering on modal decompositions. Figures 2.7 and 2.8 show four leading PCA and RPCA modes for 1% and 10% vorticity-biased corruption, respectively. The first mode corresponds to the mean flow, and the remaining modes come in energetic pairs where the corresponding coefficients α_{2j} and α_{2j+1} oscillate sinusoidally at the same frequency but $\pi/2$ out of phase, sweeping out a circle in the phase plane. Thus, we only show one mode, \mathbf{u}_{2j+1} from each of the first three energetic mode pairs. In all cases, the RPCA modes show dramatic improvement, while significant artifacts remain in the PCA modes. We also investigate the effect of increasing the amount of data, and there is a clear improvement in RPCA modes from 2 to 5 vortex shedding cycles; in contrast, the PCA modes do not improve appreciably with more data.

To quantify the improvement observed above, we compute the L_2 error between the PCA and RPCA modes of corrupted data and the PCA modes for the clean data (i.e., DNS results) as a function of the number of shedding periods. As show in Fig. 2.9, the RPCA mode velocity fields quickly converge to a small error as the amount of data is increased for

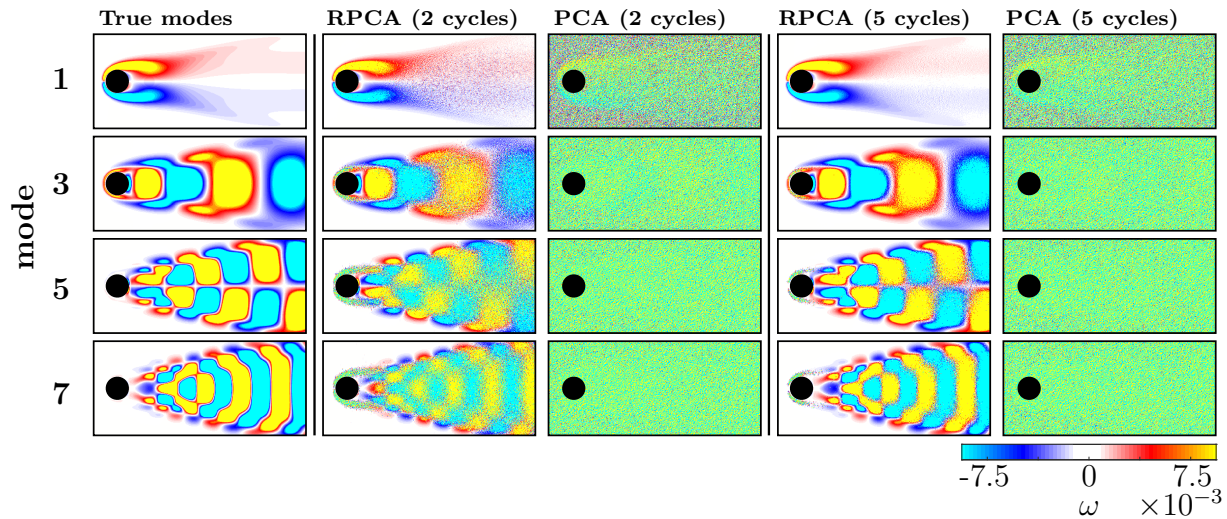


Figure 2.8: Odd PCA vorticity modes of the cylinder simulations from Fig. 2.5 with 10% of velocity measurements corrupted with a bias towards regions of high vorticity.

both the 1% and 10% corruption cases, while the PCA modes converge much more slowly and still have considerable error after 5 shedding periods are included in the analysis.

The modes for the PIV data for the cylinder flow are shown in Fig. 2.10. This figure highlights the effect of λ , the sparsity hyperparameter, which was previously discussed with respect to Fig. 2.4. In this case, we do not have a clean ground-truth data set to compare against. Although the flow field in Fig. 2.4 appears to have less corruption for $\lambda = 0.1$, here we see that all RPCA modes after the first three modes are heavily filtered, as seen in the rapid drop off in the singular values after the third mode. The corresponding modes are highly corrupt, further supporting that $\lambda = 0.1$ is not a good choice. In contrast, the RPCA modes for the theoretically optimal $\lambda = 1$ case appear to have slightly less free-stream corruption than the PCA modes. Also, as expected, for a large enough value of λ , the RPCA filtering has little effect on the modes.

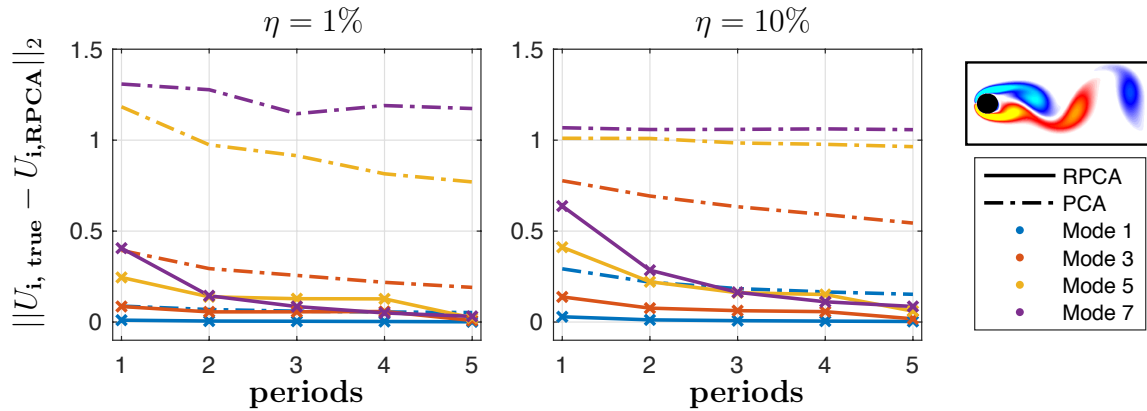


Figure 2.9: L_2 error between the true PCA modes of the clean cylinder simulation data \mathbf{X} and the RPCA and PCA modes for corrupted data with 1% (left) and 10% (right) vorticity-biased corruption.

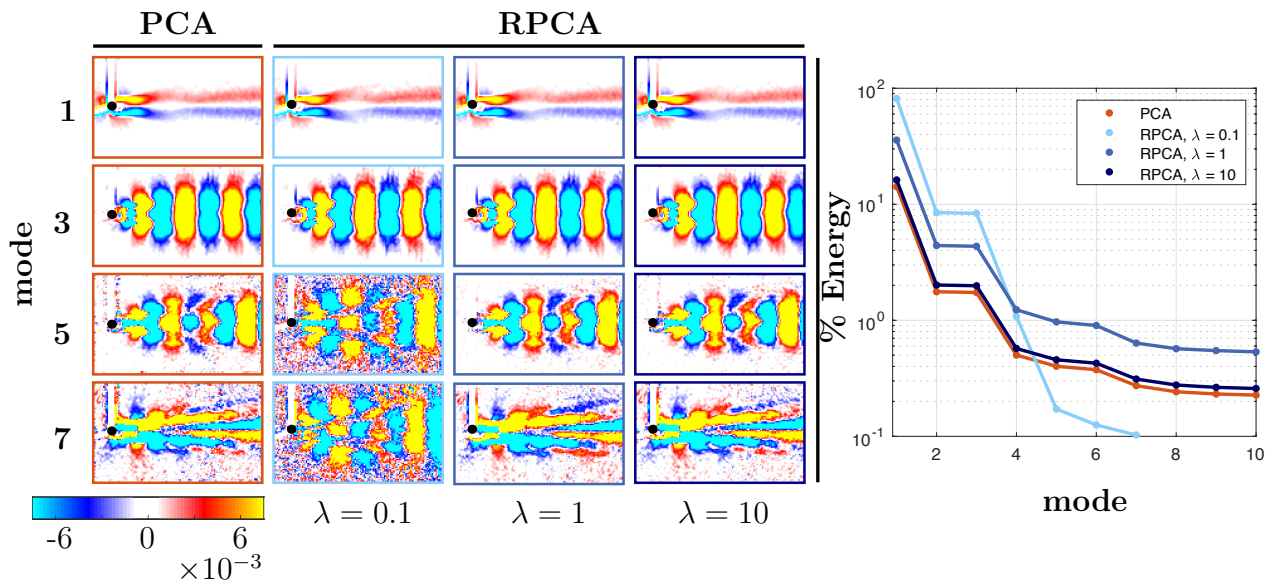


Figure 2.10: Odd PCA vorticity modes of the experimental cylinder data for PCA and RPCA at $\lambda = 0.1, 1, \text{ and } 10$, along with their singular values.

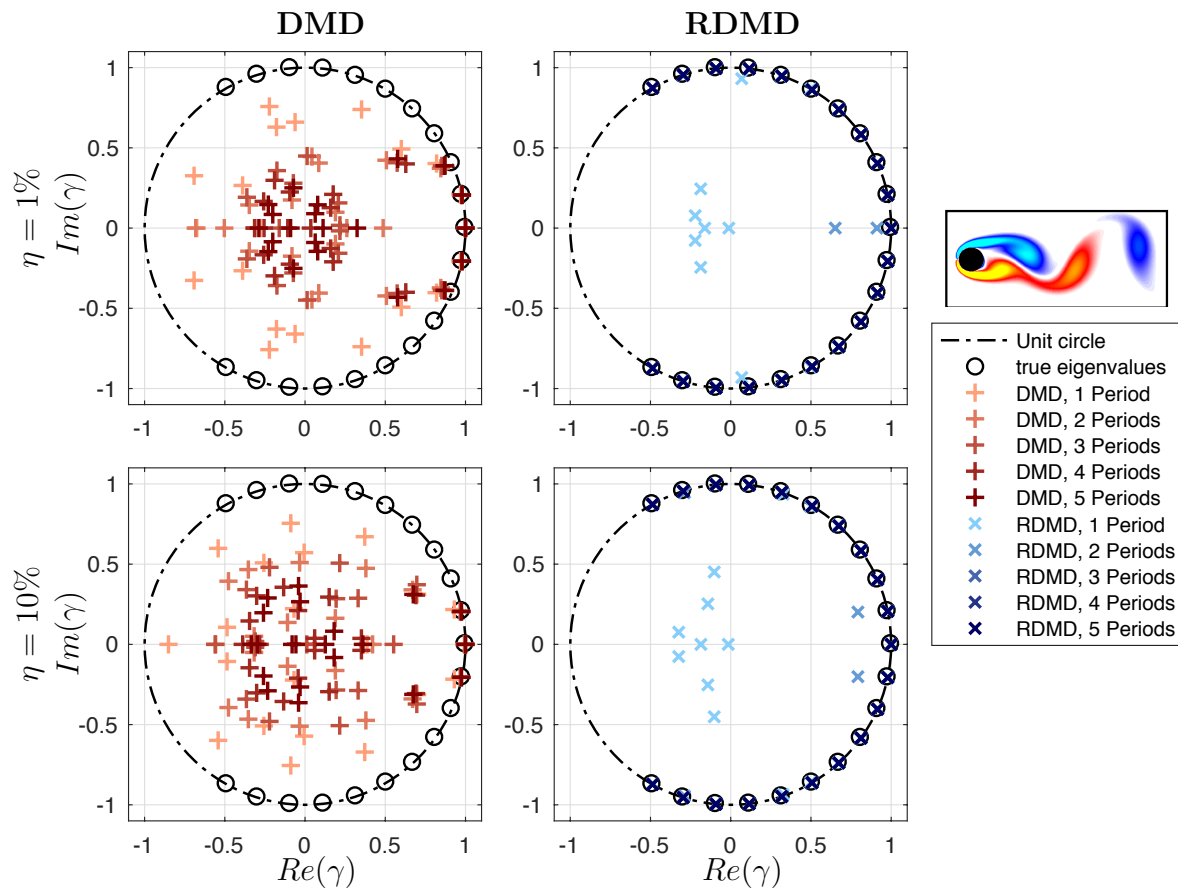


Figure 2.11: Discrete-time DMD eigenvalues for the simulated cylinder data for small and large amounts of corruption and for increasing amounts of training data. In all cases, the RPCA-filtered DMD results dramatically outperform the standard DMD results.

DMD analysis for cylinder wake flows

DMD is known to be quite sensitive to noisy data, making this a challenging test case for RPCA filtering. Figure 2.11 shows the discrete-time eigenvalues for the cylinder DNS data with vorticity-biased corruption. For the cylinder wake, the uncorrupted or true DMD eigenvalues may be computed from the noiseless data, and they are equally spaced on the unit circle in the complex plane. In all cases, the RPCA-filtered DMD (RDMD) data results in dramatically better agreement with the uncorrupted or true DMD eigenvalues compared with the corrupted DMD eigenvalues. Even with only a single period of data and $\eta = 10\%$ corruption, the RDMD values capture the first six low-frequency mode pairs; in contrast, even with five periods of data and as little as $\eta = 1\%$ corruption, corrupted DMD only captures the first two low-frequency mode pairs, and with considerably more spurious damping. To see this more clearly, we plot the eigenvalues in continuous-time in Fig. 2.12, where the x -axis is the imaginary eigenvalue component and the y -axis is the real eigenvalue component, which is a standard way to plot DMD eigenvalues [132]. Here, the best-fit parabolas for the RDMD eigenvalues and the first seven corrupted DMD eigenvalues are shown in dashed lines. The curvature of these parabolas is directly related to the noise amplitude, as in Eq. (2.14) from [13]. The same continuous-time eigenvalue plot is shown for the PIV cylinder wake data in Fig. 2.13. In both cases, we see that the parabolic eigenvalue fit for the RDMD eigenvalues has a smaller curvature than for the corrupted DMD eigenvalues, indicating a quantitative and significant reduction in noise.

2.5.2 Turbulent channel flow

The tradeoff between filtering corruption and small-scale structures is also apparent in the turbulent channel flow DNS. Unlike the cylinder wake, this flow field contains broadband turbulent phenomena across multiple spatial and temporal scales. Figure 2.14 shows RPCA filtering for various levels of corruption, sweeping across the tuning parameter λ . The corresponding turbulent kinetic energy (TKE) is shown in Fig. 2.15, providing a summary of the




Figure 2.12: Continuous-time DMD eigenvalues for the simulated cylinder data along with parabolic eigenvalue fits to estimate the error as in [13]. Here we use 5 vortex shedding periods with $\eta = 1\%$ corrupt values. The RDMD parabolic coefficient is approximately 2×10^4 times smaller than the DMD coefficient.




Figure 2.13: Continuous-time DMD eigenvalues for the PIV cylinder data along with parabolic eigenvalue fits to estimate the error as in [13]. The RDMD parabolic coefficient is approximately two times smaller than the DMD coefficient.

various scales that are filtered. The value of λ that preserves the true TKE spectrum varies with the degree of velocity field corruption. In the uncorrupted case ($\eta = 0$), we can clearly see the effect of filtering on the turbulent coherent structures, indicating that some fine-scale structures are filtered for $\lambda = 1$. As the degree of corruption increases to $\eta = 2\%$, we see that the curves for $\lambda \leq 2$ remain relatively unchanged, although the $\lambda > 2$ curves begin to exhibit spurious high-frequency spatial structures (i.e., corruption is present in \mathbf{L}). As the rate of corruption increases to $\eta = 10\%$, spurious high-frequency energy also appears for $\lambda = 2$. In this case, it is clear that the optimal filtering value λ changes with the level of corruption. For relatively limited corruption, a larger value of λ may be used, but must be decreased towards the theoretically optimal value of $\lambda = 1$ for higher levels of corruption. Finally, we note that, unlike the cylinder wake cases, it is not surprising that $\lambda = 1$ is sub-optimal because the channel flow is not fundamentally low rank which deviates from an underlying assumption of the RPCA algorithm.

2.5.3 Cross-flow turbine wake

As a final example, we consider the use of RPCA filtering to identify outliers and fill in missing PIV data collected in the wake of a cross-flow turbine, as shown in Fig. 2.16. There are several stages in the PIV processing pipeline where RPCA filtering could be applied, including after initial cross-correlation, after conventional normalized median filter vector validation, and after linear interpolation. For the cases shown here, we use $\lambda = 1.6$, which results in a velocity in the bypass flow, or lower third of the frame, that visually matches the frequency content of the unfiltered data. There are enough missing velocity vectors (23% and 20%, respectively) to degrade the effectiveness of both median filtering and interpolation. In contrast, RPCA filtering produces flow fields that capture dominant coherent structures for either cross-correlated or median-filtered fields. Finally, by investigating the standard deviation of all flow fields collected at a given turbine angular position (i.e., phase), it is clear that the RPCA filtering can be used to remove artifacts introduced by linear interpolation. This is consistent with the intuition that vector validation and interpolation should fail in

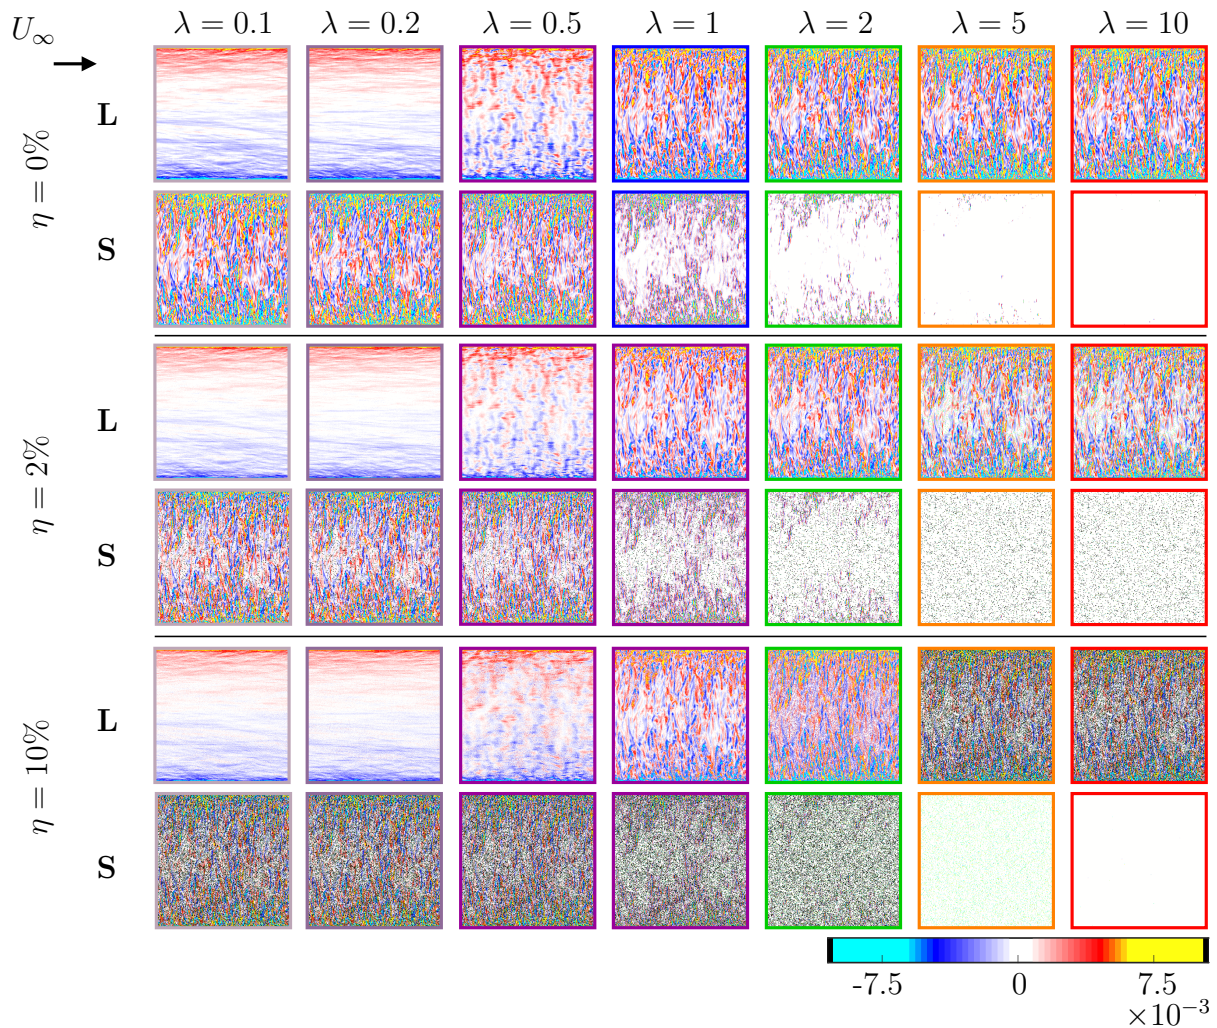


Figure 2.14: RPCA filtering for turbulent channel flow vorticity fields with various levels of added noise and tuning parameter λ . η represents the percentage of corrupted measurements in the velocity fields. The border colors match the color of the curve at the corresponding η in Fig. 2.15.




Figure 2.15: Turbulent kinetic energy (TKE) spectra for various levels of corruption and RPCA filtering. The TKE profiles provide a summary of the filtering that occurs at various scales. As corruption increases, the filtering remove more high-frequency information.

these regions where there is high density of missing data is spatially clustered.

The continuous-time DMD eigenvalues for the cross-flow turbine wake are shown in Fig. 2.17. In this plot, the parabolic fits for DMD eigenvalues computed after interpolation and RDMD eigenvalues computed after vector validation are displayed as dotted lines. For this case, the coefficient of the parabolic fit for the RDMD-based eigenvalues is six times smaller than the parabolic fit coefficient for the DMD-based eigenvalues. This demonstrates a significant quantitative improvement of the DMD spectrum using RPCA filtering to process the data.

2.6 Conclusions and discussion

In this chapter, we have demonstrated the ability of RPCA filtering to effectively recover dominant coherent structures from corrupt flow fields with missing measurements. Unlike standard POD/PCA, which is based on least squares and is susceptible to outliers and corruption, RPCA utilizes sparse optimization to separate a data matrix into a low-rank matrix containing correlated structures and a sparse matrix containing the spurious entries.

We apply RPCA filtering to several types of fluid flow data (DNS and PIV), ranging from laminar vortex shedding behind a circular cylinder, to fully turbulent channel flow DNS, and

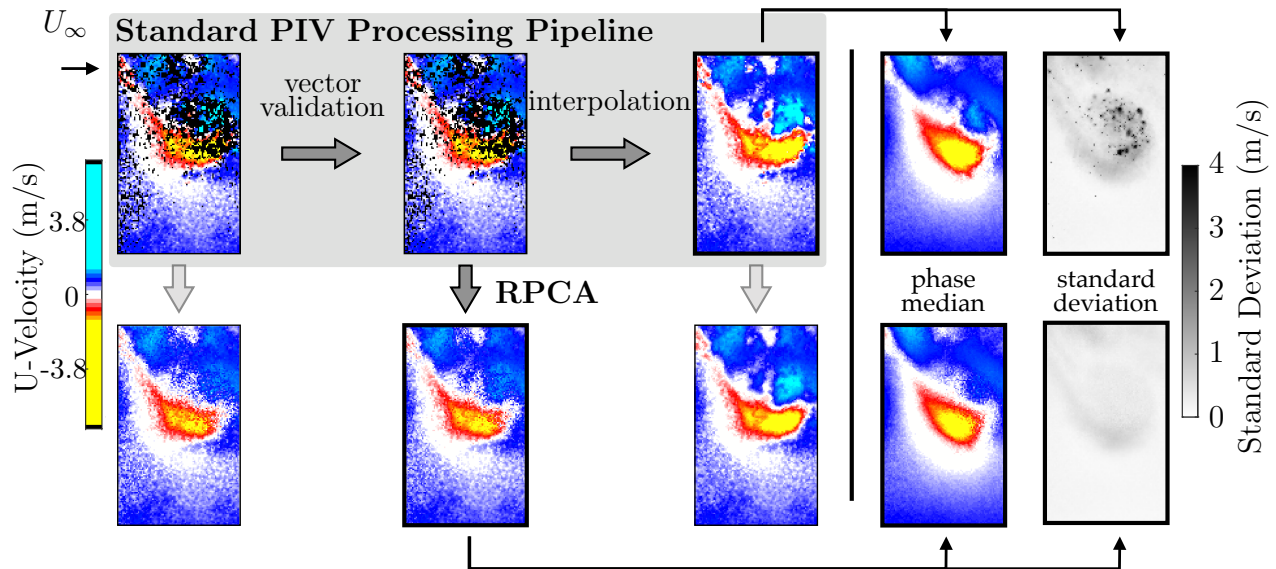


Figure 2.16: RPCA filtering of cross-flow turbine wake PIV data. The standard PIV processing pipeline (top row) includes several steps where RPCA filtering can be applied (bottom row). In the cross-correlated streamwise velocity field (top left), 23% of the velocity vectors are missing. Vector validation reduces the missing vectors to 20%. Finally, linear interpolation is used to fill in these missing vectors. In all cases, RPCA filtering captures the relevant phase-averaged coherent structures with fewer outliers and missing data, which appear as dark spots in the standard deviation plot.




Figure 2.17: Continuous-time DMD eigenvalues for the turbine wake PIV data, along with parabolic eigenvalue fits to estimate the error as in [13]. The RDMD parabolic coefficient is approximately six times smaller than the DMD coefficient.

concluding with an experimental flow past a cross-flow turbine. These flows exhibit a variety of phenomena and a range of measurement quality. The DNS examples provide us with a baseline, where it is possible to add corruption to quantitatively assess the performance of RPCA. For flow past a cylinder in DNS, RPCA filtering is extremely effective at separating the true flow field from considerable corruption, with robust recovery even in flow fields with excess of 50% of the measurements corrupted. In the experimental counterpart, RPCA is still able to remove large outliers and corruption, although there is a tradeoff between filtering the background turbulence and coherent structures in the wake. The fully turbulent channel flow DNS provides an opportunity to more fully explore this tradeoff in a controlled setting, where we can incrementally increase the corruption ratio and observe the filtering effects on various spatial frequencies. As expected, an increasingly aggressive filtering leads to degradation at higher wavenumbers, although dominant coherent structures are robustly preserved. Finally, the wake behind a cross-flow turbine provides a practical real-world flow that directly benefits from improved PIV processing. In all three wake flows we also assess the performance of RPCA filtering to yield more accurate modal decompositions. Although we do not have ground truth measurements and modal decompositions, except in the case of direct numerical simulations, we know that continuous-time DMD eigenvalues should be arranged on the imaginary axis in the complex plane for clean data, and deviations from this may be quantified using the derivation from Bagheri [13]. In all three cases, we see considerable reduction in spurious damping, indicating the de-noising effectiveness of RPCA. Based on these results, we believe that RPCA can be a valuable algorithm in the arsenal of PIV processing and filtering techniques, particularly when the processing pipeline culminates in modal analysis.

There are a number of future directions motivated by this work. First, RPCA depends on the hyperparameter λ , and a better understanding of how to objectively choose λ for different conditions is important. Because RPCA is based on sparse, non-convex optimization, it is also likely that improved optimization techniques may improve speed and robustness. Although this work considered three-dimensional flows, the data comprised two-dimensional

cross-sections, and the current analysis could be extended to flow volumes. In principle, the RPCA method should generalize, although there may be computational scaling challenges. Recent results have extended RPCA from linear subspaces to manifolds [94], so it may be possible to robustly characterize fluids data that is well-described by a low-dimensional manifold [91], rather than a low-dimensional POD subspace. Non-stationary flows may be more challenging for this method, as the bulk distribution will drift. Similarly shocks may be erroneously flagged as outliers; however, this may provide an opportunity to identify shocks in the data. Investigating these flows is an important avenue of future work. It would also be useful to extend this work to PIV measurements of other turbine configurations [114]. Finally, the quality of the RPCA filtered flow fields for additional downstream analyses should be assessed for example, in dynamical systems modeling via Galerkin projection [103] or regression [90] onto the filtered modes and in control [15].

Chapter 3

GEOMETRIC AND CONTROL OPTIMIZATION OF A TWO TURBINE ARRAY

3.1 Contributions of this chapter

In this chapter, we experimentally perform simultaneous optimization of turbine configuration and control with a two-rotor array in a constant free-stream velocity. We employ two control methods that could take advantage of both mean flow alteration and periodic flow structures to maximize the performance of the turbines in the array. This is motivated by the inherently periodic nature of cross-flow turbine fluid dynamics, in which the angle of attack encountered by the blades is continuously changing. The resulting unsteady fluid mechanics can result in leading edge vortex separation [51] and dynamic stall [27, 26]. Even at tip-speed ratios high enough to suppress dynamic stall [87, 55, 26, 50, 60], the angle of attack variation produces a phase-dependent pressure and velocity field in and around the rotor. We first consider a standard control method that varies the tip-speed ratios (ratio of tangential blade speed to free-stream velocity) of each of the turbines, exploiting mean flow alterations to maximize array power output. Throughout, we refer to this control strategy as *tip-speed ratio control*. We also introduce a type of control with a synchronous rotation rate and a phase offset between the turbines in the array. We refer to this strategy, which coordinates the operation of turbines to exploit mean *and* periodic flow field alterations, as *coordinated control*. As prior work has evaluated performance benefits for co-rotating and counter-rotating turbines [21], we consider both modalities in our experiments.




Figure 3.1: (a) Experimental set-up of the mobile, cantilevered turbine which consisted of a motor to enforce consistent tip-speed ratios and phase differences, a torque cell, an air bearing to absorb cantilever loads, and the rotor. (b) Experimental set-up of the fixed turbine, which consists of two load cells, a motor, the vacuum plate, and the rotor. (c) Photo of both turbines operating in the Bamfield Marine Science Centre flume.

3.2 Methods

Array experiments were performed in the Bamfield Marine Science Centre recirculating water channel. During these tests, one turbine was fixed in space while the other turbine was cantilevered from a robotic gantry system (Velmex BiSlide) that allowed precise control of the rotor position in the streamwise and cross-stream directions. This enabled accurate and reproducible array geometries.

The rotation rate of the fixed turbine (Fig. 3.1b) was regulated by a servomotor [109] (Yaskawa SGMCS-05B with Yaskawa SGD-05B3C41 drive). The motor had an internal encoder with over one million counts per revolution and was rigidly connected to a six-axis load cell (ATI Mini45) that measured all reaction forces and torques and was fixed to the flume structure. The turbine driveline was a 12.7 mm diameter stainless steel shaft that terminated, at the lower end, in a bearing that was fixed to a second six-axis load cell (ATI Nano25). This assembly was fixed to the bottom of the flume using a suction

plate and scroll vacuum pump (Agilent IDP3) and has been used extensively in previous work [139, 141]. For the mobile turbine, elimination of the lower load cell and vacuum plate would have introduced significant cross-talk in an upper load cell due to the orders-of-magnitude difference between the thrust-induced moment and the relatively small rotary torque. As an alternative approach, as shown in Figure 3.1a, the turbine driveline was supported by an air bearing with negligible friction (Professional Instruments Company Block-Head 4R low-inertia). For the mobile turbine, the rotation rate was regulated by a servomotor with internal encoder (Yaskawa SGMCS-02B3C41) and the reaction torque (equivalent to fluid torque) was measured by a torque cell (Futek F400). A picture of both turbines operating in the Bamfield Marine Science Centre flume is shown in Figure 3.1c.

Both turbines had two blades with strut endplates to minimize parasitic losses [141]. Each turbine had a height of $H = 0.23$ m, diameter of $D = 0.172$ m, chord length of $c = 0.04$ m, and symmetric NACA0018 blade profile. This turbine geometry has been used in past work [139, 141, 109, 78]. Due to the limited number of strut endplates on hand, one turbine had struts with a NACA0008 profile, while the other used a NACA0016 profile. This gave rise to minor performance variations (Fig. 3.2c). Both sets of struts had chord lengths that are equal to those of the blades (0.04 m).

The Bamfield Marine Science Centre recirculating water channel was 10 m long, 2 m wide, and was filled to a dynamic depth of $h = 0.73$ m. The nominal flow speed was $U_\infty = 0.6$ m/s with a turbulence intensity of approximately 2%. The freestream flow velocity was measured upstream of the array using an acoustic Doppler velocimeter (Nortek Vector) at a rate of 64 Hz. Despiking was performed using the method of Goring and Nikora [67], with spikes replaced by interpolation. The temperature was held at 18 ± 1 °C. These conditions corresponded to a chord-based Reynolds number of $Re_c = \frac{U_\infty c}{\nu} \approx 22,000 - 23,000$, where ν is the kinematic viscosity, and depth-based Froude number of $Fr = \frac{U_\infty}{\sqrt{gh}} \approx 0.22$, where g is the gravitational acceleration or 9.81 m/s^2 and h is the mean water depth. The Reynolds number placed the turbines in a transitional regime where performance varied with local Re_c [98]. The geometric confinement, defined as the turbines' projected area normalized by the area

of the channel, ranged from 2.8% (turbines directly in-line) to 5.6% (turbines side-by-side), which is relatively low [121]. The implications of this are explored for a sub-set of results in Sec. 3.3.5.

3.2.1 *Experimental design*

The test matrix of relative rotor positions sampled is given in Figure 3.2a. Arranged on a polar grid, positions ranged from side-by-side ($X = 0$) to $X = 3.61D$ in the streamwise direction, and a maximum spacing of $1.83D$ in the cross-stream direction. The minimum distance between rotor centers was $1.1D$.

The symmetry of the polar grid about the Y/D axis suggests that the same parameter space could be sampled more efficiently if only the $Y/D \geq 0$ region is probed. However, two arrays with equal X/D spacing and equal and opposite Y/D spacing do not yield identical performance, as a consequence of the asymmetric nature of a cross-flow turbine wake.

The co-rotating arrays are symmetric when the mobile turbine is at $X/D = 0$ and the Y/D values are equal and opposite (or the absolute center-to-center distances are equal). Nevertheless, those points are repeated because the two turbines and test rigs that constitute the array are not physically identical. For side-by-side counter-rotating turbines, no such symmetry exists. In this case, the two possible configurations are characterized by the direction the blades travel as they approach the space between the two turbines. In one orientation they are traveling upstream, in the other they are traveling downstream, as they approach the space between the two rotors [162, 149, 150].

3.2.2 *Performance Metrics*

The coefficient of performance for an individual turbine is the ratio of the power produced to the kinetic power in the free-stream passing through the turbine's projected area and expressed as

$$C_p = \frac{P}{\frac{1}{2}\rho U_\infty^3 HD} = \frac{\omega\tau}{\frac{1}{2}\rho U_\infty^3 HD}, \quad (3.1)$$

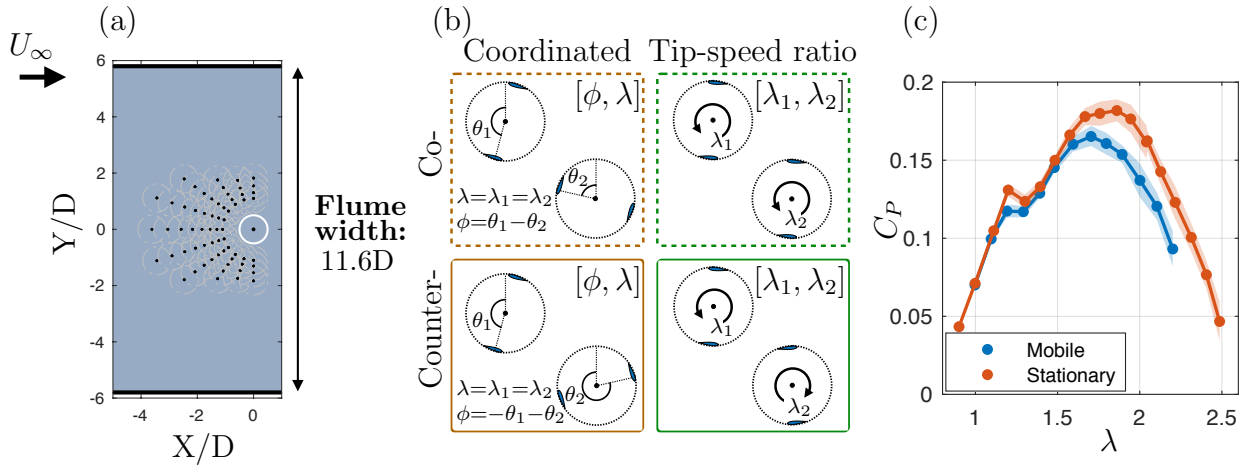


Figure 3.2: (a) Experimental matrix showing the position where each control strategy was tested. (b) Each array configuration tested, where one turbine was fixed at $X/D = Y/D = 0$ and the other turbine was tested at each prescribed polar grid location. (c) Performance curve with the interquartile range for each turbine operating in isolation.

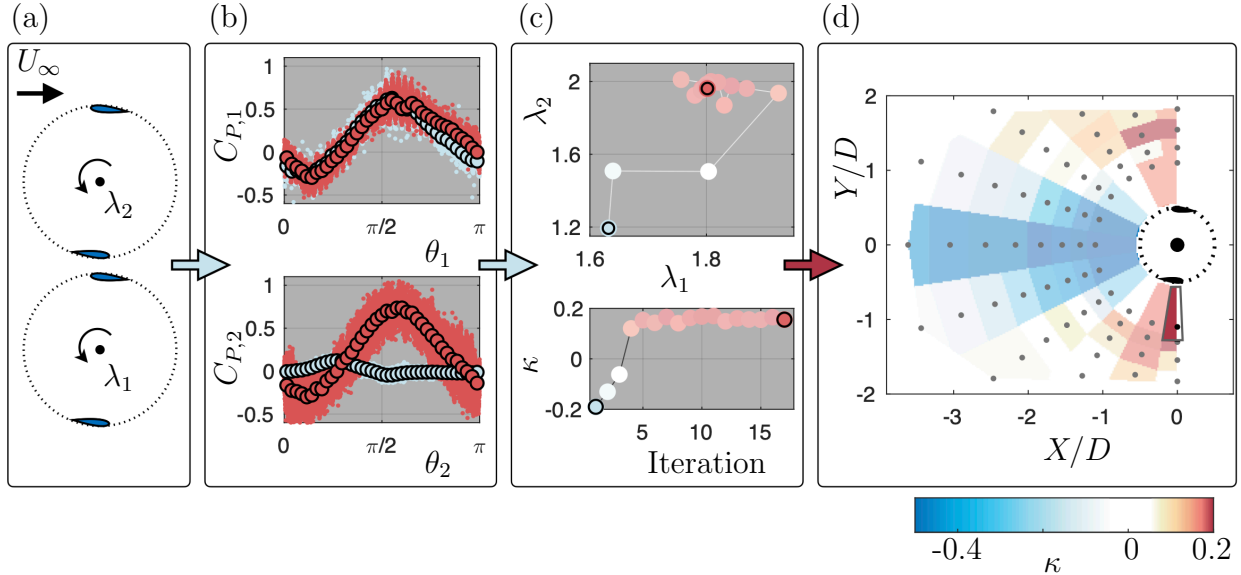


Figure 3.3: Optimization for one geometric configuration of co-rotating turbines under tip-speed ratio control contrasting initial case (light blue) with optimized case (red). From left to right: (a) the co-rotating turbines and direction of free-stream flow; (b) performance comparison of initial and optimized cases for each turbine as a function of the angular position of each turbine ($\theta = 0$ when the blade is pointing directly upstream) color encoded to denote interaction factor; (c) top shows the optimization path and bottom shows the performance evolution during optimization color encoded to denote interaction factor; (d) resulting performance of selected array geometry on a heat map corresponding to all test positions.

where P is turbine’s mechanical power, ρ is the fluid density, U_∞ is the freestream flow velocity, H is the turbine height, D is the turbine diameter, ω is the turbine rotation rate, and τ is the turbine torque. It is important to note that the undisturbed freestream flow velocity is used to normalize performance for both turbines, even when one turbine is within the wake of another. This is appropriate because we are exploring the performance of the array, not how heterogeneous inflow (e.g., rotation, shear) affects the fundamental performance of the individual turbines. In addition, particularly for cases with staggered arrays and limited inter-turbine separation, it would be difficult to uniquely define a single-valued inflow velocity or obtain sufficient information about the flow field to do so.

Our fundamental parameter of interest in this study is an “interaction factor” that describes how the performance of the two, interacting turbines compares to their performance in isolation. In experiments, the two turbine rotors were tested separately to establish a baseline performance for comparison. Their individual peak performances are denoted by $C_{P,1}^*$ and $C_{P,2}^*$, where the asterisk denotes performance at the optimized tip-speed ratio (peak of the $C_P - \lambda$ curves in Figure 3.2c). The array performance for a given configuration and control strategy is evaluated relative to the sum of the peak power that the two turbines would produce in isolation as

$$\kappa = \frac{C_{P,1} + C_{P,2}}{C_{P,1}^* + C_{P,2}^*} - 1 \quad (3.2)$$

where κ is the interaction factor ($\kappa < 0$ for detrimental interactions and $\kappa > 0$ for beneficial interactions). We note that if mean inflow was perfectly constant across all tests, κ would be equivalent to the power produced by the turbines in the array divided by the power produced by the turbines in isolation. By normalizing mechanical power by inflow velocity, we control for small fluctuations in the free stream velocity caused by the flume pumps. The method for estimating uncertainty in κ is described in Appendix B.

3.2.3 Turbine Control

At each of the locations illustrated in Figure 3.2a, two control strategies are optimized to maximize array power output for both co- and counter-rotating turbines.

Tip-speed ratio control is characterized by the tip-speed ratio or non-dimensional rotation rate, given by

$$\lambda = \frac{\omega R}{U_\infty}, \quad (3.3)$$

where R is the turbine radius. For tip-speed ratio control, the rotation rates of each of the rotors are optimized simultaneously, assuming a nominally constant free-stream velocity. As for C_P , the tip-speed ratio for both turbines (λ_1, λ_2) is normalized using the undisturbed freestream inflow velocity (U_∞). The *local* tip-speed ratio, normalized by the velocity incident on the turbine, is higher for the fixed turbine when the mobile turbine is in-line and the wake impinges directly on the fixed turbine. The objective of this control scheme is to optimize rotor operation in the mean flow field induced by the rotors. This optimization problem is non-trivial, as there is a co-dependence between turbine rotation rates, the resulting flow field, and the array power output.

In *coordinated control*, the angular velocities, or tip-speed ratios, of the two turbines are locked to the same value ($\lambda = \lambda_1 = \lambda_2$). For co-rotating turbines, the angular blade offset, or phase difference, is

$$\phi = \theta_1 - \theta_2 \quad (3.4)$$

where θ_1 and θ_2 are the angular positions of mobile and fixed turbines respectively and $\theta = 0$ when one blade is pointing directly upstream. For counter-rotating turbines, the phase difference is defined as

$$\phi = -\theta_1 - \theta_2. \quad (3.5)$$

This results in another two-parameter optimization, this time of λ and ϕ . A closed-loop controller is used to maintain a constant ϕ while testing a specific pair of parameter values.

Optimization of the control parameters is performed using the Nelder-Mead (downhill

simplex) algorithm [101]. This search algorithm is a gradient-free method that evaluates an objective function at the vertices of a simplex of in the parameter space and uses a decision tree to walk towards an optimal solution. To improve the convergence rate of the optimization, initial simplex values are chosen, in part, based on data already collected. The first simplex point is the optimum control set point found for any array configuration that has previously been tested. The second point is the optimum control set point of an array configuration with an upstream turbine position within $1D$ of the current mobile position. If that control set point is not suitably different from the first simplex point a psuedo-random second simplex vertex is chosen, ensuring sufficient distance in the parameter space. The third simplex point is always psuedo-random, but with a lower bound set so that the initial simplex volume is sufficiently large. Optimization is halted when the simplex reached a minimum volume or if 30 control set points have been tested since the last improvement in array power output. Figure 3.3 shows how the optimization was performed for each case. For a given geometry, control strategy, and rotation scheme (Fig. 3.3a), the array is tested for 25 seconds. Performance is computed (Fig. 3.3b) and the Nelder-Mead optimizer is used to determine the next point. This is repeated until convergence (Fig. 3.3c). The performance resulting from the optimized control parameters is shown as a heat map for each array geometry in Figure 3.3d.

3.3 Results and Discussion

3.3.1 Co-rotation

The interaction factor for co-rotating turbines under tip-speed ratio control and coordinated control is shown in Figure 3.4. Both turbines are rotating counterclockwise in the co-rotating arrays. Interaction factors values greater than 0 (red) represent configurations where the array is outperforming the two turbines in isolation. For both cases, when turbines are in similar Y/D locations, performance suffers, but when they are in dissimilar Y/D locations, performance is similar to or greater than isolated turbines. The co-rotating tip-speed ratio

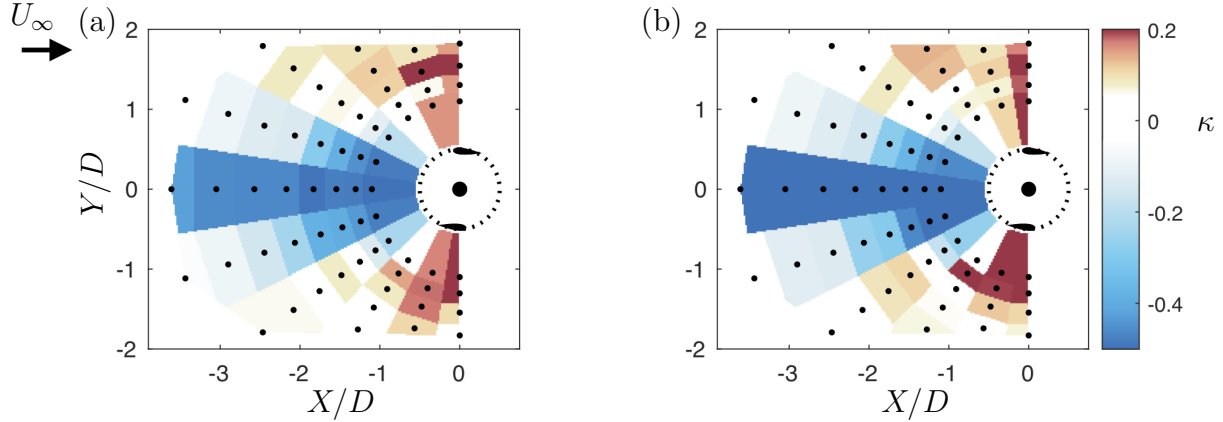


Figure 3.4: Interaction factor for co-rotating arrays under (a) tip-speed ratio control and (b) coordinated control. The fixed turbine is located at $Y/D = X/D = 0$ and the black dot denotes the location of the upstream turbine. Array performance greater than 1 (red) indicates that the array is outperforming the turbines in isolation. The color map contains a dead-band surrounding unity to highlight variations and interaction factor uncertainty of ± 0.039 .

control case had the highest interaction factor of any array geometry, control type, and rotation direction, with turbines situated side-by-side with a $1.1D$ center-to-center spacing achieving a non-dimensional array performance of 1.3 (i.e., a 30 % increase in power output relative to turbines in isolation).

To better understand the drivers for interaction factor, we first consider the optimized tip-speed ratios for the upstream and downstream turbines (Figure 3.5). For all configurations, the upstream turbine optimization converges to a relative narrow range of λ ($1.5 \leq \lambda \leq 2$). Since the optimal λ in isolation is $1.7 - 1.9$ (Fig. 3.2), this suggests that the upstream turbine is relatively unaffected by the presence of the downstream turbine. The downstream turbine optimization converges to a similar λ when the turbines are at dissimilar Y/D locations. When the downstream turbine is located at $Y/D = 0$, its optimized tip-speed ratio trends toward zero as $X/D \rightarrow 0$. This is likely a consequence of the reduction in local inflow velocity when the downstream turbine is located directly in the wake of the upstream turbine and in such close proximity relatively little wake recovery can occur [6]. In addition, this reduction in inflow velocity reduces the local Reynolds number, reducing C_P at a given λ [98].

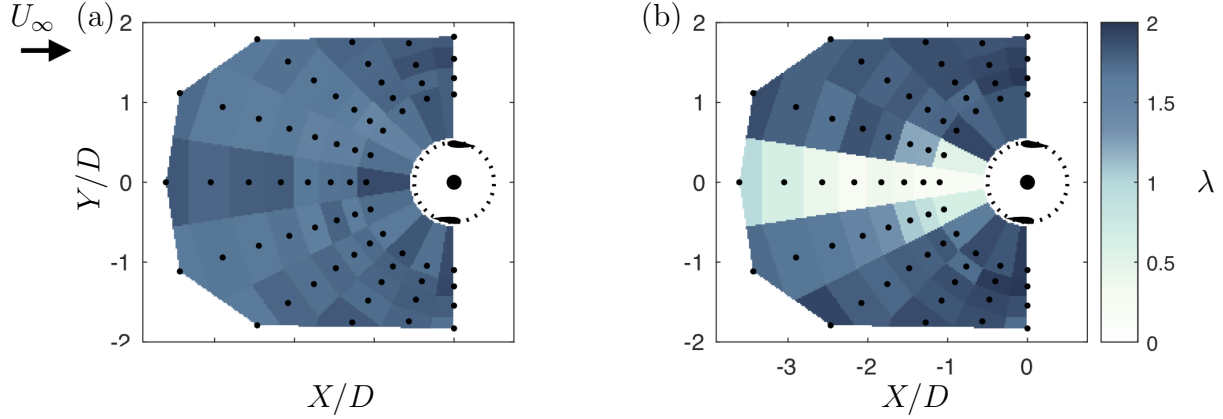


Figure 3.5: Optimized parameters for co-rotating turbines under tip-speed ratio control. (a) Upstream turbine tip-speed ratio for a given array layout and (b) downstream turbine tip-speed ratio.

The optimized parameters for coordinated control are shown in Figure 3.6. The optimal values of λ exhibit a qualitatively similar pattern to the downstream turbine under tip-speed ratio control (i.e., primarily affected when the turbines are at similar Y/D locations). There is, however, evidence of radial striations in the optimal phase. Since, at these tip-speed ratios, the upstream turbine is periodically shedding vortices [137], these radial striations could be indicative of meaningful interaction (interception or avoidance) with the coherent structures from the upstream turbine.

For the coordinated control parameters (shown in Appendix A.1, Fig. 3.6), the optimized tip-speed ratio for the array has a qualitatively similar pattern to that of the downstream turbine in tip-speed ratio control (Fig. 3.5b). However, the optimized tip-speed ratio values are greater than for tip-speed ratio control when the turbines are in similar Y/D positions. This trend matches the results for the counter-rotating tip-speed ratio control case that is explored in the next section.

3.3.2 Counter-rotation

For the counter-rotating arrays, the upstream (mobile) turbine is rotating clockwise while the downstream (fixed) turbine remains rotating counterclockwise. Interaction factor trends

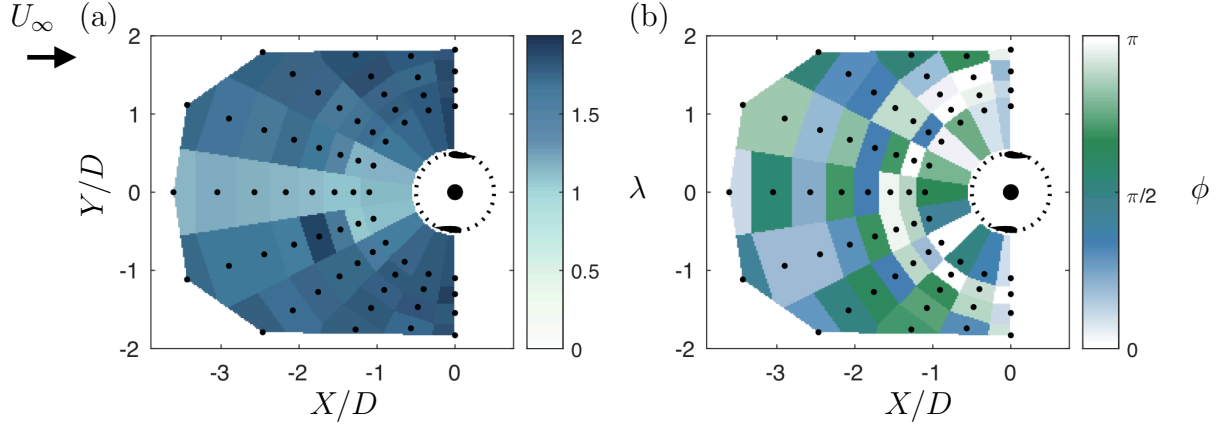


Figure 3.6: Optimized parameters for co-rotating turbines under coordinated control: (a) optimized tip-speed ratio for the turbine pair and (b) optimized phase difference.

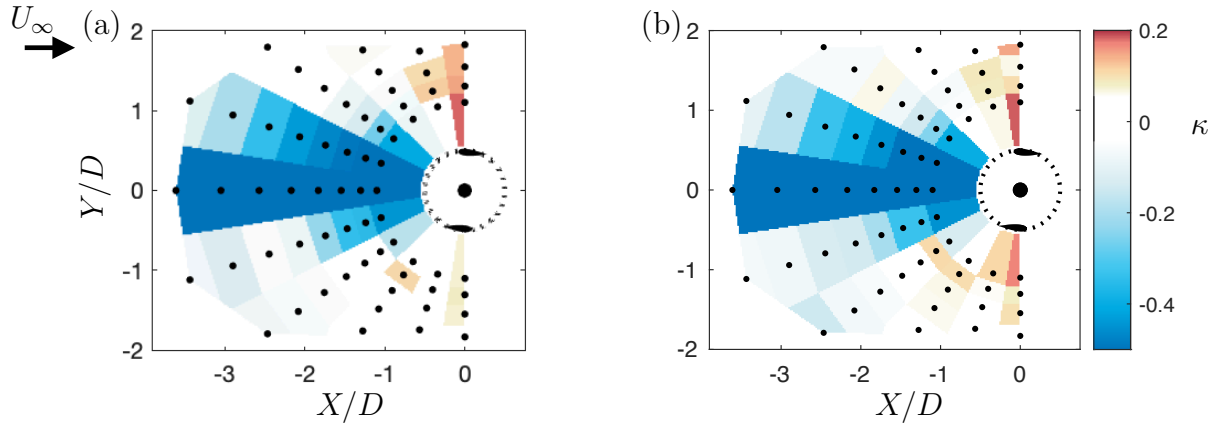


Figure 3.7: Interaction factor for counter-rotating arrays under (a) tip-speed ratio control and (b) coordinated control.

for these arrays, shown in Figure 3.7, are similar to the co-rotating arrays with interaction factors improving when the turbines are at similar X/D positions and suffering when they are at similar Y/D values.

The counter-rotating tip-speed ratio control optimization values (Appendix A.1, Fig. A.1) are also similar to the co-rotating values (Fig. 3.5). Specifically, for the upstream turbine, the optimized λ are all similar to the isolated optima. When the turbines are located in the same Y/D position, the downstream turbine is optimized to operate at a lower λ . The coordinated control values (shown in Appendix A.1, Fig. A.2) have a qualitatively similar

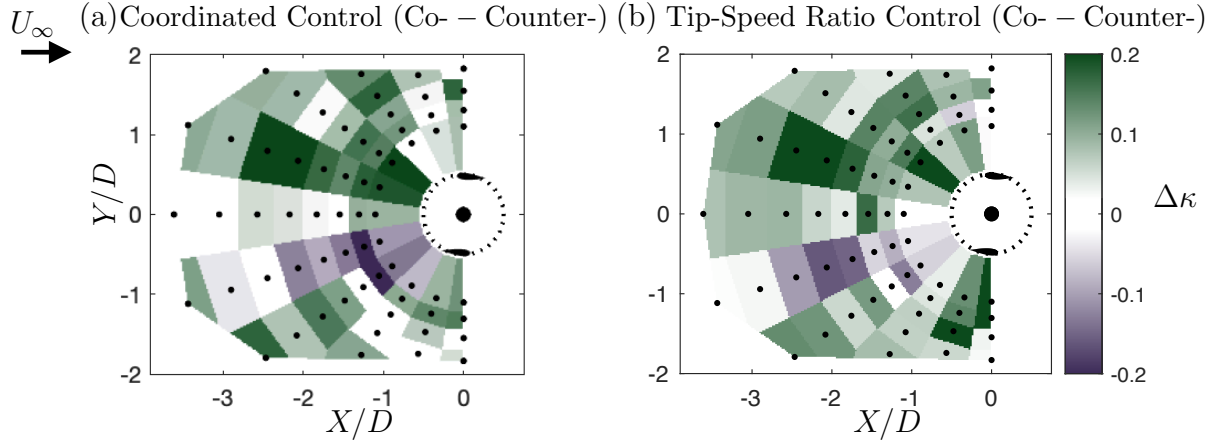


Figure 3.8: Interaction factor comparison ($\Delta\kappa = \kappa_{\text{co-}} - \kappa_{\text{counter-}}$) of co-rotating versus counter rotating arrays for (a) coordinated control and (b) tip-speed ratio control. Green denotes a co-rotating array out-performing its counter-rotating counterpart.

pattern to the co-rotating coordinated control values (Fig. 3.6). The tip-speed ratios are 1.7 – 1.9 when the turbines are in dissimilar Y/D locations, trend lower in similar Y/D locations, and the phase difference shows evidence of radial striations.

3.3.3 Counter- vs. co-rotating arrays

As shown in Figure 3.8, co-rotating arrays generally outperform counter-rotating arrays though the underlying mechanisms causing this trend are not immediately apparent. There is an exception where counter-rotating arrays outperform co-rotating arrays along a single vector. This can be explained by the turbine wake structure. In the $X - Y$ plane, cross-flow turbine wakes have an asymmetric, skewed velocity deficit, as shown in Figure 3.9a [140]. For all cases tested, the fixed turbine (at $X/D = Y/D = 0$) is rotating counterclockwise, while the upstream (mobile) turbine is rotating counterclockwise in a co-rotating array and clockwise in a counter-rotating array.

For the co- and counter-rotating cases, the upstream turbine rotates counterclockwise and clockwise, respectively. Therefore, as illustrated in Figure 3.9, the upstream turbine wake deflection is in the $+Y/D$ direction for co-rotation and in the $-Y/D$ direction for




Figure 3.9: (a) Mean wake velocity of a cross-flow turbine rotating counterclockwise with the skewed velocity deficit denoted with a red arrow [140], (b) deficit path and wake deflection direction for the co-rotating turbine array (c) deficit path and wake deflection direction for the counter-rotating turbine array.

counter-rotation. As a result, co-rotation underperforms counter-rotation in a band of $-Y/D$ upstream turbine positions (purple in Figure 3.8) while counter-rotation underperforms co-rotation for the mirrored $+Y/D$ upstream turbine positions (dark green in Figure 3.8). For both cases, this is consistent with alternating directions of wake velocity skew for counterclockwise and clockwise rotation of the upstream turbine. As turbine separation increases beyond $X/D = -2$, this effect is reduced by wake mixing.

3.3.4 Tip-speed ratio control vs. coordinated control

In Figure 3.10, the two control strategies are compared for co-rotating and counter-rotating arrays. For most array configurations, the differences between the strategies are limited and without obvious spatial trends. However, regardless of rotation direction for the upstream turbine, when the turbines are aligned in the streamwise direction ($Y/D = 0$), tip-speed ratio control generally outperforms coordinated control. Since the downstream turbine operates directly in the wake of the upstream turbine, tip-speed ratio control allows the downstream turbine to operate at a rotation rate with a more favorable *local* tip-speed ratio (i.e., λ defined by the local inflow instead of the freestream flow). Conversely in coordinated control, the turbines are operating at equal rotation rate, which means that the local λ for both turbines is far from optimal.

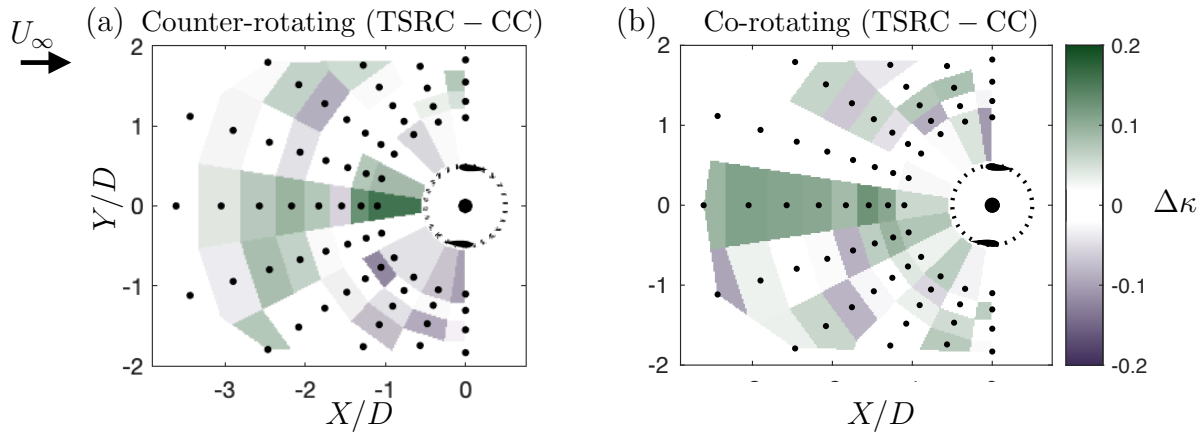


Figure 3.10: Interaction factor comparison ($\Delta\kappa = \kappa_{\text{TSRC}} - \kappa_{\text{CC}}$) of tip-speed ratio control (TSRC) minus coordinated control (CC) for (a) counter-rotating and (b) co-rotating arrays.

3.3.5 Inter-Array and Array-Channel Interactions

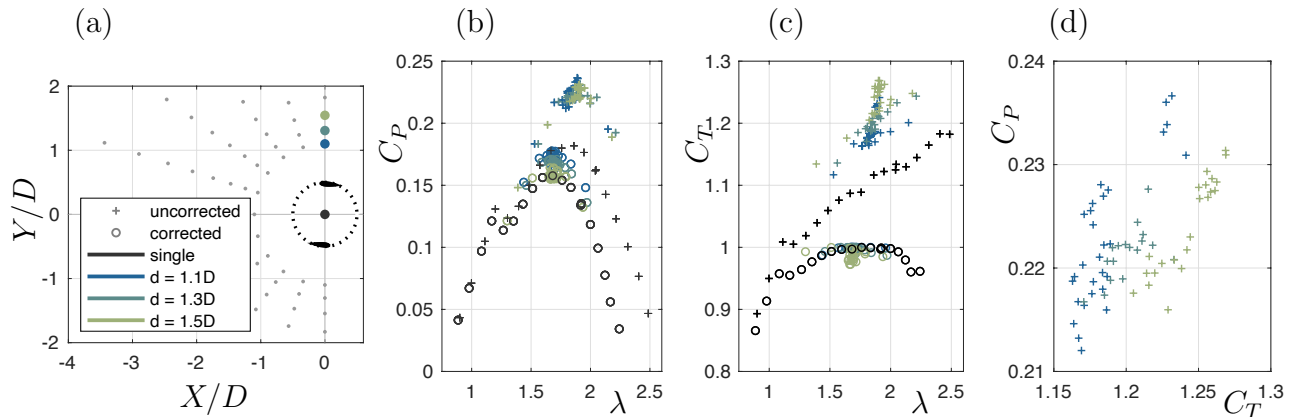


Figure 3.11: Coefficient of performance (b) and coefficient of thrust (c) for the downstream turbine in three side-by-side array configurations (a) for coordinated control optimization. The point clouds correspond to variations in λ and ϕ during the control optimization. C_P , C_T , and λ are corrected according to the method of Barnsley and Wellicome [14].

Our results demonstrate a clear influence of control strategy and array geometry on turbine performance, as quantified by the interaction factor. They do not, however, allow us to fully describe the underlying hydrodynamics. While optical techniques, such as in-rotor particle

image velocimetry [137] might, in the future, provide definitive insight, we can use theoretical arguments to establish the presence of inter-array interactions and their coupling with array-channel energetics. It is well understood that, for a single turbine, geometric confinement increases maximum C_P and the associated λ , as well as the thrust coefficient, C_T . C_T is defined as

$$C_T = \frac{T}{\frac{1}{2}\rho U_\infty^3 HD} \quad (3.6)$$

where T is the thrust force on the turbine in the streamwise direction. The coupling between geometric confinement and C_T augmentation can be explained by linear momentum actuator disc theory, and, with additional approximation, so can changes to C_P and λ [121].

We consider a sub-set of our results in this context: co-rotating coordinated coordinated control with both turbines at $X/D = 0$ and lateral separation, d of $1.1D$, $1.3D$, and $1.5D$ (Figure 3.11a). For these cases, the geometric confinement, defined as the turbines' cross-section area relative to the channel cross-section, is constant at 5.6%. Figure 3.11d shows C_T and C_P for each step in the optimization procedure, with each point representing a discrete value of λ and ϕ for the array. For this specific case ($X/D = 0$, $Y/D > 0$, co-rotation), as separation is increased, C_T decreases for the fixed turbine. This clearly demonstrates an interaction between the rotors that alters the thrust on the fixed turbine at constant geometric confinement. The nature of this thrust-versus-separation interaction varies from case-to-case (co- or counter-rotation, $Y/D >$ or < 0).

To gain further insight into the inter-turbine and array-channel interactions, we apply the method of Barnsley and Wellicome [14] to these data to estimate the unconfined C_P , C_T , and λ for the fixed turbine. This method has been previously shown to predict, with reasonable accuracy, the unconfined performance of a single cross-flow turbine [121]. As shown in Figure 3.11b,c, when the correction is applied to the baseline performance curves (black markers), even modest geometric confinement (2.8%) is predicted to meaningfully alter performance. To apply this same method to the two-turbine array, we assume that the thrust experienced by the two turbines is similar, such that C_T for the fixed turbine

is representative of C_T for the array. The corrected thrust coefficients are estimated to be nearly equal to the isolated turbine (colored circles compared to black circles in Figure 3.11c). We caution against drawing strong conclusions from this result given that, by the nature of the underlying linear momentum theory, unconfined C_T cannot exceed unity, such that we are operating at the margin of theory. Further, given that the asymmetric wake skew for the co-rotating turbines (Figure 3.9), it is possible that the thrust on the mobile turbine differs from the fixed turbine.

This same correction is applied to C_P in Figure 3.11b. As for thrust, the effect of geometric confinement is substantial and, in the absence of any inter-turbine interaction, this would elevate C_P when moving from a single turbine at 2.8% confinement to a two-turbine array at 5.6% confinement. However, the maximum blockage-corrected C_P for the fixed turbine (colored circles) still exceeds blockage-corrected C_P^* (black circles) when this turbine is operated in isolation. This is to say that confinement contributes to the increase in performance, but does not entirely explain it. Further, for this case, there is a clear striation in unconfined $C_P(\phi, \lambda)$ with inter-turbine separation distance. This suggests that the interaction between the two turbines changes their thrust *and* torque. Because the blockage correction attributes an increase in thrust to confinement alone, C_P is corrected farther downward as separation increases due to the rise in C_T , despite the constant confinement.

From this, we can conclude that when turbines are in close proximity, there are important and beneficial mean and periodic interactions that cannot be explained by a simple energetic coupling between the array and channel. Similarly, blockage corrections cannot provide insight into interaction mechanisms for cases other than $X/D = 0$ as they are formulated for rows of turbines with uniform inflow.

Finally, we note that, for these tests, the downstream (fixed) turbine is located at the center of the flume, such that the upstream (mobile) turbine has variable proximity to the flume wall depending on configuration. While this means that wall proximity was not held constant, given that the flume wall is located at $Y/D \approx \pm 6$, variations in the boundary proximity are unlikely to affect turbine performance [62].

3.4 Conclusions

Experimental control and configuration optimization are performed to maximize array power for a pair of cross-flow turbines. The best-performing arrays have a power coefficient 1.3 times greater than for the turbines operating in isolation. Performance is augmented by three main mechanisms: mean flow alteration (e.g., faster bypass flow harnessed by the neighboring rotor), periodic flow alteration, and blockage. Sometimes they are present individually, but often two or more factors simultaneously affect performance.

This work inspires future studies in several directions. First, there has been some visualization of the interaction mechanisms between cross-flow turbines in dense arrays [21], but they have not been fully visualized and compared to measured performance. In conjunction with visualization, modeling using recent advances in data-driven methods [144, 118, 104] would advance understanding of this system. Second, while our results can be provisionally extended to larger arrays, additional interaction mechanisms may occur when the number of rotors increases beyond two. In many applications, arrays of more than two turbines may be desirable. While these experiments consider a two-turbine array, some of the results can be extended to a larger numbers of rotors. Trends are particularly durable for “fences” of adjacent turbines if the free-stream flow speed and direction are uniform across an array. Since performance differences between the two strategies are minimal, coordinated control may be a promising option for this array design because multiple rotors could be mechanically coupled to a single generator, reducing cost. However, if the array spans heterogeneous (e.g. sheared) inflow, tip-speed ratio control would likely be preferred and could be achieved without an inflow measurement using a non-linear control strategy [57]. For homogeneous inflow, in arrays with multiple rows, we see possibilities for extrapolation by taking advantage of periodic interactions that will occur at each row by holding tip-speed ratios constant and repeating phase differences.

Third, in this study we present an online experimental optimization for the control strategy. This approach could be applied to other fluid-structure interactions and flow control

scenarios [23] to explore broad parameter spaces in a time-efficient manner. Similarly, “deep” arrays with multiple rows of turbines could benefit from the experimental approach in field optimization (i.e., each additional turbine adds either a tip-speed ratio or phase to the control scheme). In general, compared to brute force experimentation, an optimization scheme similar to what is described in this work becomes critically important as the number of control variables increases.

Chapter 4

NEAR WAKE DYNAMICS OF A TWO TURBINE ARRAY

As demonstrated in Chapter 3, the best-performing cases from the array optimization corresponded to a side-by-side configuration. This does not, however, provide fundamental insight into the interaction mechanisms present for this case, nor provide a comprehensive understanding of how these interactions are affected by turbine rotation direction and phase offset. Here, we experimentally evaluate the performance and wake of a two-turbine array under a “coordinated control” strategy with a controlled mean phase difference and constant rotation rate. These experiments comprehensively address the potential operational modes and correlations between power production and wake evolution. Specifically, turbine pairs are operated co-rotating, counter-rotating with the blades advancing at the mid-line of the array, and counter-rotating with the blades retreating at the mid-line of the array. First we present the array performance for each of the rotation schemes over the full range of phase differences. We then explore mean wake dynamics across rotation schemes for selected phase differences and compare phase resolved dynamics of the counter-rotating arrays.

4.1 *Methods*

This section outlines the techniques used in this work, including information on turbine dynamics and experimental methods. This builds on the fundamental definitions of turbine performance presented in Chapter 3. The experimental methods describe both performance and particle image velocimetry (PIV) measurements.




Figure 4.1: Diagram of turbine showing the advancing and retreating directions within a rotation.

4.1.1 Cross-flow turbine performance and dynamics

To evaluate the performance of a single turbine, we use the coefficient of performance (C_P), as defined in Chapter 3. To evaluate the performance of the array we utilize the array mean. To evaluate the fluctuations within the array we use a mean subtracted term:

$$\Delta C_P = C_P - \langle C_P \rangle . \quad (4.1)$$

Here, $\langle C_P \rangle$ is the average turbine performance over phase.

Cross-flow turbines have high-dimensional, periodic dynamics [107, 139]. As a turbine rotates, the angles of attack of the blades are constantly changing and structures are periodically shed from the rotor. A single rotation can be bisected into the advancing portion, where the blade is traveling upstream ($\theta = 270^\circ - 90^\circ$), and retreating portion, where the blade is traveling downstream ($\theta = 90^\circ - 270^\circ$). The advancing and retreating directions are labeled in Figure 4.1.

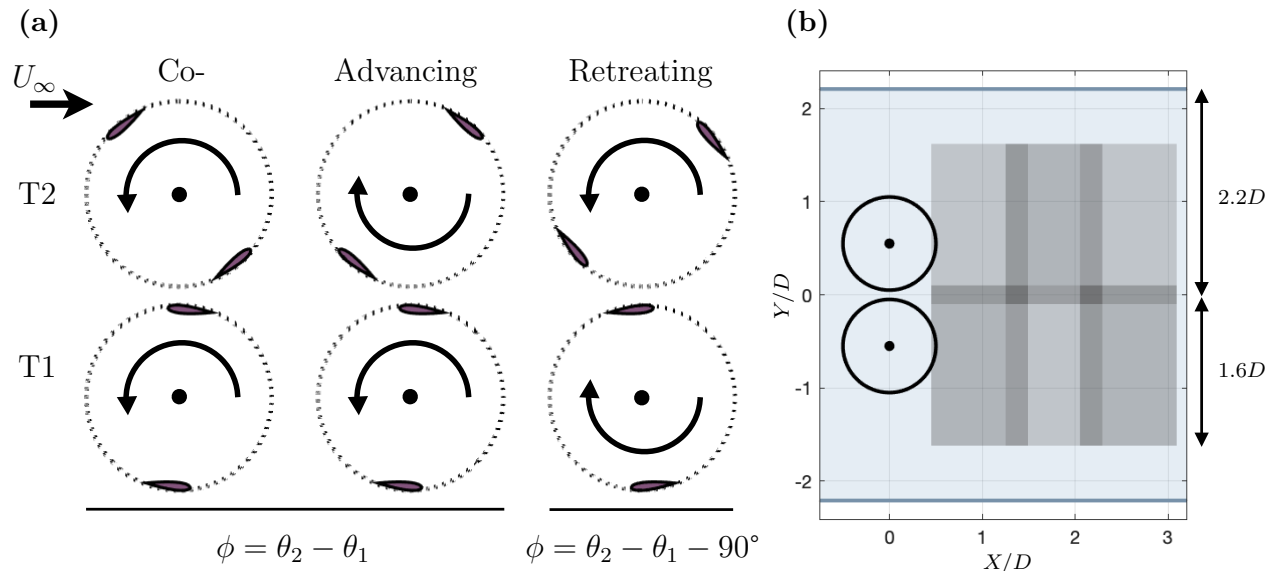


Figure 4.2: (a) Rotation schemes, co-, retreating, and advancing shown with phase difference calculations. Turbine 1 is located at a Y/D position of -0.55 and turbine 2 is located at a Y/D position of 0.55 . For all cases, turbine 1 is at 0° and the phase difference shown is 45° , (b) turbine positions and PIV fields shown in the flume coordinates.

4.1.2 Array Configurations, Control, & Symmetry

The turbines were tested co-rotating (rotating in the same direction) and counter-rotating (rotating in opposite directions). Within a single rotation of a cross-flow turbine, exists inherent asymmetry. For half of the cycle, the blade is advancing and for the other half the blade is retreating. Given the bi-directional travel of a blade over a single rotation, there are two possible counter-rotating fence array configurations. These cases are defined by which direction the blades are traveling at the array midline. For the “retreating” counter-rotating case, the blades are traveling in the same direction as the freestream flow or retreating when they are approaching the array mid-line ($Y/D = 0$). Conversely, for the “advancing” case, the blades are traveling against the direction of the freestream flow or advancing as they approach the mid-line of the array. The three rotation schemes, co-rotation, advancing, and retreating are shown in Figure 4.2(a). The phase difference or the average difference in phase between the two turbines are explicitly defined in this figure and all arrays are shown with

Turbine 1 at a phase of 0° and a phase difference of 45° .

For each of the rotation schemes, performance was measured across phase differences from $\phi = 0^\circ - 180^\circ$.

4.1.3 *Experimental methods*

The data presented in this work was collected in the Alice C. Tyler flume at the University of Washington. This recirculating water channel has a width of 76 cm. The test conditions consisted of a water depth of 52.5 cm, temperature of 23.5°C , and freestream flow velocity (U_∞) of approximately 0.80 m/s. These conditions result in a chord-based Reynolds number of 3.5×10^4 .

Both turbines had two blades with strut end-plates to minimize parasitic losses [141]. Each turbine had a height of $H = 0.23$ m, diameter of $D = 0.172$ m, chord length of $c = 0.04$ m, and symmetric NACA0018 blade profile. For all cases, the center-to-center separation of the rotors is 1.1D.

The turbine assemblies used in this work were identical to the stationary turbine detailed in Chapter 3. The freestream flow velocity was measured using an acoustic Doppler velocimeter (Nortek Vectrino) positioned $X/D = 5$ upstream of the center of the rotors at the array midlines and sampling at a 100 Hz.

4.1.4 *PIV Methods*

The flow fields comprise of two-components of velocity, stream-wise and cross-stream in a 2D plane downstream of the turbine shown in Fig. 4.2b.

The laser used to illuminate the flow is a dual-cavity 30mJ per pulse Nd:YLF laser from Continuum Terra PIV. This laser produced a light sheet with 2 mm thickness. The camera used to acquire the image pairs is a high-speed Vision Research Phantom v641 with a resolution of 2560 x 1600 pixels. The lens used is a 50 mm at an F# of 2. The flow was seeded with 10 μm silver-coated hollow-glass beads which resulted in a particle size of \sim pixels in the images. DaVis LaVision (version 10.1.1) was used for processing the images.

Post-processing the images included background subtraction, masking, cross-correlation, and vector validation. Cross-correlation was performed using a multi-pass algorithm with 75% overlap.

The PIV data field is located at $X/D \approx 0.4 - 3$ and $Y/D \approx -1.6 - 1.6$. Due to the limitations in laser intensity and camera resolution, the data was acquired in a two-by-three grid. The PIV data acquisition was phase-locked with the turbine rotation. Acquisition was triggered when turbine 1 was pointing directly upstream ($\theta = 0^\circ$). Subsequently, 34 image pairs were taken over the course of a full turbine 1 rotation. The presented PIV data is taken at the midspan of the turbines' blades. To produce the composite fields, the phase-averaged fields from the six individual phase-locked acquisitions are stitched together. Motorized rails were used to precisely translate cameras in the cross-stream direction. To acquire the downstream fields of view, the turbine assemblies were moved upstream.

Given the data intensive nature of PIV, the wakes were measured for a subset of phase differences. For both counter-rotating cases, the phases measured were $\phi = 0^\circ, 45^\circ, \& 90^\circ$. For the co-rotating cases, the phases measured were $\phi = 0^\circ, 45^\circ, 90^\circ, \& 135^\circ$. Due to the inherent symmetry of the two-bladed turbine array, the counter-rotating $\phi = 135^\circ$ cases are identical to the $\phi = 45^\circ$ cases mirrored vertically and were not measured.

The presented vorticity fields were calculated from the streamwise and cross-stream velocity fields using a second-order difference method [4].

4.1.5 Array Performance

To quantify array performance, we are using the average turbine performance. In past work, array performance is normalized by individual turbine performance. We are not performing this normalization because the channel energetics change significantly when increasing the blockage by adding the additional rotor.

To determine the optimal array tip-speed ratio for the PIV data acquisition, we tested a coarse grid of phase differences at increasing tip-speed ratios. These cases are shown in Figure 4.3. Phase differences are tested in the range of $\phi = 0 - 180^\circ$. Given the inherent

symmetry of a two-bladed cross-flow turbine, the $\phi = 180 - 360^\circ$ range is identical to the tested range. The change in optimal tip-speed ratio between an individual turbine and the array is due to the change in blockage of the channel between the two cases [120]. The individual turbine occupies approximately 10% of the channel cross-section and its peak efficiency is at a tip-speed ratio of $\lambda = 2$. The array occupies 20% and Figure 4.3 shows that the optimal average array performance occurs at $\lambda = 2.4$. This higher tip-speed ratio of $\lambda = 2.4$ was used for the wake measurements and fine resolution performance measurements across phases presented in this chapter. For a single turbine, blockage tends to increase the optimal tip-speed ratio in addition to increasing the turbine performance [121]. We demonstrate a similar trend when increasing blockage by adding an additional rotor.




Figure 4.3: $C_P - \phi$ plot for the retreating case at $\lambda = 2.1 - 2.6$.

4.2 Results and Discussion

4.2.1 Array Performance Results

The turbine array performance is shown in Figure 4.4. The array average performance for each rotation scheme is shown in Figure 4.4a and mean subtracted performance for individual turbines and the arrays are shown in Figure 4.4b-d for each of the rotation schemes at all phase differences. The phase differences affect both individual turbine and array average performance. Since the counter-rotating cases are equal and opposite about $\phi = 90^\circ$ we focus

on variations in the $\phi = 0 - 90^\circ$ range. On the whole, we see increased absolute performance for the retreating case over the advancing and co-rotating cases. One possible mechanism for this performance increase is a bulk flow interaction with the blades. In the retreating case, the flow is accelerated in the direction of motion of the blades, likely resulting in a higher torque when the blades are traveling towards the array midline. This increased torque would then increase the array-average efficiency. The performance discrepancy shown here is somewhat mirrored in the results in Chapter 3. In the optimization, we saw that the performance for the advancing case dropped off with increased rotor separation. However, the vastly different performance results between the counter-rotation cases seen in this chapter require further exploration. Possible differences that could contribute to performance variations between these experimental campaigns include differences in confinement, non-uniform inflow in the channels, and experimental uncertainty relating to measurement. The retreating rotation scheme has the most pronounced phase dependence for both individual turbines and the array average. The optimal performance is at a phase difference of 90° and the $0 - 5^\circ$ range is the poorest performing. Whereas for an individual turbine (Turbine 1), the performance peaks at approximately 110° and is the lowest at approximately 20° .

The other two cases, co-rotating and advancing have lower magnitude array performance variations due to phase. The advancing case has fairly flat average array performance. However, for the individual turbines we see performance variation at a phase difference of approximately 15° . At $\phi = 15^\circ$ Turbine 1 performance increases and the performance for Turbine 2 decreases, balancing each other out. The co-rotation case shows fairly flat average performance with slight performance increases between $\phi = 110^\circ$ and 150° and slight performance decreases between 20° and 60° . These individual turbine performance variations demonstrate that the phase invariant average performance is a consequence of the individual turbine performances balancing one another, not a lack of interaction between turbines.

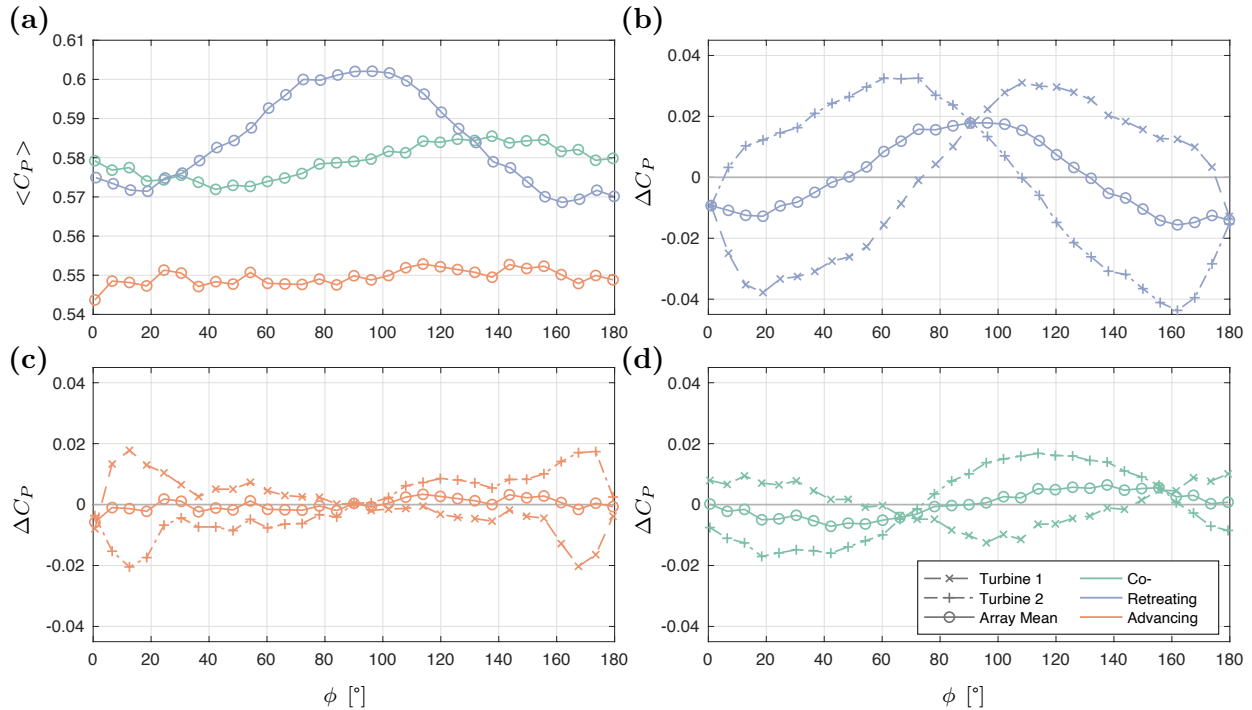


Figure 4.4: (a) Array average performance for all rotation schemes (b) Mean-subtracted average performance for the retreating rotation scheme (c) Mean-subtracted performance for the advancing rotation scheme (d) Mean-subtracted performance for the co-rotation scheme.

4.2.2 Time-average flow field evolution

The time-average wake velocity magnitude for each of the rotation schemes and the subset of phase differences are showed in Figure 4.5 a-j and wake profiles at specific cross-sections ($X/D = 1, 2,$ and 3) are shown in Figure 4.5 k-f. Broadly, each of the wakes has similar structure. As demonstrated by the profiles, they are comprised of (1) a wake deficit that skews moderately towards the advancing side of the rotor, (2) moderately accelerated flow at the array mid-line ($Y/D = 0$), and (3) strongly accelerated flow in the bypass regions. The flow acceleration is due to the bluff-body affects where the flow is diverted around the turbine rather than passing through. This acceleration is augmented by the channel blockage of the array (22%). In all cases, the largest velocity deficit is present at approximately $2.5D$ downstream. The location of maximum deficit is consistent with the past literature [6].

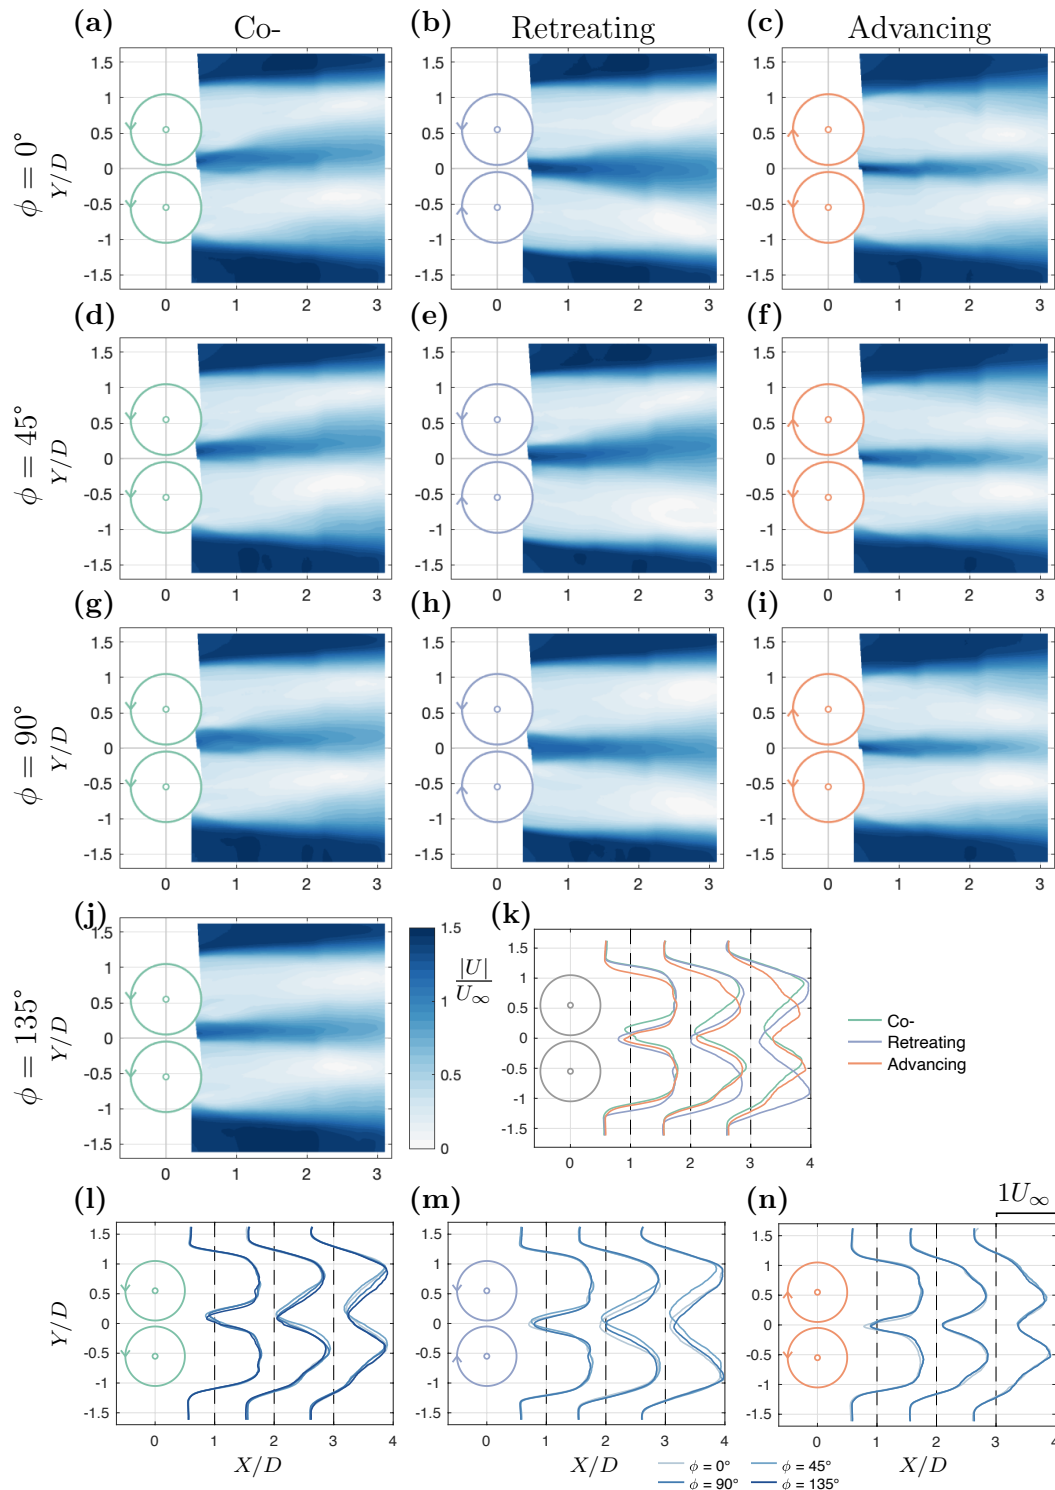


Figure 4.5: (a-j) Average flow field magnitude for each rotation scheme and phase difference measured (k-n) Average flow field magnitude contours for a 90° phase difference for each rotation scheme (k), all co-rotating phase differences (l), all retreating counter-rotating phase differences (m), and all advancing counter-rotating phase differences (n); for each of the profiles, the contour represents the velocity magnitude present at $X/D = 1, 2, \text{ or } 3$. The curves are located such that the dotted lines are at $\frac{|U|}{U_\infty} = 1$ and the spacing between the dotted lines ($1 X/D$) represents $\frac{|U|}{U_\infty} = 1$.

Figure 4.5k shows wake profiles for each of the rotation schemes averaged over all the phases differences. In this plot, there are evident differences in magnitude of the impingement region (i.e., the accelerated flow between the turbines) and its evolution as the wake mixes and skews. The relative locations of the peaks and troughs of the contours are the clearest indicator of the skew direction and are most evident at $X/D = 3$.

The commonalities in general structure across cases can be contrasted with differences both between rotation schemes and phase differences. The width, magnitude, and diffusion of the impingement jet, which is nearly independent of rotation direction and phase difference at $X/D = 1$, notably varies further downstream.

The advancing case (Fig. 4.5c, f, i), produces the narrowest jet and the jet diffuses most rapidly. The wake profile contours (Fig. 4.5k, n), show convergent characteristics where the local minima in wake velocity converges towards $Y/D = 0$ with increasing X/D . The mean wake evolution is independent of phase difference. This uniformity is likely a consequence of the absence of direct interaction between the energetic wake structures on the retreating sides of the individual turbines. The variations in the less energetic structures on the advancing sides of the turbines are not significantly affecting these averaged fields. Figure 4.6 shows the smaller scale structures interacting on the advancing side and minimal interaction on the retreating side of the rotation for the two turbines.

For the retreating case, the impingement jet (Fig. 4.5b, e, h) is widest and diffuses more slowly than the advancing case. The retreating (Fig. 4.5k, m) profiles show divergent characteristics where both turbine wakes skew away from the array midline. Unlike the advancing case, the mean wake evolution depends on the phase difference. At the 0° and 90° phase differences, the wake evolution is symmetric about the midplane, which is consistent with the inherent cross-stream symmetry of the array. Conversely, in the 45° case, the wake skews towards the positive Y/D direction. If we tested the 135° case, we would expect the skew to go toward the negative Y/D direction because the wake dynamics should be equal and opposite about the Y/D axis. This skew can also be seen in the phase averaged fields in Figure 4.6.

The wake evolution in the co-rotating case is an amalgam of the characteristics of the advancing and retreating counter-rotating cases. The magnitude of the impingement jet and its diffusion is similar to the retreating case, but narrower in width. The contours of the co-rotation case (Fig. 4.5k, l) shows both of the wakes skewed in parallel towards the positive Y/D direction. Variation in phase difference produces changes in the wake evolution at $X/D > 2$, but these are less pronounced than for the retreating case.

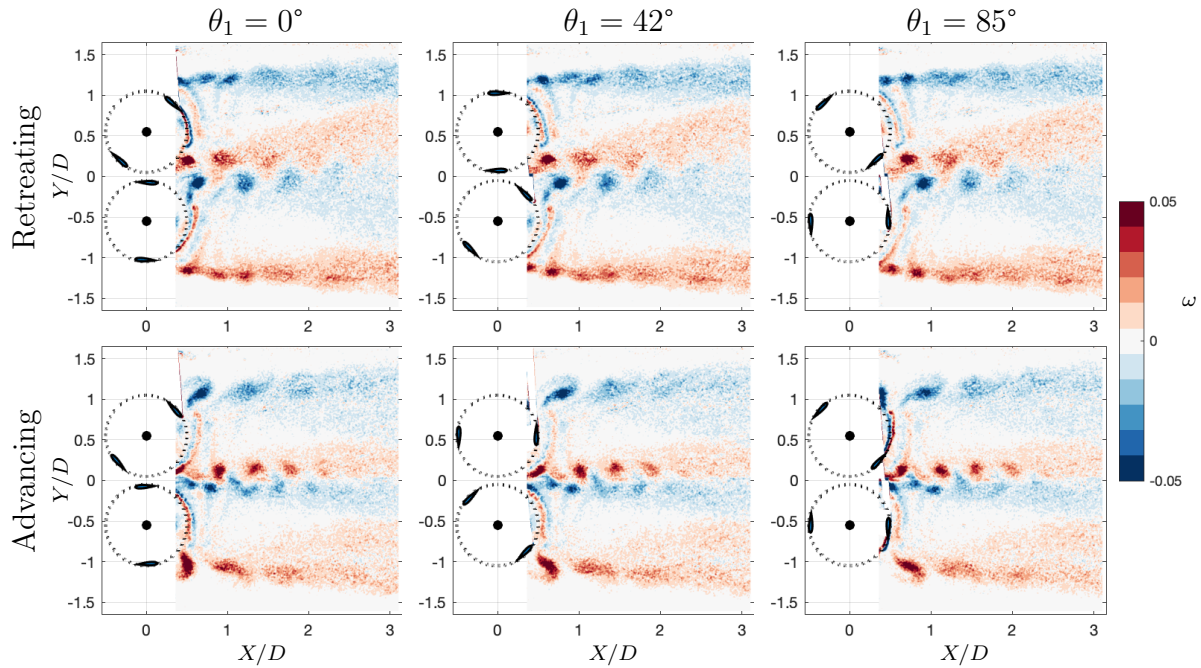


Figure 4.6: Selected phase averaged frames for the advancing and retreating counter-rotating cases at turbine 1 azimuthal positions of $\theta_1 \approx 0^\circ$, 42° , or 85° .

Further insight is gained by considering the phase-average flow fields. Figure 4.6 shows vorticity at three blade positions for the two counter-rotating schemes, both operating with a phase differences of 45° . These frames highlight alterations in the vortex shedding patterns between rotation schemes. For the advancing case, the power stroke (the region where the turbine produces most of its power) is occurring adjacent to the bypass wake region. For the retreating cases, the power stroke is occurring closer to the array mid-line, where the rotors are directly interacting. These interactions and how they vary with phase difference have a

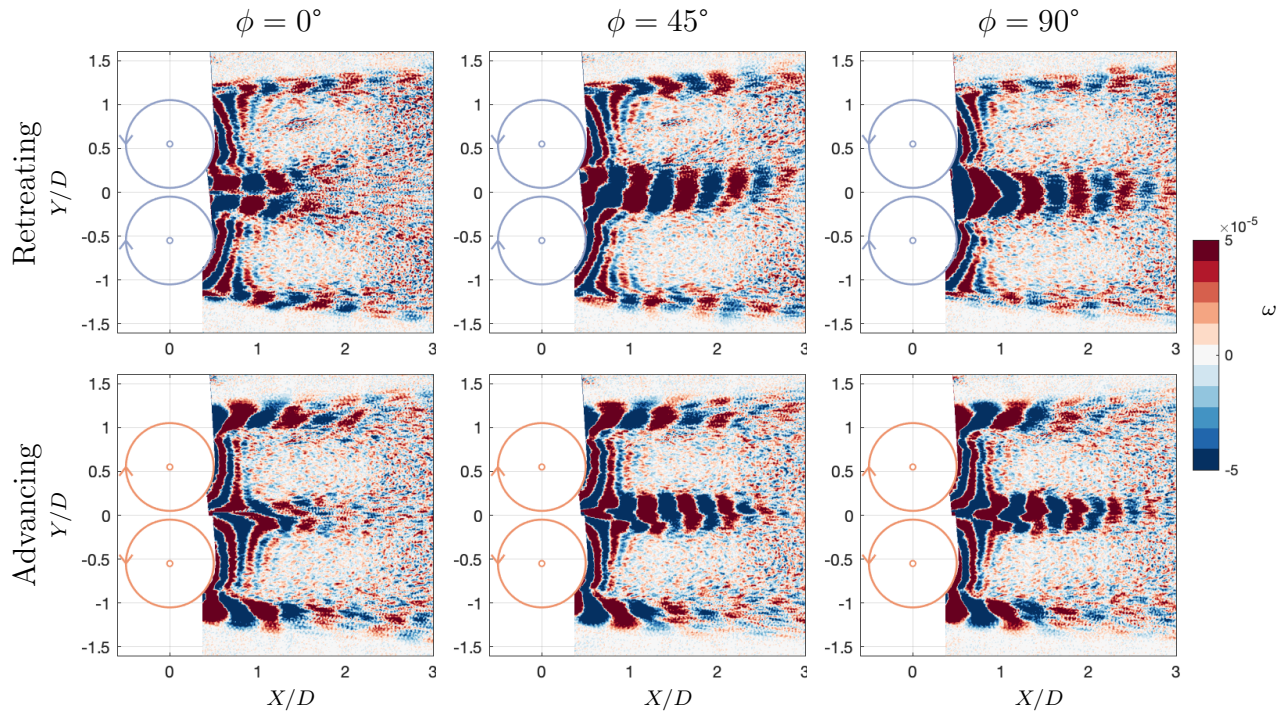


Figure 4.7: Second POD mode for the advancing and retreating counter-rotating cases at each phase difference.

clear dependence on performance. However, it would be speculative to attribute performance variations to variation in wake structure.

These wake structures are mirrored in the second proper orthogonal decomposition (POD) modes shown in Figure 4.7 (POD is reviewed extensively in Chapter 2). In this figure, the retreating case midline structures show the larger scale advected structure and the advancing case shows that the dominant dynamics at the array midline are the blade wake advecting downstream. Additionally, it can be seen in both cases that the structures evolve with phase difference, albeit in a non-linear manner, with a more pronounced difference between the 0 and 45 ° phase offsets than the 45 and 90 ° offsets. For both rotation directions, at a phase offset of 0°, the most energetic coherent structures from each are equal and opposite and shed simultaneously (centerline for retreating case, wake bypass for advancing case), whereas the at 90° phase offset, equal structures are shed alternately.

4.3 Conclusions

In this chapter, we present comprehensive performance and wake measurements for the three configurations of cross-flow turbine pairs. For each rotation scheme, the turbines are tested across phase differences. We find that the main driver of dynamics and performance variations is the rotation scheme. For each rotation scheme, the dominant dynamics persist across different phase differences. The case in which counter-rotating blades are retreating or traveling downstream at their closest point of approach has the largest phase-difference dependent variations in performance and wake dynamics. For the counter-rotating case where the blades are advancing or traveling upstream at their closest point of approach or co-rotating blades, there are only small phase perturbations in mean array performance and wake evolution. However, in the case of performance, phase variations do contribute to varied power outputs from individual turbines that are balanced out in aggregation at the array level.

Past work has shown similar mean flow-field dynamics in dual-rotor arrays [111]. It remains a challenge to draw connections between the wake evolution and performance variations. We speculate that there are two main mechanisms causing variation between rotation schemes and across phases. These mechanisms include bulk flow interactions between the array and channel and blade level flow-field alterations, such as the azimuthal position where stall occurs. These two mechanisms are intertwined in these results. Future work on a single-bladed turbine array utilizing in-rotor flow field measurements and simultaneous force measurements could illuminate some of these dependencies. Additionally, the wake dynamics could be further explored to understand more of the time and phase dependent fluid mechanisms present and how they change with rotation scheme and phase difference. Three-dimensional effects could also be examined by measuring spanwise velocity as well as collecting tomographic measurements to understand edge effects in the wake.

Chapter 5

CONCLUSION

Cross-flow turbine arrays are a promising solution for renewable energy generation in wind, marine, and riverine currents. Data-driven methods can be used to better understand, model, and interpret their dynamics. This is first demonstrated using RPCA to improve flow field fidelity. In Chapter 2, we show that RPCA can identify and replace spurious and missing vectors in flow fields more reliably than existing methods. This allows improved flow-field modeling and interpretation. We find that RPCA can be a valuable algorithm in the arsenal of PIV processing and filtering techniques, particularly when the processing pipeline culminates in modal analysis

In Chapter 3, we demonstrate that with well-considered geometries and control strategies, cross-flow turbines can produce more power in dense arrays than equivalent, non-interacting turbines. Array performance is shown to depend on geometry and control strategy. Turbine arrays perform best when they are arranged in a fence configuration (i.e., side-by-side turbines). Even when corrected for variations in flow confinement, performance is still augmented relative to non-interacting turbines. In addition, while the downstream turbine is able to interact with coherent structures shed by the upstream turbine, any benefits of such interaction are likely outweighed by reductions in the mean flow velocity at the downstream turbine.

Chapter 3 explores the best-performing geometric configuration in greater detail, comprehensively evaluating performance the wake evolution instead of seeking an optimal performance point. We find that cross-flow turbine performance and wake dynamics depend on rotation scheme and phase difference. When the turbine blades are retreating as they approach the array midline, the time-average array efficiency is highest when the turbines

are perfectly out of phase with one another. For the retreating rotation scheme, the large scale wake structures shed from the individual turbines during dynamic stall occur close to the array midline where the turbines blades are directly interacting. This interaction is hypothesized to cause the increased variability in both array performance and coherent flow structures propagating in the wake. Additionally, the second PCA mode for the counter-rotating retreating case shows that the flow patterns at the midline are dominated by the large shed vortices. Conversely, for all measured phase differences of the advancing case (i.e., blades advancing as they approach the array midline) the structures at the midline are dominated by the blade wake advecting downstream. The flow fields and performance for the advancing case have less variation across phase likely due to these less energetic turbine structures interacting. Similarly, the co-rotation case sees little variability in performance and mean wake across phases.

5.1 Future work

There are endless opportunities for data-driven methods to improve interpretation and acquisition of data from fluid mechanics. Exploration and uniform sampling of parameter spaces is an enormous source of bottleneck in research. When we employ optimizations we speed up acquisition but we do not come away with parametric interdependence resolved. For example, the experimental time required to combine the detailed flow field and performance studies in Chapter 3 with the range of array geometries presented would be entirely infeasible without a data-driven approach. Active learning and parameter modeling provide promising alternatives to uniform sampling and optimizations. These methods could provide an increase in efficiency in experimental acquisition in a wide array of other fields of study. Knowledge about the impact of different variables on underlying dynamics can be as valuable as the optimal solution for system. Recent work has shown that Gaussian process regression can be used to explore system dynamics in a variety of fields. These include, vortex-induced vibrations [53], phytoplankton sampling in the ocean [59], spectroscopy [34], and learning partial differential equations [116].

There is also potential for additional analysis on the turbine array wakes. A promising direction that builds on the initial results presented in this dissertation would be to apply Dynamic Mode Decomposition (DMD) to better understand the inherent flow patterns present. These flow fields could also be evaluated for their coherent structures and the vortex formation patterns could be characterized, similar to analysis previously presented for a single turbine [142]. Additionally, to better understand how the wakes of the individual turbines interact with one another, the forward and backward finite-time Lyapunov exponents could be calculated to uncover the convergent and divergent manifolds. The work presented in this thesis utilizes a single rotor geometry. The array interactions presented may be generalizable to other rotor shapes (e.g. chord-to-radius ratio, blade number, blade shape) and flow conditions (e.g. Reynolds number or blockage). Alternatively, changing the flow conditions could unearth dynamics that have yet to be identified.

More broadly there are additional improvements that can be made to how we interpret experimentally acquired data. There are a wealth of techniques being developed to uncover dynamics. These techniques are largely tested on canonical examples and simulations. Additional care and data processing often must be taken when extending these methods to experimentally collected data.

BIBLIOGRAPHY

- [1] R. J. Adrian. Particle-imaging techniques for experimental fluid mechanics. *Annual review of fluid mechanics*, 23(1):261–304, 1991.
- [2] R. J. Adrian. Twenty years of particle image velocimetry. *Exp. Fluids*, 39(2):159–169, 2005.
- [3] R. J. Adrian and J. Westerweel. *Particle image velocimetry*. Cambridge University Press, 2011.
- [4] Ronald J Adrian and Jerry Westerweel. *Particle image velocimetry*. Number 30. Cambridge University Press, 2011.
- [5] Mojtaba Ahmadi-Baloutaki, Rupp Carriveau, and David SK Ting. A wind tunnel study on the aerodynamic interaction of vertical axis wind turbines in array configurations. *Renewable Energy*, 96:904–913, 2016.
- [6] Daniel B Araya, Tim Colonius, and John O Dabiri. Transition to bluff-body dynamics in the wake of vertical-axis wind turbines. *Journal of Fluid Mechanics*, 813:346–381, 2017.
- [7] Travis Askham and J Nathan Kutz. Variable projection methods for an optimized dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 17(1):380–416, 2018.
- [8] Nadine Aubry, Philip Holmes, John L Lumley, and Emily Stone. The dynamics of coherent structures in the wall region of a turbulent boundary layer. *Journal of Fluid Mechanics*, 192:115–173, 1988.

- [9] Peter Bachant and Martin Wosnik. Characterising the near-wake of a cross-flow turbine. *Journal of Turbulence*, 16(4):392–410, 2015.
- [10] Peter Bachant and Martin Wosnik. Effects of reynolds number on the energy conversion and near-wake dynamics of a high solidity vertical-axis cross-flow turbine. *Energies*, 9(2):73, 2016.
- [11] Peter Bachant, Martin Wosnik, Budi Gunawan, and Vincent S Neary. Experimental study of a reference model vertical-axis cross-flow turbine. *PloS one*, 11(9):e0163799, 2016.
- [12] Shervin Bagheri. Koopman-mode decomposition of the cylinder wake. *Journal of Fluid Mechanics*, 726:596–623, 2013.
- [13] Shervin Bagheri. Effects of weak noise on oscillating flows: Linking quality factor, Floquet modes, and Koopman spectrum. *Physics of Fluids*, 26(9):094104, 2014.
- [14] MJ Barnsley and JF Wellicome. Final report on the 2nd phase of development and testing of a horizontal axis wind turbine test rig for the investigation of stall regulation aerodynamics. carried out under etsu agreement e. *Carried Out Under ETSU Agreement No. E. A*, 5, 1990.
- [15] Zachary P Berger, Patrick R Shea, Matthew G Berry, Bernd R Noack, Sivaram Gogineni, and Mark N Glauser. Active flow control for high speed jets with large window piv. *Flow, Turbulence and Combustion*, 94(1):97–123, 2015.
- [16] Gal Berkooz, Philip Holmes, and John L Lumley. The proper orthogonal decomposition in the analysis of turbulent flows. *Ann. Rev. Fluid Mech.*, 25(1):539–575, 1993.
- [17] Matthieu Boudreau and Guy Dumas. Comparison of the wake recovery of the axial-flow and cross-flow turbine concepts. *Journal of Wind Engineering and Industrial Aerodynamics*, 165:137–152, 2017.

- [18] T. Bouwmans and E. H. Zahzah. Robust PCA via principal component pursuit: A review for a comparative evaluation in video surveillance. *Comp. Vis. & Image Under.*, 122:22–34, 2014.
- [19] J Bremseth and K Duraisamy. Computational analysis of vertical axis wind turbine arrays. *Theoretical and Computational Fluid Dynamics*, 30(5):387–401, 2016.
- [20] Ian D Brownstein, Matthias Kinzel, and John O Dabiri. Performance enhancement of downstream vertical-axis wind turbines. *Journal of Renewable and Sustainable Energy*, 8(5):053306, 2016.
- [21] Ian D Brownstein, Nathaniel J Wei, and John O Dabiri. Aerodynamically interacting vertical-axis wind turbines: Performance enhancement and three-dimensional flow. *Energies*, 12(14):2724, 2019.
- [22] S. L. Brunton and J. N. Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2019.
- [23] S. L. Brunton and B. R. Noack. Closed-loop turbulence control: Progress and challenges. *Applied Mechanics Reviews*, 67:050801–1–050801–48, 2015.
- [24] S. L. Brunton, J. L. Proctor, J. H. Tu, and J. N. Kutz. Compressed sensing and dynamic mode decomposition. *Journal of Computational Dynamics*, 2(2):165–191, 2015.
- [25] Steven L. Brunton, Bernd R. Noack, and Petros Koumoutsakos. Machine learning for fluid mechanics. *to appear in Annual Review of Fluid Mechanics (arXiv preprint arXiv:1905.11075)*, 52, 2020.
- [26] Abel-John Buchner, Julio Soria, Damon Honnery, and Alexander J Smits. Dynamic stall in vertical axis wind turbines: scaling and topological considerations. *Journal of Fluid Mechanics*, 841:746–766, 2018.

- [27] AJ Buchner, MW Lohry, L Martinelli, J Soria, and AJ Smits. Dynamic stall in vertical axis wind turbines: comparing experiments and computations. *Journal of Wind Engineering and Industrial Aerodynamics*, 146:163–171, 2015.
- [28] Peter A Businger and Gene H Golub. Algorithm 358: singular value decomposition of a complex matrix [f1, 4, 5]. *Communications of the ACM*, 12(10):564–565, 1969.
- [29] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3):11, 2011.
- [30] Kevin Carlberg, Matthew Barone, and Harbir Antil. Galerkin v. least-squares petrov–galerkin projection in nonlinear model reduction. *Journal of Computational Physics*, 330:693–734, 2017.
- [31] Travis J Carrigan, Brian H Dennis, Zhen X Han, and Bo P Wang. Aerodynamic shape optimization of a vertical-axis wind turbine using differential evolution. *ISRN Renewable Energy*, 2012, 2012.
- [32] Marco Raciti Castelli and Ernesto Benini. Effect of blade inclination angle on a darrieus wind turbine. *Journal of turbomachinery*, 134(3):031016, 2012.
- [33] Chein-Chang Chen and Cheng-Hsiung Kuo. Effects of pitch angle and blade camber on flow characteristics and performance of small-size darrieus vawt. *Journal of Visualization*, 16(1):65–74, 2013.
- [34] Tao Chen, Julian Morris, and Elaine Martin. Gaussian process regression for multivariate spectroscopic calibration. *Chemometrics and Intelligent Laboratory Systems*, 87(1):59–71, 2007.
- [35] Wei-Hsin Chen, Ching-Ying Chen, Chun-Yen Huang, and Chii-Jong Hwang. Power output analysis and optimization of two straight-bladed vertical-axis wind turbines. *Applied energy*, 185:223–232, 2017.

- [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] Thomas C Corke and Flint O Thomas. Dynamic stall in pitching airfoils: aerodynamic damping and compressibility effects. *Annual Review of Fluid Mechanics*, 47:479–505, 2015.
- [38] John O Dabiri. Potential order-of-magnitude enhancement of wind farm power density via counter-rotating vertical-axis wind turbine arrays. *Journal of renewable and sustainable energy*, 3(4):043104, 2011.
- [39] Mukul Dave and Jennifer A Franck. Comparison of rans and les for a cross-flow turbine in confined and unconfined flow. *Journal of Renewable and Sustainable Energy*, 13(6):064503, 2021.
- [40] Mukul Dave, Benjamin Strom, Abigale Snortland, Owen Williams, Brian Polagye, and Jennifer A Franck. Simulations of intracycle angular velocity control for a crossflow turbine. *AIAA Journal*, 59(3):812–824, 2021.
- [41] Scott TM Dawson, Maziar S Hemati, Matthew O Williams, and Clarence W Rowley. Characterizing and correcting for the effect of sensor noise in the dynamic mode decomposition. *Experiments in Fluids*, 57(3):1–19, 2016.
- [42] Stefano Discetti, Andrea Natale, and Tommaso Astarita. Spatial filtering improved tomographic PIV. *Experiments in fluids*, 54(4):1505, 2013.
- [43] Minh N Doan, Yuriko Kai, and Shinnosuke Obi. Twin marine hydrokinetic cross-flow turbines in counter rotating configurations: A laboratory-scaled apparatus for power measurement. *Journal of Marine Science and Engineering*, 8(11):918, 2020.

- [44] J Duncan, D Dabiri, J Hove, and M Gharib. Universal outlier detection for particle image velocimetry (PIV) and particle tracking velocimetry (PTV) data. *Measurement Science and Technology*, 21(5):057002, 2010.
- [45] Reeve Dunne and Beverley J McKeon. Dynamic stall on a pitching and surging airfoil. *Experiments in Fluids*, 56(8):157, 2015.
- [46] Karthik Duraisamy, Gianluca Iaccarino, and Heng Xiao. Turbulence modeling in the age of data. *To appear in Annual Reviews of Fluid Mechanics*, 2018.
- [47] Karthikeyan Duraisamy and Vinod Lakshminarayan. Flow physics and performance of vertical axis wind turbine arrays. In *32nd AIAA Applied Aerodynamics Conference*, page 3139, 2014.
- [48] Thomas Duriez, Steven L Brunton, and Bernd R Noack. *Machine learning control-taming nonlinear dynamics and turbulence*, volume 116. Springer, 2017.
- [49] Naveed Durrani, Ning Qin, Harris Hameed, and Shahab Khushnood. 2-d numerical analysis of a vawt wind farm for different configurations. In *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, page 461, 2011.
- [50] Jonathan M Edwards, Louis Angelo Danao, and Robert J Howell. Piv measurements and cfd simulation of the performance and flow physics and of a small-scale vertical axis wind turbine. *Wind Energy*, 18(2):201–217, 2015.
- [51] Jeff D Eldredge and Anya R Jones. Leading-edge vortices: mechanics and modeling. *Annual Review of Fluid Mechanics*, 51:75–104, 2019.
- [52] Sandra Eriksson, Hans Bernhoff, and Mats Leijon. Evaluation of different turbine concepts for wind power. *renewable and sustainable energy reviews*, 12(5):1419–1434, 2008.

- [53] Dixia Fan, Gurvan Jodin, TR Consi, L Bonfiglio, Y Ma, LR Keyes, George E Karniadakis, and Michael S Triantafyllou. A robotic intelligent towing tank for learning complex fluid-structure dynamics. *Science Robotics*, 4(36), 2019.
- [54] Marie Farge and Kai Schneider. Coherent vortex simulation (CVS), a semi-deterministic turbulence model using wavelets. *Flow, Turbulence and Combustion*, 66(4):393–426, 2001.
- [55] Carlos Simao Ferreira, Gijs Van Kuik, Gerard Van Bussel, and Fulvio Scarano. Visualization by piv of dynamic stall on a vertical axis wind turbine. *Experiments in Fluids*, 46(1):97–108, 2009.
- [56] FE Fish and GV Lauder. Passive and active flow control by swimming fishes and mammals. *Annu. Rev. Fluid Mech.*, 38:193–224, 2006.
- [57] Dominic Forbush, Robert J Cavagnaro, and Brian Polagye. Power-tracking control for cross-flow turbines. *Journal of Renewable and Sustainable Energy*, 11(1):014501, 2019.
- [58] JW Foreman Jr, EW George, and RD Lewis. Measurement of localized flow velocities in gases with a laser doppler flowmeter. *Applied Physics Letters*, 7(4):77–78, 1965.
- [59] Trygve O Fossum, Glaucia M Fragoso, Emlyn J Davies, Jenny E Ullgren, Renato Mendes, Geir Johnsen, Ingrid Ellingsen, Jo Eidsvik, Martin Ludvigsen, and Kanna Rajan. Toward adaptive robotic sampling of phytoplankton in the coastal ocean. *Science Robotics*, 4(27), 2019.
- [60] Nobuyuki Fujisawa and Satoshi Shibuya. Observations of dynamic stall on darrieus wind turbine blades. *Journal of Wind Engineering and Industrial Aerodynamics*, 89(2):201–214, 2001.
- [61] Damien Garcia. A fast all-in-one method for automated post-processing of PIV data. *Experiments in fluids*, 50(5):1247–1259, 2011.

- [62] Olivier Gauvin-Tremblay and Guy Dumas. Two-way interaction between river and deployed cross-flow hydrokinetic turbines. *Journal of Renewable and Sustainable Energy*, 12(3):034501, 2020.
- [63] M. Gavish and D. L. Donoho. The optimal hard threshold for singular values is $4/\sqrt{3}$. *IEEE Transactions on Information Theory*, 60(8):5040–5053, 2014.
- [64] G. H. Golub and C. Reinsch. Singular value decomposition and least squares solutions. *Numerical Mathematics*, 14:403–420, 1970.
- [65] Gene Golub and William Kahan. Calculating the singular values and pseudo-inverse of a matrix. *Journal of the Society for Industrial & Applied Mathematics, Series B: Numerical Analysis*, 2(2):205–224, 1965.
- [66] Gene H Golub and Charles F Van Loan. An analysis of the total least squares problem. *SIAM Journal on Numerical Analysis*, 17(6):883–893, 1980.
- [67] Derek G Goring and Vladimir I Nikora. Despiking acoustic doppler velocimeter data. *Journal of Hydraulic Engineering*, 128(1):117–126, 2002.
- [68] J Graham, K Kanov, XIA Yang, M Lee, N Malaya, CC Lalescu, R Burns, G Eyink, A Szalay, RD Moser, and C Meneveau. A web services accessible database of turbulent channel flow and its use for testing a new integral wall model for LES. *Journal of Turbulence*, 17(2):181–215, 2016.
- [69] Frank E Grubbs. Procedures for detecting outlying observations in samples. *Technometrics*, 11(1):1–21, 1969.
- [70] F. Gueniat, L. Mathelin, and L. Pastur. A dynamic mode decomposition approach for large and arbitrarily sampled systems. *Physics of Fluids*, 27(2):025113, 2015.
- [71] Douglas P Hart. PIV error correction. *Experiments in fluids*, 29(1):13–22, 2000.

- [72] Maziar S Hemati, Clarence W Rowley, Eric A. Deem, and Louis N. Cattafesta. De-biasing the dynamic mode decomposition for applied Koopman spectral analysis. *Theoretical and Computational Fluid Dynamics*, 31(4):349–368, 2017.
- [73] TC Hohman, L Martinelli, and AJ Smits. The effects of inflow conditions on vertical axis wind turbine wake structure and performance. *Journal of Wind Engineering and Industrial Aerodynamics*, 183:1–18, 2018.
- [74] P. J. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley. *Turbulence, coherent structures, dynamical systems and symmetry*. Cambridge Monographs in Mechanics. Cambridge University Press, 2nd edition, 2012.
- [75] Harold Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
- [76] H Huang, D Dabiri, and Morteza Gharib. On errors of digital particle image velocimetry. *Measurement Science and Technology*, 8(12):1427, 1997.
- [77] Po-Sen Huang, Scott Deeann Chen, Paris Smaragdis, and Mark Hasegawa-Johnson. Singing-voice separation from monaural recordings using robust principal component analysis. In *IEEE ICASSP*, pages 57–60, 2012.
- [78] Aidan Hunt, Carl Stringer, and Brian Polagye. Effect of aspect ratio on cross-flow turbine performance. *Journal of Renewable and Sustainable Energy*, 12(5):054501, 2020.
- [79] Guoqing Jin, Zhi Zong, Yichen Jiang, and Li Zou. Aerodynamic analysis of side-by-side placed twin vertical-axis wind turbines. *Ocean Engineering*, 209:107296, 2020.
- [80] Mihailo R Jovanović, Peter J Schmid, and Joseph W Nichols. Sparsity-promoting dynamic mode decomposition. *Physics of Fluids*, 26(2):024103, 2014.

- [81] K Karhunen. Über lineare methoden in der wahrscheinlichkeitsrechnung, vol. 37. *Annales Academiæ Scientiarum Fennicæ, Ser. A. I*, 1947.
- [82] Matthias Kinzel, Daniel B Araya, and John O Dabiri. Turbulence in vertical axis wind turbine canopies. *Physics of Fluids*, 27(11):115102, 2015.
- [83] Matthias Kinzel, Quinn Mulligan, and John O Dabiri. Energy exchange in an array of vertical-axis wind turbines. *Journal of Turbulence*, 13(1):N38, 2012.
- [84] J. N. Kutz, S. L. Brunton, B. W. Brunton, and J. L. Proctor. *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM, 2016.
- [85] J Nathan Kutz. Deep learning in fluid dynamics. *Journal of Fluid Mechanics*, 814:1–4, 2017.
- [86] HF Lam and HY Peng. Measurements of the wake characteristics of co-and counter-rotating twin h-rotor vertical axis wind turbines. *Energy*, 131:13–26, 2017.
- [87] A Laneville and P Vittecoq. Dynamic stall: the case of the vertical axis wind turbine. *Journal of Solar Energy Engineering*, 108(2):140–145, 1986.
- [88] J. Lei, S. Liu, X. Wang, and Q. Liu. An image reconstruction algorithm for electrical capacitance tomography based on robust principle component analysis. *Sensors*, 13(2):2076–2092, 2013.
- [89] Zhouchen Lin, Minming Chen, and Yi Ma. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055*, 2010.
- [90] J.-C. Loiseau and S. L. Brunton. Constrained sparse Galerkin regression. *Journal of Fluid Mechanics*, 838:42–67, 2018.

- [91] J.-C. Loiseau, B. R. Noack, and S. L. Brunton. Sparse reduced-order modeling: sensor-based dynamics to full-state estimation. *Journal of Fluid Mechanics*, 844:459–490, 2018.
- [92] E. N. Lorenz. Empirical orthogonal functions and statistical weather prediction. Technical report, Massachusetts Institute of Technology, December 1956.
- [93] J. L. Lumley. *Stochastic Tools in Turbulence*. Academic Press, 1970.
- [94] He Lyu, Ningyu Sha, Shuyang Qin, Ming Yan, Yuying Xie, and Rongrong Wang. Manifold denoising by nonlinear robust principal component analysis. In *Advances in Neural Information Processing Systems*, pages 13390–13400, 2019.
- [95] MA Mendez, M Raiola, A Masullo, S Discetti, A Ianiro, R Theunissen, and J-M Buchlin. Pod-based background removal for particle image velocimetry. *Experimental Thermal and Fluid Science*, 80:181–192, 2017.
- [96] Igor Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1-3):309–325, 2005.
- [97] Igor Mezic. Analysis of fluid flows via spectral properties of the Koopman operator. *Annual Review of Fluid Mechanics*, 45:357–378, 2013.
- [98] Mark A Miller, Subrahmanyam Duvvuri, Ian Brownstein, Marcus Lee, John O Dabiri, and Marcus Hultmark. Vertical-axis wind turbine experiments at full dynamic similarity. *Journal of Fluid Mechanics*, 844:707–720, 2018.
- [99] Michael Miller, Jennifer Cardona, Leanne Block, Kenta Kondo, Michael Lee, Rebecca Lorick, Michael Manning, Isabel Scherl, Filip Simeski, Arriane Spaulding, et al. Results from the field test of two 1 kw oscillating hydrofoil generators in a tidal canal. In *APS Division of Fluid Dynamics Meeting Abstracts*, pages M2–009, 2016.

- [100] Vincent S Neary, Michael Lawson, Mirko Previsic, Andrea Copping, Kathleen C Hallett, Alison Labonte, Jeff Rieks, Dianne Murray, et al. Methodology for design and economic analysis of marine energy conversion (mec) technologies. 2014.
- [101] John A Nelder and Roger Mead. A simplex method for function minimization. *The computer journal*, 7(4):308–313, 1965.
- [102] Michele Nini, Valentina Motta, Giampiero Bindolino, and Alberto Guardone. Three-dimensional simulation of a complete vertical axis wind turbine using overlapping grids. *Journal of Computational and Applied Mathematics*, 270:78–87, 2014.
- [103] 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.
- [104] Bernd R Noack, Marek Morzynski, and Gilead Tadmor. *Reduced-order modelling for flow control*, volume 528. Springer Science & Business Media, 2011.
- [105] Bernd R Noack, Witold Stankiewicz, Marek Morzynski, and Peter J Schmid. Recursive dynamic mode decomposition of a transient cylinder wake. *Journal of Fluid Mechanics*, 809:843–872, 2016.
- [106] J Nogueira, A Lecuona, and PA Rodriguez. Data validation, false vectors correction and derived magnitudes calculation on PIV data. *Meas. Sci. and Tech.*, 8(12):1493, 1997.
- [107] Colin M Parker and Megan C Leftwich. The effect of tip speed ratio on a vertical axis wind turbine at high reynolds numbers. *Experiments in Fluids*, 57(5):74, 2016.
- [108] K. Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2(7–12):559–572, 1901.

- [109] Brian Polagye, Ben Strom, Hannah Ross, Dominic Forbush, and Robert J Cavignaro. Comparison of cross-flow turbine performance under torque-regulated and speed-regulated control. *Journal of Renewable and Sustainable Energy*, 11(4):044501, 2019.
- [110] Steven J Portugal, Tatjana Y Hubel, Johannes Fritz, Stefanie Heese, Daniela Trobe, Bernhard Voelkl, Stephen Hailes, Alan M Wilson, and James R Usherwood. Upwash exploitation and downwash avoidance by flap phasing in ibis formation flight. *Nature*, 505(7483):399–402, 2014.
- [111] Antonio Posa. Wake characterization of coupled configurations of vertical axis wind turbines using large eddy simulation. *International Journal of Heat and Fluid Flow*, 75:27–43, 2019.
- [112] Antonio Posa. Wake characterization of paired cross-flow turbines. *Renewable Energy*, 2022.
- [113] Antonio Posa, Colin M Parker, Megan C Leftwich, and Elias Balaras. Wake structure of a single vertical axis wind turbine. *International Journal of Heat and Fluid Flow*, 61:75–84, 2016.
- [114] Antonio Posa, Colin M Parker, Megan C Leftwich, and Elias Balaras. Wake structure of a single vertical axis wind turbine. *International Journal of Heat and Fluid Flow*, 61:75–84, 2016.
- [115] M. Raffel, C. E. Willert, F. Scarano, C. J. Kähler, S. T. Wereley, and J. Kompenhans. *Particle image velocimetry: a practical guide*. Springer, 2018.
- [116] Maziar Raissi and George Em Karniadakis. Hidden physics models: Machine learning of nonlinear partial differential equations. *Journal of Computational Physics*, 357:125–141, 2018.

- [117] Bernardo Luiz R Ribeiro, Yunxing Su, Quentin Guillaumin, Kenneth S Breuer, and Jennifer A Franck. Wake-foil interactions and energy harvesting efficiency in tandem oscillating foils. *Physical Review Fluids*, 6(7):074703, 2021.
- [118] Matthew P Rockwood, Kunihiko Taira, and Melissa A Green. Detecting vortex formation and shedding in cylinder wakes using lagrangian coherent structures. *AIAA Journal*, 55(1):15–23, 2017.
- [119] V Rolin and Fernando Porté-Agel. Wind-tunnel study of the wake behind a vertical axis wind turbine in a boundary layer flow using stereoscopic particle image velocimetry. In *Journal of physics: conference series*, volume 625, page 012012. IOP Publishing, 2015.
- [120] Hannah Ross and Brian Polagye. Blockage effects on current turbine performance and wake characteristics. *European Wave and Tidal Energy Conference*, 1(2 (Nov)), 2017.
- [121] Hannah Ross and Brian Polagye. An experimental assessment of analytical blockage corrections for turbines. *Renewable Energy*, 152:1328–1341, 2020.
- [122] C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D.S. Henningson. Spectral analysis of nonlinear flows. *J. Fluid Mech.*, 645:115–127, 2009.
- [123] Clarence W Rowley and Scott TM Dawson. Model reduction for flow analysis and control. *Annual Review of Fluid Mechanics*, 49:387–417, 2017.
- [124] Kevin J Ryan, Filippo Coletti, Christopher J Elkins, John O Dabiri, and John K Eaton. Three-dimensional flow field around and downstream of a subscale model rotating vertical axis wind turbine. *Experiments in Fluids*, 57(3):38, 2016.
- [125] A Sagharichi, M Zamani, and Am Ghasemi. Effect of solidity on the performance of variable-pitch vertical axis wind turbine. *Energy*, 161:753–775, 2018.
- [126] SH Salter. Are nearly all tidal stream turbine designs wrong? In *4th International Conference on Ocean Energy*, pages 1–7, 2012.

- [127] Fulvio Scarano. Iterative image deformation methods in PIV. *Measurement science and technology*, 13(1):R1, 2001.
- [128] Isabel Scherl, Benjamin Strom, Steven L Brunton, and Brian L Polagye. Geometric and control optimization of a two cross-flow turbine array. *Journal of Renewable and Sustainable Energy*, 12(6):064501, 2020.
- [129] Isabel Scherl, Benjamin Strom, Jessica K Shang, Owen Williams, Brian L Polagye, and Steven L Brunton. Robust principal component analysis for modal decomposition of corrupt fluid flows. *Physical Review Fluids*, 5(5):054401, 2020.
- [130] Frank Scheurich and Richard E Brown. Modelling the aerodynamics of vertical-axis wind turbines in unsteady wind conditions. *Wind Energy*, 16(1):91–107, 2013.
- [131] Frank Scheurich, Timothy M Fletcher, and Richard E Brown. Effect of blade geometry on the aerodynamic loads produced by vertical-axis wind turbines. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 225(3):327–341, 2011.
- [132] P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, August 2010.
- [133] FFJ Schrijer and Fulvio Scarano. Effect of predictor–corrector filtering on the stability and spatial resolution of iterative PIV interrogation. *Exp. Fluids*, 45(5):927–941, 2008.
- [134] Sina Shamsoddin and Fernando Porté-Agel. Large eddy simulation of vertical axis wind turbine wakes. *Energies*, 7(2):890–912, 2014.
- [135] Jessica K Shang. *Flexibility and curvature effects on vortex dynamics and fluid-structure interactions*. PhD thesis, Princeton University, 2015.
- [136] L. Sirovich. Turbulence and the dynamics of coherent structures, parts I-III. *Q. Appl. Math.*, XLV(3):561–590, 1987.

- [137] Abigale Snortland, Brian Polagye, and Owen Williams. Influence of near-blade hydrodynamics on cross-flow turbine performance. In *Proceedings of the 13th European Wave and Tidal Energy Conference*, pages 1–9, 2019.
- [138] Andrews Sobral, Thierry Bouwmans, and El-hadi Zahzah. Lrslibrary: Low-rank and sparse tools for background modeling and subtraction in videos. In *Robust Low-Rank and Sparse Matrix Decomposition: Applications in Image and Video Processing*. CRC Press, Taylor and Francis Group., 2015.
- [139] B. Strom, S. L. Brunton, and B. Polagye. Intracycle angular velocity control of cross-flow turbines. *Nature Energy*, 2(17103):1–9, 2017.
- [140] Benjamin Strom. *Cross-Flow Turbine Fluid Mechanics: Experimental Optimization and Analysis*. PhD thesis, University of Washington, 2019.
- [141] Benjamin Strom, Noah Johnson, and Brian Polagye. Impact of blade mounting structures on cross-flow turbine performance. *Journal of Renewable and Sustainable Energy*, 10(3):034504, 2018.
- [142] Benjamin Strom, Brian Polagye, and Steven L Brunton. Near-wake dynamics of a vertical-axis turbine. *Journal of Fluid Mechanics*, 935, 2022.
- [143] K. Taira and T. Colonius. The immersed boundary method: a projection approach. *Journal of Computational Physics*, 225(2):2118–2137, 2007.
- [144] Kunihiko Taira, Steven L Brunton, Scott TM Dawson, Clarence W Rowley, Tim Colonius, Beverley J McKeon, Oliver T Schmidt, Stanislav Gordeyev, Vassilios Theofilis, and Lawrence S Ukeiley. Modal analysis of fluid flows: An overview. *Aiaa Journal*, pages 4013–4041, 2017.
- [145] Kunihiko Taira, Maziar S Hemati, Steven L Brunton, Yiyang Sun, Karthik Duraisamy, Shervin Bagheri, Scott Dawson, and Chi-An Yeh. Modal analysis of fluid flows: Applications and outlook. *arXiv preprint arXiv:1903.05750*, 2019.

- [146] Aaron Towne, Oliver T Schmidt, and Tim Colonius. Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *Journal of Fluid Mechanics*, 847:821–867, 2018.
- [147] J. H. Tu, C. W. Rowley, J. N. Kutz, and J. K. Shang. Spectral analysis of fluid flows using sub-Nyquist-rate PIV data. *Exp. Fluids*, 55(9):1805, 2014.
- [148] 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.
- [149] A Vergaerde, T De Troyer, J Kluczevska-Bordier, N Parneix, F Silvert, and MC Runacres. Wind tunnel experiments of a pair of interacting vertical-axis wind turbines. In *Journal of Physics: Conference Series*, volume 1037, page 072049. IOP Publishing, 2018.
- [150] Antoine Vergaerde, Tim De Troyer, Lieven Standaert, Joanna Kluczevska-Bordier, Denis Pitance, Alexandre Immas, Frédéric Silvert, and Mark C Runacres. Experimental validation of the power enhancement of a pair of vertical-axis wind turbines. *Renewable Energy*, 146:181–187, 2020.
- [151] Siddhartha Verma, Guido Novati, and Petros Koumoutsakos. Efficient collective swimming by harnessing vortices through deep reinforcement learning. *Proceedings of the National Academy of Sciences*, 115(23):5849–5854, 2018.
- [152] J. Westerweel. Efficient detection of spurious vectors in particle image velocimetry data. *Exp. Fluids*, 16(3-4):236–247, 1994.
- [153] J. Westerweel, G. E. Elsinga, and R. J. Adrian. Particle image velocimetry for complex and turbulent flows. *Annual Review of Fluid Mechanics*, 45:409–436, 2013.
- [154] J. Westerweel and F. Scarano. Universal outlier detection for PIV data. *Experiments in fluids*, 39(6):1096–1100, 2005.

- [155] Robert W Whittlesey, Sebastian Liska, and John O Dabiri. Fish schooling as a basis for vertical axis wind turbine farm design. *Bioinspiration & biomimetics*, 5(3):035005, 2010.
- [156] C.E. Willert and M. Gharib. Digital particle image velocimetry. *Exp. Fluids*, 10(4):181–193, 1991.
- [157] Ryan H Wiser, Mark Bolinger, Ben Hoen, Dev Millstein, Joseph Rand, Galen L Barbose, Naim R Darghouth, Will Gorman, Seongeun Jeong, Andrew D Mills, et al. Land-based wind market report: 2021 edition. Technical report, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States), 2021.
- [158] Jeffrey Wong, Chris Colburn, Elijah Meeks, and Shankar Vedaraman. Rad–outlier detection on big data, February 2015.
- [159] Theodore Yaotsu Wu. Fish swimming and bird/insect flight. *Annual Review of Fluid Mechanics*, 43:25–58, 2011.
- [160] Y Yeh and HZ Cummins. Localized fluid flow measurements with an he–ne laser spectrometer. *Applied Physics Letters*, 4(10):176–178, 1964.
- [161] Xiaoming Yuan and Junfeng Yang. Sparse and low-rank matrix decomposition via alternating direction methods. *preprint*, 12, 2009.
- [162] Stefania Zanforlin and Takafumi Nishino. Fluid dynamic mechanisms of enhanced power generation by closely spaced vertical axis wind turbines. *Renewable Energy*, 99:1213–1226, 2016.

Appendix A

APPENDIX

A.1 Appendix: Optimized parameters

The optimized parameters for the counter-rotating array operating under coordinated control and tip-speed ratio control are shown in Figure 3.6 and Figure A.1, respectively. These results are similar to the co-rotating coordinated control and tip-speed ratio control cases explored in Section 3.3.1.

A.2 Appendix: Uncertainty analysis

To evaluate uncertainty in the interaction factor (κ), we apply propagation of uncertainty. Interaction factor is defined as

$$\kappa = \frac{C_{P,1} + C_{P,2}}{C_{P,1}^* + C_{P,2}^*} - 1. \quad (\text{A.1})$$

For mathematical convenience, we define

$$\kappa' = \kappa + 1 = \frac{C_{P,1} + C_{P,2}}{C_{P,1}^* + C_{P,2}^*} = \frac{A}{B} \quad (\text{A.2})$$

The propagation of uncertainty in the numerator and denominator are, respectively, given as

$$\delta A = \left((\delta C_{P,1})^2 + (\delta C_{P,2})^2 \right)^{\frac{1}{2}} \quad (\text{A.3})$$

and

$$\delta B = \left((\delta C_{P,1}^*)^2 + (\delta C_{P,2}^*)^2 \right)^{\frac{1}{2}}. \quad (\text{A.4})$$

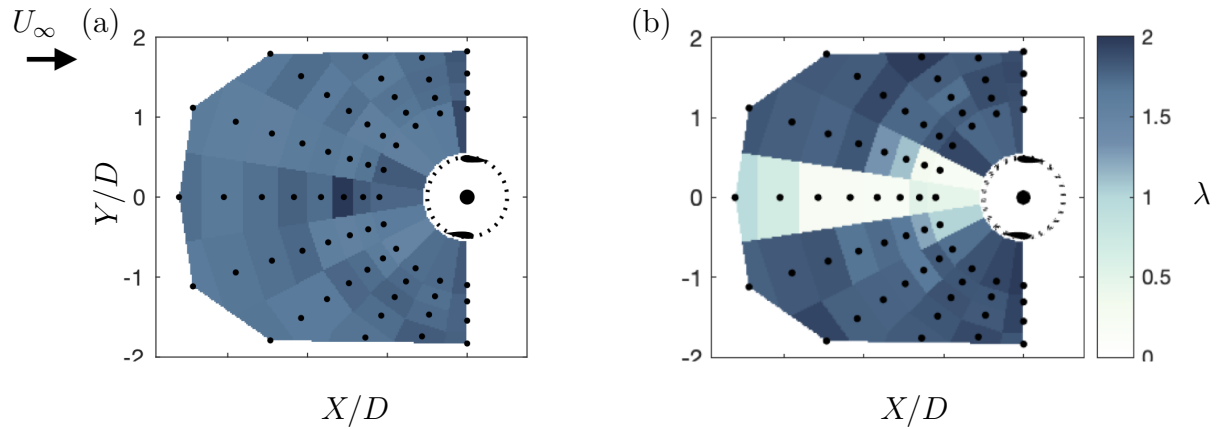


Figure A.1: Optimized parameters for counter-rotating turbines under tip-speed ratio control. (a) Upstream turbine tip-speed ratio for a given array layout and (b) downstream turbine tip-speed ratio for the same.

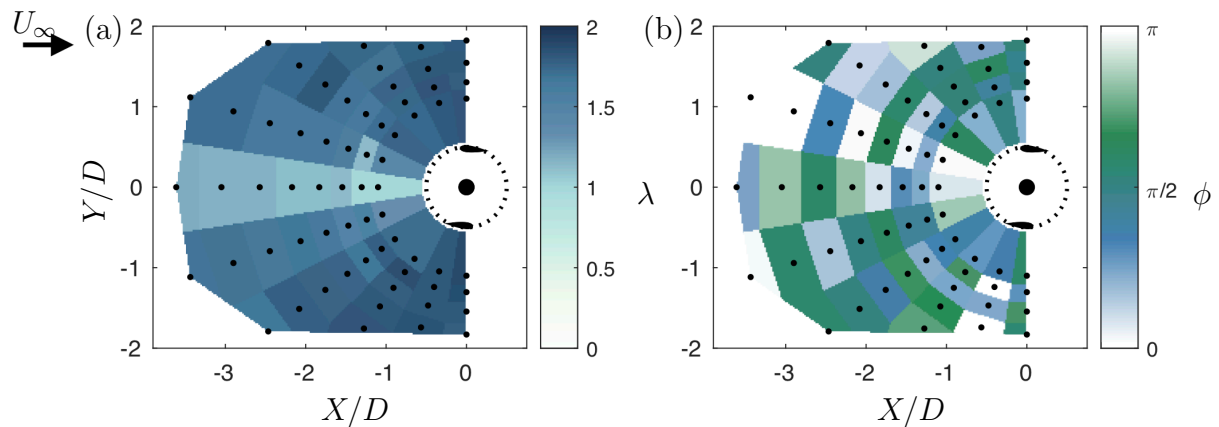


Figure A.2: Optimized parameters for counter-rotating turbines under coordinated control: (a) optimized tip-speed ratio for the turbine pair and (b) optimized phase difference.

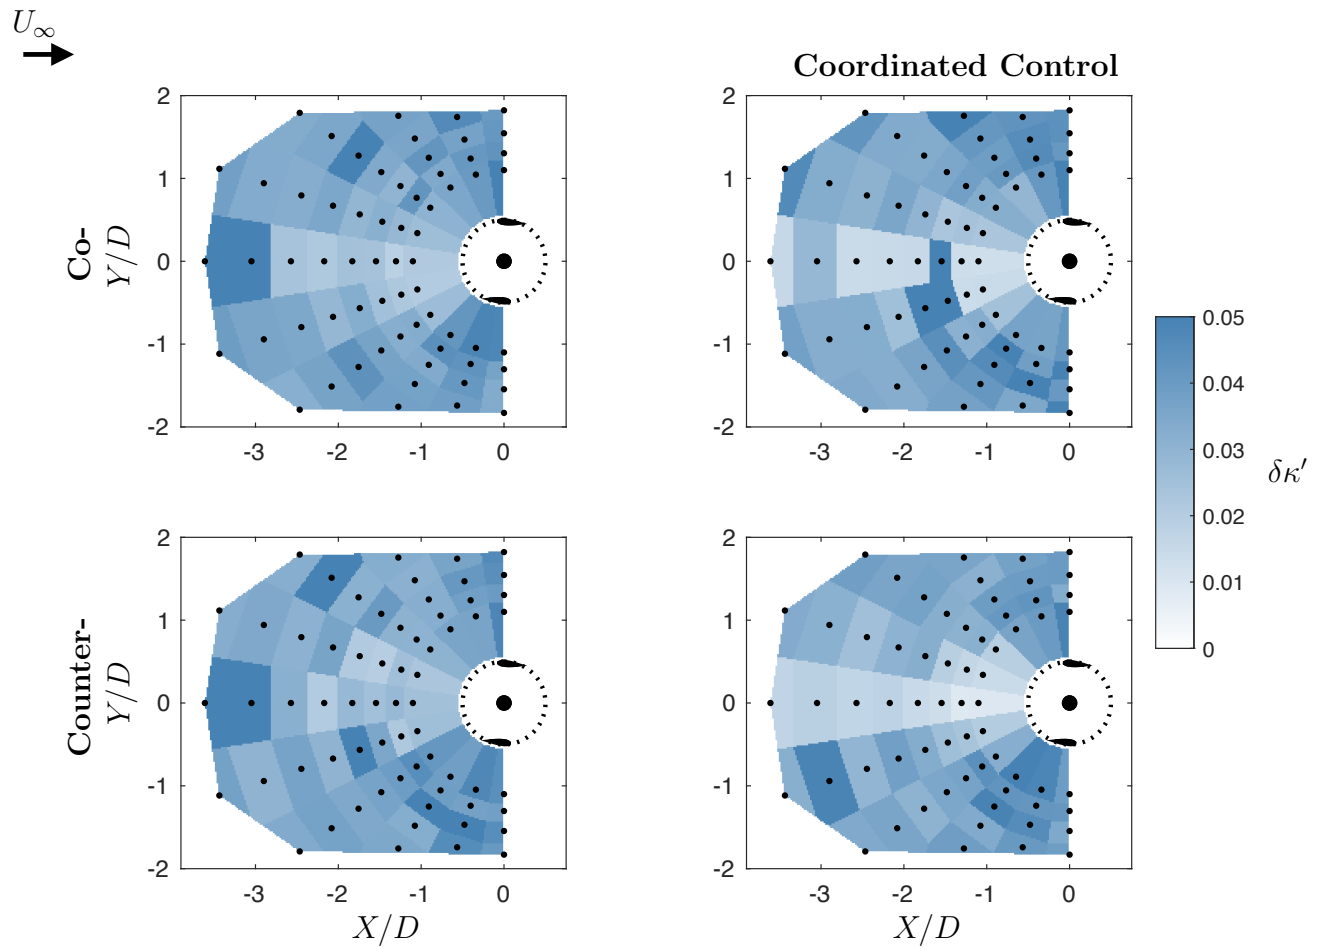


Figure A.3: Uncertainty values ($\delta\kappa'$) calculated from Eq. 11 for both rotation directions and control strategies.

Consequently, the uncertainty in κ' is given by fractional quadrature as

$$\frac{\delta\kappa'}{|\kappa'|} = \left(\left(\frac{\delta A}{A} \right)^2 + \left(\frac{\delta B}{B} \right)^2 \right)^{\frac{1}{2}}. \quad (\text{A.5})$$

Finally, we note that

$$\delta\kappa = \delta\kappa', \quad (\text{A.6})$$

such that $\delta\kappa$ can be calculated for each turbine position and control strategy.

This does, however, require an estimate for C_P . Depending on ω , each test consists of $\approx 12 - 24$ rotations and the values of cycle-average C_P . We use the inter-quartile range of cycle-average C_P within each test as a proxy for δC_P for each test. Uncertainty for each of the four cases are presented in Figure A.3. The “dead band” on the performance figures corresponds to the approximate average uncertainty of all measurement points.