

Reconstruction and Estimation of Flows Using Resolvent Analysis and Data-Assimilation

Thesis by
Sean Pearson Symon

In Partial Fulfillment of the Requirements for the
degree of
Ph.D. Aeronautics

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2018
Defended May 11, 2018

© 2018

Sean Pearson Symon
ORCID: 0000-0001-9085-0778

All rights reserved

To my parents and sister

ACKNOWLEDGEMENTS

First I would like to thank my family, particularly my parents who decided to leave their friends and family behind in Scotland for a new life in the U.S. This was such a brave decision on their part, particularly since my sister and I were only two and four years old, respectively, at the time. I cannot overstate the number of doors that have been opened for me as a result of their decision. My family in Scotland has also been incredibly supportive and it always feels like paradise when I can leave behind my world here and fly across the ocean to see them.

Next, I would like to thank my thesis committee: Professors Tim Colonius, Anthony Leonard, and Dan Meiron. Each member of my committee has helped me much more than they probably think they have by asking insightful questions during conferences/presentations or offering seemingly simple advice which turned out to be incredibly useful. I next would like to thank my fellow G1's especially Mélanie Delapierre, Simon Lapointe, Neel Nadkarni, Chris Roh, and Kevin Rosenberg. I think every GALCIT student appreciates that the friendship you develop with your fellow first years is quite special as most grad students don't have to take on five classes as soon as they arrive at Caltech. My choice to study abroad at Ecole Polytechnique meant that I took quals a year later than everyone else. My fellow G2's, as I like to call them, were extremely welcoming when I returned from France, particularly Noah Braun, Ian Brownstein, Ilana Gat, David & Nikki Huynh, and Jason & Susie Schlup. Qualls were not the most pleasant experience for me but I was able to get through them largely in part due to their help.

There are several more GALCIT individuals who I would like to mention. First and foremost is Christine Ramirez, who was my first point of contact with Caltech. She has always been incredibly helpful and welcoming to not only me but also the rest of the incoming class. I would like to thank both Dimity Nelson and Peggy Blue, the department administrators, for their incredibly efficiency with running GALCIT. With respect to the faculty, I would like to acknowledge Ravi, who I met during visit day. His personality impressed me so much that it ended up being one of the primary reasons I chose to attend Caltech. I would also like to thank Professors Pullin and Dimotakis, who taught Ae101. Despite having very different teaching styles, I learned a significant amount in this course and it fundamentally changed my perspectives and approach to science. It was also nice to go full circle and TA for Professor Pullin's Ae 101 course on two separate occasions.

I also thank GALCIT for allowing me to participate in the Ecole Polytechnique exchange program, where I met my friends Charu Datt and Keshav Raja. I still enjoy talking about politics and science on a regular basis with Keshav. During my research internship at ESPCI, I had the very good fortune of being paired up with my mentor, Anaïs Gauthier. As an inexperienced scientist, I learned a great deal from her and other members of the lab directed by David Quéré and Christophe Clanet. This brings me finally to my advisor in France, David Quéré, who I would like to thank for allowing me to join the group based exclusively on my enthusiasm for the projects in his lab.

I would like to thank my collaborators Rémi Carmigniani, Peter Schmid, and Denis Sipp. Rémi was one of my fellow first years and teammate for Ae104c along with Simon Lapointe for the bouncing oil drops experiment. It was a great pleasure to work with both of them on this project although it took months for me to finally convince myself that my hands were rid of all traces of the silicone oil we used! I met Peter briefly before I headed over to France as he was on sabbatical and spent some of his time at Caltech. Peter has been an incredibly positive influence on my Ph.D. in a number of ways including the control of shear flows course he taught at Polytechnique and our collaboration on the experimental data-assimilation project. He also co-organized a workshop in the heart of Tuscany where I met Denis, who also played a very positive role on my Ph.D. research both in terms of the data-assimilation project as well as resolvent analysis. I wish to thank him in particular for making his codes available at the workshop and assisting me with setting up the resolvent code on a Caltech cluster.

There are some friends outside of research who I would like to thank. First is Harrison Chau with whom I have traveled on many occasions around the world. He also graciously hosted me in both San Francisco and Seattle on many occasions in addition to visiting me in Los Angeles. Next is my friend Hannah Wang who I met in undergrad. She has always taken a keen interest in my academic progress and always makes an effort to visit me whether I'm in Seattle, D.C., Paris, or Vienna. Finally, I would like to give a shout out to my Hebron crew and my housemates, Paul Magyar and Roberta Poceviciute, with whom I enjoyed sharing Caltech gossip and attending concerts at the Hollywood Bowl & Walt Disney.

Speaking of music, there are many excellent people at the performing and visual arts department who more than deserve to be acknowledged. First and foremost are Bill and Delores Bing, who essentially founded the music department and transformed it

into the enormous success it is today. I thank maestro Allen Gross, chamber music director Maia Jasper White, and PVA manager Cindy De Mesa for their diligence and dedication to Caltech's music program. I also want to thank Dorothy Pan and Zach Erickson with whom I had the pleasure of playing Beethoven's Spring Sonata.

After returning from France, I had the extremely good fortune of being a member of a piano quartet alongside Joe Iverson, Kelly Kim, and Ian Wong. Sarah Jeoung later joined us when we played Dvorak's piano quintet. We developed a high level of trust in each other allowing us to take on Beethoven Opus 132, one of my all time favorites. I will especially miss laughing hysterically during our sightreading trainwrecks and being coached by Martin Chalifour, the principal concertmaster of the LA Philharmonic. He is undoubtedly one of the kindest individuals I have ever met and I had the opportunity to meet other members of the LA Phil as well as one of my favorite pianists, Mitsuko Uchida.

I need to thank my 4-1 crew: Nick Dou, Reid Kawamoto, Amber Lin, Andrew Robbins, Kevin Rosenberg, Dingyi Sun, and Ishan Tembhekar. Despite my monthly bills skyrocketing every now and then, I thoroughly enjoyed our time together whether it was in Los Angeles, San Francisco, or Tokyo (to be clear this list is not complete). I will also miss our weekly tennis nights, which varied substantially in quality week-to-week, and the 1055 parties.

I want to thank every member of the McKeon group including Ben Barthel, Scott and Lucy Nicholls, Reeve Dunne, Subrahmanyam Duvvuri, Morgan Hooper, Esteban Hufstedler, David Huynh, Angeliki Laskari, Mitul Luhar, Ryan McMullen, Rashad Moarref, Jonathan Morgan, Kevin Rosenberg, Tess Saxton-Fox, Maysam Shamai, Professor Ati Sharma, and Simon Tödtli. Morgan, Esteban, and Maysam helped me setup experiments in NOAH, Kevin and Simon assisted in setting up various pieces of code on the clusters, Ati provided valuable feedback with respect to the cylinder resolvent work, and Jonathan organized many good bad movie nights. A special thanks to Angeliki, Scott, and Lucy, who organized the Happy Hour Club. Lucy had a knack for choosing the right places since we ended up running into the former prime minister of Australia in one of the fancier establishments!

Finally (I always save the best for last), I would like to thank Jamie and Beverley. Jamie was the administrative assistant for our group and, in addition to being incredibly effective at her job, was someone I just enjoyed talking to. I invented excuses to go to her office so that we could catch up. I could easily write a separate thesis on what makes Beverley a good advisor, but I will sum it up in two words:

leadership and trust. Leadership is a unique skill which has little to do with exercising authority as most people believe. Instead, it is the ability to inspire other to realize their full potential. I also believe the extremely friendly and supportive group environment is testament to her exemplary leadership. Finally, I trust Beverley because I know she gives advice and makes decisions based on what is best for me and the group.

Being a member of the McKeon group has, without a doubt, been the highlight of my graduate career. I wish Beverley and all of my fellow group members the best of luck at Caltech and beyond.

Financial support for this thesis comes from a National Science Foundation (NSF) Graduate Fellowship, the Air Force Office of Scientific Research (AFOSR), under Grant No. FA 9550-16-1-0361, and the Army Research Office (ARO), under Grant No. W911NF-17-1-0306.

ABSTRACT

A flow reconstruction methodology is presented for incompressible, statistically stationary flows using resolvent analysis and data-assimilation. The only inputs necessary for the procedure are a rough approximation of the mean profile and a single time-resolved measurement. The objective is to estimate both the mean and fluctuating states of experimental flows with limited measurements which do not include pressure. The input data may be incomplete, in the sense that measurements near a body are difficult to obtain with techniques such as particle image velocimetry (PIV), or contaminated by noise. The tools developed in this thesis are capable of filling in missing data and reducing the amount of measurement noise by leveraging the governing equations. The reconstructed flow is capable of estimating fluctuations where time-resolved data are not available and solving the flow on larger domains where the mean profile is not known.

The first part of the thesis focuses on how resolvent analysis of the mean flow selects amplification mechanisms. Eigenspectra and pseudospectra of the mean linear Navier-Stokes (LNS) operator are used to characterize amplification mechanisms in flows where linear mechanisms are important. The real parts of the eigenvalues are responsible for resonant amplification and the resolvent operator is low-rank when the eigenvalues are sufficiently separated in the spectrum. Two test cases are studied: low Reynolds number cylinder flow and turbulent channel flow. The latter is studied by considering well-known turbulent structures while the former contains a marginally stable eigenvalue which drowns out the effect of other eigenvalues over a large range of temporal frequencies. There is a geometric manifestation of this dominant mode in the mean profile, suggesting that it leaves a significant footprint on the time-averaged flow that the resolvent can identify. The resolvent does not provide an efficient basis at temporal frequencies where there is no separation of singular values. It can still be leveraged, nevertheless, to identify coherent structures in the flow by approximating the nonlinear forcing from the interaction of highly amplified coherent structures.

The second part of the thesis extends the framework of Foures et al. (2014), who data-assimilated the mean cylinder wake at very low Reynolds numbers. The contributions presented here are to assess the minimum domain for successfully reconstructing Reynolds stress gradients, modifying the algorithm to assimilate mean pressure, determining whether weighting input measurements contributes to im-

proved performance, and adapting the method to experimental data at higher Reynolds numbers. The results from data-assimilating the mean cylinder wake at low Reynolds numbers suggest that the measurement domain needs to coincide with the spatial support of the Reynolds stress gradients while point weighting has a minimal impact on the performance. Finally, a smoothing procedure adapted from Foures et al. (2014) is proposed to cope with data-assimilating an experimental mean profile obtained from PIV data. The data-assimilated mean profiles for an idealized airfoil and NACA 0018 airfoil are solved on a large domain making the mean profile suitable for global resolvent analysis. Data-assimilation is also able to fill in missing or unreliable vectors near the airfoil surface.

The final piece of the thesis is to synthesize the knowledge and techniques developed in the first two parts to reconstruct the experimental flow around a NACA 0018 airfoil. Preliminary results are presented for the case where $\alpha = 0^\circ$ and $Re = 10250$. The mean profile is data-assimilated and used as an input to resolvent analysis to educe coherent structures in the flow. The resolvent operator for non-amplified temporal frequencies is forced by an approximated nonlinear forcing. The amplitude and phase of the modes are obtained from the discrete Fourier-transform of a time-resolved probe point measurement. The final reconstruction contains less measurement noise compared to the PIV snapshots and obeys the incompressible Navier-Stokes equations (NSE). The thesis concludes with a discussion of how elements of this methodology can be incorporated into the development of estimators for turbulent flows at high Reynolds numbers.

PUBLISHED CONTENT AND CONTRIBUTIONS

Symon, S., Dovetta, N., McKeon, B. J., Sipp, D., and Schmid, P. J. (2017). "Data assimilation of mean velocity from 2D PIV measurements of flow over an idealized airfoil". In: *Experiments in Fluids* 58.

<http://resolver.caltech.edu/CaltechAUTHORS:20170426-063702442>

Performed the experiments, solved the forward and adjoint equations, implemented adjoint looping procedure on cluster, analyzed the results, created the figures, and was the primary author of the paper.

Symon, S., Rosenberg, K., Dawson, S. T. M., and McKeon, B. J. (2018). "Non-normality and classification of amplification mechanisms in stability and resolvent analysis". In: *Physical Review Fluids* 3.

<http://resolver.caltech.edu/CaltechAUTHORS:20180517-082629521>

Ran the simulations, performed the global resolvent analysis, devised the model operators, created most of the figures, and was the primary author of the paper.

TABLE OF CONTENTS

Acknowledgements	iv
Abstract	viii
Published Content and Contributions	x
Table of Contents	xi
List of Illustrations	xv
List of Tables	xix
Nomenclature	xx
Chapter I: Introduction	1
1.1 Motivation	1
1.2 A Musical Analogy	2
1.2.1 Concertos and Arranging	2
1.2.2 Structure and Statistics	3
1.2.3 Beethoven Violin Concerto	5
1.3 Reduced-Order Modeling	7
1.3.1 Stability of the Mean Flow	8
1.3.2 Resolvent Analysis	9
1.3.3 Non-normality	10
1.3.4 Instabilities and Pseudoresonance	10
1.3.5 Data-Driven Methods	10
1.4 Data-Assimilation	11
1.4.1 Variational Methods	12
1.4.2 Incorporating Pressure and Extension to Experiments	14
1.5 Flow Reconstruction	14
1.6 Approach and Outline of Thesis	16
Chapter II: Methods	19
2.1 Governing Equations	19
2.2 Numerical Methods for Cylinder Flow	19
2.2.1 2D Cylinder	19
2.2.2 3D Cylinder	20
2.2.3 Results and Validation	20
2.2.4 Mesh Size and Resolution	21
2.3 Experimental Methods	22
2.3.1 Idealized Airfoil	22
2.3.2 NACA 0018 Airfoil	24
2.3.3 Vector Post-Processing	25
2.4 Modal Analysis Methods	25
2.4.1 Stability Analysis	26
2.4.2 Resolvent Analysis	29
2.4.3 Link to Pseudo-Spectrum	31

2.4.4	Flows with Homogeneous Directions	32
2.4.5	Dynamic Mode Decomposition	34
2.5	Data-Assimilation of the Mean Flow	34
2.5.1	Computational Model	35
2.5.2	Fitting Criterion	36
2.5.3	Cost Functional and Adjoint Equations	36
2.5.4	Mean Pressure Reconstruction	38
2.5.5	Amplitude Calibration	39
2.6	Smoothing Procedure	40
2.7	Adjoint Looping: Implementation of Data-Assimilation Algorithm	42
2.8	Numerical Details of Experimental Data-Assimilation	44
2.9	Evaluation of Agreement Between Experiment and Data-Assimilation	45
2.10	Summary	45
Chapter III: Amplification Mechanisms in Stability and Resolvent Analysis		47
3.1	Resonance and Pseudoresonance	47
3.2	Amplification Mechanisms Highlighted by a Model Operator	48
3.2.1	Eigenvectors, singular functions, and non-normality of the operator	48
3.2.2	Eigenvectors, singular functions, and self-adjointness of the operator	50
3.2.3	Mean flow advection and the Orr mechanism	51
3.2.4	Resolvent (approximate) wavemaker	51
3.3	Lower Rank Approximations of the Resolvent Operator	54
3.3.1	Dyad expansion of the Resolvent Operator	54
3.3.2	The Relationship Between Spectral Radius and Spectral Norm for Approximately Low-Rank Operators	55
3.4	Resonance curve and resolvent norm	57
3.5	Application to Cylinder Flow	60
3.5.1	Base flow velocity profile	60
3.5.2	Mean velocity profile	65
3.5.3	Resolvent modeling of (low Reynolds number) cylinder flow	67
3.6	Application to Wall-Bounded Turbulence	68
3.6.1	Near-wall cycle	69
3.6.2	VLSM	70
3.6.3	Stationary disturbances	73
3.6.4	Influence of spatial wavenumber and wave speed	75
3.7	Comparison of Flows	76
3.8	Summary and Contribution	78
Chapter IV: Scaling Characteristics and Nonlinear Forcing of Cylinder Flow		79
4.1	Scaling of the Mean Profile and Resolvent Norm	79
4.2	Scaling and Convection Velocity of Resolvent Modes	82
4.3	The Wavemaker and Recirculation Length	85
4.4	Singular Values and Mode Branches	87
4.5	Comparing Resolvent and DMD Modes	90
4.6	Parasitic Modes and Triadic Interactions	92

4.7	Nonlinear Forcing	96
4.8	Summary and a Closed Loop	99
Chapter V: Data-Assimilation of Numerical Data		103
5.1	Full Domain	103
5.1.1	Implications for experimental data	107
5.1.2	Stabilizing feedback	108
5.2	Mean Pressure Correction	109
5.3	Domain Truncation and Impact on Reconstructed Forcing	110
5.3.1	Capturing the Wavemaker	111
5.3.2	Nonlinear Forcing and Maximum Pressure Discrepancy	111
5.4	Weighting Measurement Points	113
5.5	Extension to 3D Wake	114
5.6	Summary	115
Chapter VI: Data-Assimilation of Experimental Data		118
6.1	Base Flow and PIV Results	118
6.2	Full-Field Information	119
6.3	Minimum Resolution and Domain	122
6.4	Residual Discrepancy	126
6.4.1	Three-dimensionality	126
6.4.2	Model simplifications	127
6.4.3	Choice of cost function	128
6.5	Zero Angle of Attack Case	129
6.5.1	$Re = 10250$	129
6.5.2	$Re = 20700$	131
6.6	$\alpha = 10^\circ$ Case	133
6.7	Summary	135
Chapter VII: Flow Reconstruction		137
7.1	Reconstruction Procedure	137
7.2	Resolvent Analysis Frequency Sweep	139
7.3	Comparison Between Resolvent and DMD Modes	142
7.4	Nonlinear Interactions for Low and High Frequency Modes	144
7.5	Probe Measurements for Velocity Fluctuations	147
7.6	Unsteady Flow Reconstruction	148
7.6.1	Selection of Probe Point and Point Reconstructions	149
7.6.2	Full Field Approximation Using Rank-1 Frequencies Only	151
7.6.3	Adding ‘Forced’ Resolvent Modes	152
7.7	Discussion	154
7.7.1	Error Metric	154
7.7.2	Influence of Mode Position and Amplitude	156
7.7.3	Influence of Velocity Component	156
7.8	Mean Pressure Reconstruction	157
7.9	Convergence of the Method	158
7.10	Summary	159
Chapter VIII: Conclusions and Future Work		161
8.1	Summary and Conclusions	161

8.1.1	Fundamentals of Resolvent Analysis and Application to Flows Around Bodies	161
8.1.2	Data-Assimilation of Time-Averaged Simulation and Ex- perimental Data	163
8.1.3	Flow Reconstruction from Very Limited Measurements . . .	164
8.2	Future Work	164
8.3	Implications for Estimating Turbulent Flows	166
	Bibliography	167
	Appendix A: Derivation of Adjoint RANS Equations	178
	Appendix B: NACA 0018 Airfoil With $\alpha = 10^\circ$	182
	B.1 Resolvent Frequency Sweep	182
	B.2 Power Spectrum of Shear Layer and Wake	183
	B.3 Resolvent Response Modes	183
	B.4 Resolving the Forcing Modes	185

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
1.1 Closed loop control and data-assimilation	2
1.2 First two lines of Schubert's Impromptu No. 3 in G-flat Major Opus 90	4
1.3 Opening measures of Beethoven's violin concerto	6
1.4 Another excerpt of Beethoven's violin concerto	7
1.5 Schematic of flow reconstruction approach	18
2.1 Shedding frequency ω_s as a function of Reynolds number Re	21
2.2 Time-averaged 2D DNS results for $Re = 150$	22
2.3 Schematic of the idealized airfoil	23
2.4 Experimental setup of the flow around an idealized airfoil	24
2.5 Subdivision of the domain into cells	40
2.6 Block diagram representation of the adjoint looping procedure	43
3.1 Cartoon of resolvent modes for various amplification mechanisms . .	53
3.2 Pseudospectra and resolvent norm for example operators	58
3.3 Projection of first resolvent mode onto eigenvectors	60
3.4 Spectrum and pseudospectrum of the cylinder base flow at $Re = 47$.	61
3.5 Stability and resolvent modes for the cylinder base flow at $Re = 47$.	63
3.6 The adjoint eigenmode, forward eigenmode, and wavemaker for cylinder flow at subcritical Reynolds numbers	64
3.7 Spectrum and pseudospectrum of the cylinder mean flow at $Re = 100$	65
3.8 Resolvent norm for subcritical base flows and supercritical mean flows	66
3.9 Wavemakers computed from stability and resolvent modes	68
3.10 Turbulent mean velocity and mean shear profiles for channel flow at $Re_\tau = 2000$	69
3.11 $ \hat{\phi}_1(y) $ and $ \hat{\psi}_1(y) $ for the near-wall cycle mode where $(k_x, k_z, c^+) = (4\pi, 40\pi, 14)$	70
3.12 Spectrum and pseudospectrum of the near-wall cycle mode	71
3.13 First 20 singular values of the near-wall cycle mode	71
3.14 $ \hat{\phi}_1(y) $ and $ \hat{\psi}_1(y) $ for the near-wall cycle mode where $(k_x, k_z, c^+) = (\pi/9, 2\pi/3, 22)$	72
3.15 Spectrum and pseudospectrum of the VLSM mode	72
3.16 First 20 singular values of the VLSM mode	73

3.17	$ \hat{\phi}_1(y) $ and $ \hat{\psi}_1(y) $ for the stationary disturbance mode where $(k_x, k_z, c^+) = (0, 2\pi/3, 0)$	74
3.18	Spectrum and pseudospectrum of the stationary disturbance mode	75
3.19	First 20 singular values of the stationary disturbance mode	75
4.1	First two singular values for mean cylinder wake at Reynolds numbers where flow is 2D	80
4.2	Schematic defining recirculation length l_m	80
4.3	Unscaled and scaled mean profiles at recirculation point for 2D cylinder flow	81
4.4	Unscaled and scaled mean profiles at recirculation point for square cylinder and 3D cylinder flows	82
4.5	Unscaled and scaled base profiles at supercritical Reynolds numbers	82
4.6	Unscaled and scaled profiles of the first resolvent modes along the centerline	83
4.7	Mean profile on centerline and vortex convection velocity	85
4.8	Recirculation bubble length, shedding frequency, and position of vortex cores	86
4.9	First two singular values for $Re = 100$ mean cylinder wake	87
4.10	$\hat{\psi}_1$ and $\hat{\psi}_2$ at various ω for $Re = 100$ mean cylinder wake	89
4.11	Unscaled and scaled profiles on centerline for resolvent modes at various ω for $Re = 100$ mean cylinder wake	90
4.12	Comparison of resolvent and DMD modes for $Re = 100$ cylinder flow	91
4.13	Discrepancy between resolvent and DMD modes for $Re = 100$ cylinder flow	92
4.14	Nonlinear forcing at $2\omega_s$ vs. self-interaction of $\hat{\psi}_1$	93
4.15	Velocity response modes generated by forcing the $\mathcal{H}(2\omega_s)$ resolvent operator with $\hat{\psi}_1(\omega_s) \cdot \nabla \hat{\psi}_1(\omega_s)$	94
4.16	Discrepancy between the forced resolvent and DMD modes for $Re = 100$ cylinder flow	94
4.17	Velocity response modes generated by forcing the $\mathcal{H}(3\omega_s)$ resolvent operator with an approximated $\hat{f}(3\omega_s)$	97
4.18	Nonlinear forcing at the shedding frequency	98
4.19	Projection of nonlinear forcing onto resolvent forcing mode	99
4.20	Instability regimes in cylinder flow	100
4.21	Triadic interactions in $Re = 100$ cylinder flow	102
5.1	$\mathcal{E}(\bar{\mathbf{u}})$ for various cases of cylinder data-assimilation	105

5.2	Reconstructed pressure and components of mean forcing for measurements on the whole domain	106
5.3	Absolute pressure discrepancy for initial guess, assimilation without pressure measurements, and assimilation with surface pressure measurements	107
5.4	Least stable eigenvalue at various stages of the data-assimilation . . .	109
5.5	Pressure correction factor predicted by resolvent modes compared to DNS	110
5.6	Reconstructed mean forcing for various truncated domains	112
5.7	Data-assimilation results for 3D cylinder flow	115
5.8	f_x comparison for 3D cylinder flow	115
5.9	$\nabla \times \mathbf{f}$ comparison for 3D cylinder flow	116
6.1	Experimental and base flow comparison for idealized airfoil	119
6.2	Full-field mean velocity reconstruction for idealized airfoil	120
6.3	Reconstructed \mathbf{f} for idealized airfoil	121
6.4	Reconstructed $\nabla \times \mathbf{f}$ for idealized airfoil	122
6.5	Effect of resolution on mean velocity reconstruction for idealized airfoil	123
6.6	Effect of resolution on forcing reconstruction for idealized airfoil . .	124
6.7	Effect of domain on mean velocity reconstruction for idealized airfoil	125
6.8	Effect of domain on forcing reconstruction for idealized airfoil	126
6.9	Divergence of idealized airfoil velocity fields	127
6.10	Residual discrepancy of idealized airfoil	128
6.11	Mean velocity reconstruction for NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$	129
6.12	Reconstructed \mathbf{f} for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$	130
6.13	Reconstructed $\nabla \times \mathbf{f}$ for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$	131
6.14	Mean velocity reconstruction for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 20700$	132
6.15	Reconstructed $\nabla \times \mathbf{f}$ for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 20700$	132
6.16	Mean velocity reconstruction for the NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$	133
6.17	Reconstructed $\nabla \times \mathbf{f}$ for the NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$	134
7.1	Schematic of flow reconstruction procedure	138

7.2	First three singular values of the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$	140
7.3	Comparison of $\hat{\phi}_1$ for $\omega_g = 12.24$ of the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$	141
7.4	DMD eigenvalues of the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$ 142	
7.5	Comparison of $\hat{\psi}_1$ and its DMD counterpart for the globally most amplified frequency of the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$	143
7.6	Resolvent and DMD mode comparison at various frequencies for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$	144
7.7	Resolvent modes using an approximated nonlinear forcing at $\omega = 24.49$	145
7.8	Resolvent modes using an approximated nonlinear forcing at $\omega = 2.90$ 146	
7.9	Probe points for NACA 0018 at $\alpha = 0^\circ$ and $Re = 10250$ flow reconstruction	147
7.10	Power spectrum at points P4 and P7 in the airfoil wake	148
7.11	Time traces of the reconstructed v -component when the probe point is located at P4	149
7.12	Time traces of the reconstructed v -component when the probe point is located at P7	150
7.13	Reconstructed flow at $t = 0$ using rank-1 modes for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$	151
7.14	Fluctuating v -component at Point P7	153
7.15	Reconstructed flow for various times using all modes for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$	153
7.16	Reconstructed u -velocity using the flow reconstruction procedure and various numbers of frequencies	157
7.17	\bar{p} for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$	158
B.1	First two singular values of the NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$	182
B.2	Points where the power spectrum is computed for the $\alpha = 10^\circ$ flow	183
B.3	PSD of two points in the $\alpha = 10^\circ$ flow	184
B.4	Resolvent response modes computed for the flow around a NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$	184
B.5	Resolvent forcing modes computed for the flow around a NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$	185

LIST OF TABLES

<i>Number</i>	<i>Page</i>
2.1 Operator form of the equations for stability and resolvent analyses . . .	29
2.2 Summary of flows considered	46
3.1 Quantification of non-normality for cylinder and channel flow	76
5.1 Cylinder data-assimilation results	104
6.1 Idealized airfoil data-assimilation results.	119
7.1 Points used in the NACA0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$ flow reconstruction	147
7.2 Comparison of the global error for different input points, velocity components, and range of frequencies	155

NOMENCLATURE

- $(\cdot)^+$. pseudo-inverse.
- α . angle of attack.
- α_s . stall angle.
- β . step size in new descent direction.
- $\Delta\bar{u}$. discrepancy velocity between experiment and simulation.
- $\Delta\bar{p}$. discrepancy pressure field between experiment and simulation.
- γ . eigenvector of companion matrix.
- Λ . diagonal matrix of eigenvalues.
- Φ . resolvent forcing modes or right singular vectors of the resolvent operator.
- Ψ . resolvent response modes or left singular vectors of the resolvent operator.
- Σ . diagonal matrix containing the singular values of the resolvent operator.
- v . DMD eigenvector.
- A . LNS operator with respect to the base flow.
- A^* . adjoint LNS operator with respect to the base flow.
- B . projector to velocity space.
- C . restriction operator.
- E . a perturbation matrix.
- f . mean forcing or the divergence of the Reynolds stress tensor.
- f' . nonlinear forcing.
- f_s . solenoidal component of mean forcing field.
- G . stiffness matrix for 1D case.
- I . identity matrix.
- J . forcing operator.
- K . response operator.
- k . wavenumber vector.

- L . LNS operator with respect to the mean flow.
- L^* . adjoint LNS operator with respect to the mean flow.
- M . 2-by-2 model operator.
- Q . generic, nonsingular operator.
- R . Reynolds stress tensor.
- r . vector containing the ratio of singular vector inner products.
- S . companion matrix in the DMD algorithm.
- T . linear mapping between snapshot matrices.
- U . left singular vectors of an arbitrary matrix.
- u . velocity field.
- u' . fluctuating velocity field.
- u^\dagger . adjoint velocity field.
- U_0 . base flow velocity field.
- V . matrix of eigenvectors.
- W . right singular vectors of an arbitrary matrix.
- X . snapshot matrix.
- $X^\#$. displaced snapshot matrix.
- $\check{(\cdot)}$. time-resolved measurement.
- χ . complex amplitude of a resolvent response mode.
- Δt . non-dimensional time step.
- η . wall-normal vorticity.
- $\hat{(\cdot)}$. resolvent analysis.
- $\hat{\phi}$. a resolvent forcing mode.
- $\hat{\psi}$. a resolvent response mode.
- κ . condition number.
- Λ . set of eigenvalues or spectrum of operator.
- λ . eigenvalue of LNS operator.

- Λ_ϵ . pseudo-spectrum of an operator.
- λ_{ls} . least stable eigenvalue.
- \langle, \rangle . scalar product associated with the energy in the whole domain.
- $\langle \rangle_k$. Fourier component associated with wavenumber vector.
- \mathcal{D} . wall-normal derivative operator.
- \mathcal{E} . data-assimilation cost function.
- $\mathcal{H}(\omega)$. resolvent operator.
- \mathcal{L} . full OS-SQ system.
- \mathcal{L}_{OS} . Orr-Sommerfeld operator.
- \mathcal{L}_{SQ} . Squire operator.
- \mathcal{N} . quadratic nonlinearity operator.
- \mathcal{P} . operating representing the smoothing procedure.
- \mathcal{Q} . agreement metric between PIV data and the data-assimilation.
- \mathcal{R} . resolvent of a generic, nonsingular operator.
- \mathcal{S} . projection operator which artificially truncates PIV data.
- \mathcal{W} . wavemaker region.
- μ . eigenvalue of companion matrix.
- Ω . flow domain.
- ω . temporal frequency.
- ω_s . vortex shedding frequency.
- $\bar{\mathbf{u}}$. mean velocity field.
- $\bar{\mathbf{u}}_{exp}$. experimental mean velocity field.
- $\bar{\chi}$. complex weight or expansion coefficient of a resolvent response mode.
- \bar{p} . mean pressure field.
- ρ . density.
- σ . singular value of the resolvent operator.
- σ_1 . resolvent norm or the first singular value of the resolvent operator.

- τ_w . wall shear stress.
- Θ . spectral radius.
- $(\tilde{\cdot})$. stability analysis.
- $\tilde{\mathbf{g}}$. left eigenvector of a generic, nonsingular operator.
- $\tilde{\mathbf{g}}$. right eigenvector of a generic, nonsingular operator.
- $\varphi(x)$. area averaging operator.
- ξ . scalar associated with divergence part of mean forcing field.
- ζ . weight associated with a measurement in the data-assimilation algorithm.
- a . vortex acceleration.
- A_s . amplitude of shedding mode.
- AR . aspect ratio.
- C . a complex constant.
- c . wavespeed.
- D . cylinder diameter.
- d . mean shear in model operator.
- E . global error of flow reconstructed field.
- h . channel half-height.
- $I(t)$. instantaneous error for flow reconstruction.
- k_x . streamwise wavenumber.
- k_z . spanwise wavenumber.
- L . characteristic length scale.
- l_m . recirculation bubble length.
- m . complex eigenvalues of model operator.
- N . number of mesh points.
- n_s . number of stall cells.
- p . pressure field.
- p^\dagger . adjoint pressure field.

- P_0 . base flow pressure field.
- q . spatial wavelength of resolvent mode.
- Re . Reynolds number.
- Re_{chord} . chord-based Reynolds number.
- Re_{crit} . critical Reynolds number.
- U . characteristic velocity scale.
- U_c . vortex convection velocity.
- U_∞ . free-stream velocity.
- u_τ . friction velocity.
- z . generic complex number.
- LES**. Large-eddy simulation.
- LNS**. linear Navier-Stokes.
- OS**. Orr-Sommerfeld.
- SQ**. Squire.

Chapter 1

INTRODUCTION

1.1 Motivation

Practical implementation of closed-loop control of complex fluid systems has so far remained elusive due to a number of challenges. These include a lack of understanding of the dominant physics, the nonlinearity of the governing equations, and the wide range of scales inherent to the problem. Despite these obstacles, several studies such as Bewley et al. (2001) have successfully relaminarized wall-bounded turbulent flows using adjoint-based optimal control theory. The control algorithm in Bewley et al. (2001) had access to full state information or knowledge of all flow quantities at every grid point. This is not feasible experimentally so an estimator is necessary to deduce states which are not directly measured. Figure 1.1(a) is a schematic of a typical closed-loop flow control configuration where the estimator measures the flow and the controller actuates on it. The design of an effective estimator, however, is not straightforward as the questions of where to measure in the flow and which flow states should be measured have no clear answers. They are frequently flow dependent and may be sensitive to small changes in the flow configuration.

This thesis focuses on obtaining the maximum amount of information from the flow from the fewest measurements possible. To meet this challenge, fundamental questions about reduced-order modeling and data-assimilation are addressed. Reduced-order models capture the essential features of a flow and the advantage of such an approach is to gain insight of the flow physics in order to identify mechanisms for controlling the flow (Rowley and Dawson, 2017). If a suitable basis for the flow can be determined, then it can be reconstructed with reasonable fidelity using a small subset of basis functions. Another approach to flow reconstruction is data-assimilation whereby experimental measurements are merged with computational fluid dynamics (CFD) to improve prediction of real-world flows (Hayase, 2015) as seen in Figure 1.1(b). The underlying principle is to complement CFD, which lack full fidelity, with limited experimental measurements so that the simulation reflects the dynamics observed in the laboratory. Reduced-order modeling, data-assimilation, and the combination of both are addressed in this thesis.

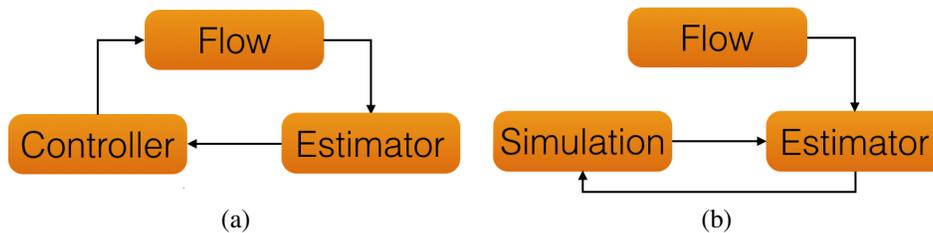


Figure 1.1: Comparison of the elements in a closed loop control arrangement (a) with data-assimilation (b).

1.2 A Musical Analogy

Before introducing background literature on the technical aspects of this thesis, the author would like to offer a musical analogy for reduced-order modeling and flow reconstruction. No musical knowledge is necessary to understand this analogy. In addition to being able to visualize music by looking at the notes written on a page or watching a pianist's hands move, for example, music can be heard. In fact, it is easier to judge whether the notes written on a page sound 'right' by hearing them being played. Thus, music has an edge over fluid mechanics when deciding whether or not a reduced-order model captures the essential components of a piece or song. One might ask when reduced-order modeling is necessary in the first place since there are no pieces written by composers which cannot be performed as long as there are enough musicians to cover the parts and they have the technical skills necessary to play it well. It turns out that reduced-order modeling is used all the time and the particular example that will be discussed below is arranging the orchestra accompaniment to a concerto for piano.

1.2.1 Concertos and Arranging

A concerto is a musical composition written for a solo instrument with the orchestra accompanying. The solo part tends to be virtuosic and technically demanding in order to showcase the soloist's mastery of the instrument. It also contains the principal themes and melodies of the piece. The orchestra, which is made up of many instruments, plays a supporting role by providing harmonies and background atmosphere. Concertos can be written for any instrument although the piano is the most common followed by the violin. Since the orchestra consists of as many as eighty people, it is not practical for a soloist auditioning on a piece to have full orchestra accompaniment. This would involve coordinating the schedule of many people, including a conductor, all of whom would probably expect payment for

practicing their parts and turning up to rehearsal! Instead, it is far more practical for the orchestra part to be arranged for piano and for the soloist to audition with a piano accompanist. As one might guess, the pianist cannot play the orchestra part to full fidelity so a reduced-order model is necessary to arrange the music such that it is playable by a single pianist.

Accomplishing such a task is not straightforward and requires pretty advanced knowledge of not only music theory but also the piano. To begin with, the pianist has only two hands which have a limited amount of reach. While it is possible to play many notes at once, the hand can maybe reach between an interval of nine keys on the piano. There are physical limitations, therefore, that the arranger needs to consider. The musical constraints are less obvious and are harder to satisfy by a brute force approach. For instance, if the arranger were to simply take all the orchestra parts and write them on the same page at once, it would immediately become obvious that some instruments are playing the same notes at the same time. Thus a part is being doubled and so the arranger has already accounted for two instruments. Occasionally the same notes at different frequencies are being played simultaneously — think of men and women singing Happy Birthday at an office party. They are singing the ‘same’ notes but at different pitches since men tend to have deeper voices than women. In physics, one would think of a fundamental frequency f_1 and higher harmonics $2f_1$, etc. In music, these are called octaves since $2f_1$ would be the eighth note after f_1 in a musical scale.

1.2.2 Structure and Statistics

Instead of turning to physics and governing equations such as Navier-Stokes, the arranger makes use of music theory. Figure 1.2 is the beginning of Schubert’s Impromptu No. 3 Opus 90 in G-flat major. This piece is for solo piano and one reads the music from left to right like a book. The pianist reads two lines together: the right-hand plays the upper staff, which is a collection of five horizontal lines, while the left-hand plays the bottom staff. The two staves are connected by vertical lines denoting measures, or musical subunits, and the notes appear on individual lines or spaces with respect to the staff. A high placement of a note on the staff denotes a high pitch and vice versa for low notes.

Three different colored boxes appear in Figure 1.2 to denote the three ‘structures’ which appear in the music. The part or voice in the solid red box is the melody which is played by the ring finger and pinky of the pianist’s right hand. The melody

The image shows the first two lines of Schubert's Impromptu No. 3 in G-flat Major, Opus 90. The music is in 3/4 time and G-flat major. The first line of music is marked 'Andante'. The melody is contained in a thick red solid box, the bass line in a dashed blue box, and the ornamentation in a thin brown dotted box. The thickness of the boxes indicates the relative volume at which each of these lines should be played.

Figure 1.2: The first two lines of Schubert's Impromptu No. 3 in G-flat Major Opus 90. The melody is contained in the red solid box, the bass line in the dashed blue box, and the ornamentation in the brown dotted box. The thickness of the boxes indicates the relative volume at which each of these lines should be played.

is the most important voice and should be played the loudest which is why the red box has the thickest lines. The bass is the voice contained in the dashed blue box and is played by the left hand. It is the second most important voice by supporting the melody and enriching the overall sound of the music. Last, and certainly least in terms of importance, is the ornamentation in the dotted brown box which has the thinnest lines. The purpose of the third voice is to add atmosphere and movement creating a shimmering effect. Incidentally, this voice is what makes the piece challenging since the ornamentation is played by the strong fingers (thumb, pointer, and middle) of the right hand while the melody is in the weaker fingers (ring and pinky).

The structure in complex flows is also quite important, as they typically account for a large amount of the fluctuating energy. While the structures in Figure 1.2 are relatively easy to identify given their spatial separation on the staff, identifying them in flows tends to be more challenging, warranting the use of modal decomposition techniques which are introduced in Section 1.3. Some of these make use of the statistics of the flow such as the time-averaged velocity field. The 'statistics' in music are the key and time signatures which are located at the very beginning of the piece. The key signature is made up of \sharp and \flat symbols, of which there are six of the latter in Figure 1.2, and denotes the key or scale the piece is based on. There are chords (combinations of notes) which are particular to each key and these often play an accompanying role in music. The time signature in Figure 1.2 is indicated by the two bold C's which are 'cut' by vertical lines and sets how many beats are

in a measure. Implicitly, this suggests what types of rhythms are likely to appear as well as which beats should be emphasized more than others.

In short, the voices in music can be likened to coherent structures in flows which have varying degrees of importance. The time and key signatures are similar to statistics which, using music theory or the governing equations, may yield useful information about the actual piece or flow.

1.2.3 Beethoven Violin Concerto

The arranger considers both the structure and statistics of the piece when attempting to build a reduced-order model of the orchestration. Two fragments of the Beethoven Violin Concerto are reproduced below to observe how the arranger took music which was originally written for full orchestra and simplified it for piano. In Figure 1.3, the parts for each instrument appear all together and the piano arrangement appears as the last two lines. The arranger has, in fact, written the original instruments next to the parts the piano is playing. In Figure 1.3, the piano is able to play almost all the notes played by the orchestra. Even though the reduced-order model (piano arrangement) very faithfully reconstructs the original orchestration, it loses the distinct sound, or timbre, of each instrument such as the oboe vs. the clarinet. It is impossible, therefore, for the piano to capture everything.

Another fragment of the piece is reproduced in Figure 1.4 where the piano arrangement is the final two lines and the orchestration is the top seven lines. Here, the arranger needs to make some choices since the pianist cannot play all the notes from the original orchestration. The first notes of lines 1 and 2, for example, are the same notes at different frequencies or octaves. The arranger has decided, therefore, to choose just line 1 which is played by the right-hand. The arranger has also simplified Lines 4-7 since the pianist's left-hand cannot play all four notes together at the speed written. To make the music playable, Line 4 is omitted since it is covered by the right-hand and the left-hand alternates between Lines 6/7 and Line 5. All the notes from the orchestration are being played, just at a slightly lower temporal frequency. This type of trick makes it much easier for the pianist since the playing of repeated sixteenth notes can be avoided.

There are many videos online which visualize arrangements of music for solo piano while they are being played. Thus one can hear and see a reduced-order model although a true assessment of its fidelity is to study the music. The rest of the thesis focuses on the tools and methods applied in fluid mechanics which rely on the same

Allegro, ma non troppo.
TUTTI.

Flauto.

Oboi.

Clarineti in A.

Fagotti.

Corni in D.

Trombe in D.

Timpani in D.A.

Violino principale.

Violino I.

Violino II.

Viola.

Violoncello e Basso.

Klavierauszug von AUG. GÖLLNER.

PIANO.

Figure 1.3: The opening measures of Beethoven's violin concerto. The instruments which play each line are labeled on the left-hand side. The piano arrangement of the orchestration is placed at the bottom and consists of two lines. The left-hand plays the bottom of these two lines and the right-hand plays the top line.

The image displays a musical score for a section of Beethoven's Violin Concerto. The score is divided into three main systems. The top system consists of two staves, both marked 'dolce' and 'Tutti.', representing the violin and viola parts. The middle system consists of seven staves, representing the orchestration, with a 'p' (piano) dynamic marking. The bottom system consists of two staves, representing the piano arrangement, with a 'p dolce' dynamic marking. The music is in 3/4 time and features a melodic line in the upper parts and a more rhythmic accompaniment in the lower parts.

Figure 1.4: A second excerpt of Beethoven's violin concerto where the top seven lines are the orchestration and the bottom two are the piano arrangement.

principles. Perhaps one advantage fluid mechanics has over music is that success or optimality can be measured mathematically. It will be seen, however, that the way optimal is defined might not yield the best answer so it might be necessary to make use of other tricks akin to the one used by the arranger to avoid repeated sixteenth notes.

1.3 Reduced-Order Modeling

There are numerous modal decomposition techniques which have been applied to analyze flows. In the overview of Taira et al. (2017), techniques which use flow data as an input are classified as *data-based* while techniques which rely on a more theoretical framework or discrete operators from the Navier-Stokes Equations (NSE) are classified as *operator-based*. One of the major advantages of the latter category is that large amounts of data are not necessary. This makes them an attractive choice

for flow reconstruction if only a small number of time-resolved measurements are available. It becomes important, consequently, to understand the conditions under which an operator-based technique identifies meaningful structures which actually appear in the flow. This section reviews recent progress and challenges with respect to coherent structure eduction with an emphasis on resolvent analysis using the time-averaged mean as an input.

1.3.1 Stability of the Mean Flow

A time-averaged flow, or mean, that is statistically stationary can often be defined and leveraged using the eigenvalue spectrum of the governing NSE to educe the frequencies, i.e. the imaginary part of the eigenvalues, and shapes of coherent structures which appear in the flow. Recent studies have demonstrated the success of mean flow stability analysis for a variety of flows including thermosolutal convection (Turton et al., 2015), turbulent jets (Gudmundsson and Colonius, 2011; Oberleithner et al., 2014; Schmidt et al., 2017a), and flow over a backward facing step (Beneddine et al., 2016). There is also a significant body of work discussing stability analysis of the mean cylinder wake which was shown by Barkley (2006) to correctly identify the frequency of the globally unstable flow above the critical Reynolds number of $Re = 47$ (Provansal et al., 1987; Sreenivasan et al., 1987; Noack and Eckelmann, 1994). Notably, classical linear stability analysis of the base flow, which is an equilibrium solution of the NSE, at supercritical Reynolds numbers does not predict the correct observed frequency. The base (laminar) and mean (time-average of the fluctuating velocity field) profiles are differentiated because of the importance of nonlinearity in sustaining the latter.

The two-dimensional von-Kármán vortex street becomes unstable to three-dimensional perturbations at a Reynolds number of $Re = 189$ (Barkley and Henderson, 1996; Williamson, 1996). A global stability analysis of the span-averaged mean wake continues to identify the shedding frequency as demonstrated by Leontini et al. (2010). Recent work has endeavored to explain why and when mean stability analysis is valid. Barkley (2006) suggested that success corresponds to cases where the Reynolds stresses are unperturbed at order ϵ when considering infinitesimal perturbations $\epsilon \tilde{\mathbf{u}}(x, y) \exp(\lambda t)$ to the mean flow solution. This was confirmed by Sipp and Lebedev (2007) who determined that the nonlinear interaction of the leading global mode with its conjugate, i.e. the contribution to the mean Reynolds stresses, significantly outweighed the interaction of the mode with itself leading to higher frequency harmonics. As shown by Mantič-Lugo et al. (2014) the Reynolds stresses

can be approximated with the leading global mode and its conjugate.

1.3.2 Resolvent Analysis

Sipp and Lebedev (2007) used open cavity flow as a counter example to the validity of mean stability analysis where the predicted frequencies do not match direct numerical simulation (DNS) of the flow. This discrepancy can be attributed to the non-normality of the flow which leads to non-orthogonality of the global modes and sensitivity of the spectrum to perturbation of the operator (Trefethen et al., 1993). The behavior of these systems can be more accurately characterized by the pseudospectrum of the LNS operator using resolvent analysis, e.g., Trefethen et al. (1993) and Schmid and Henningson (2001), rather than the spectrum alone. Jovanović and Bamieh (2005) formulated the linearized problem for laminar channel flow in input-output terms, where the resolvent operator constitutes the transfer function between them, considering the component-wise transfer from harmonic exogenous disturbance or forcing (input) to velocity response (output). There is also a broad literature considering stochastic forcing, e.g., Farrell and Ioannou, 1993, and the initial condition, transient growth problem, e.g., Butler and Farrell, 1992.

McKeon and Sharma (2010) and Hwang and Cossu (2010) considered the resolvent reformulated with respect to the turbulent mean flow for canonical turbulent wall flows. The latter authors employed an eddy viscosity to account for the action of the Reynolds stresses, while the former analysis extends the approach to include the nonlinear terms as the input forcing to the linear operator, i.e. closing the feedback loop. McKeon and Sharma (2010) performed a singular value decomposition of the resolvent to identify the inputs giving rise to the most amplified responses which are ranked by their gain (singular value). The approach has been extended to non-parallel flows, e.g., Lu and Papadakis (2014), Beneddine et al. (2016), Jeun et al. (2016), and Schmidt et al. (2017b). Beneddine et al. (2016) concluded that mean stability analysis was valid when the dominant singular value of the resolvent operator was significantly greater than the others at a given frequency and that this condition holds for flows where there is a dominant convective instability mechanism and an eigenvalue which is nearly marginally stable. In such circumstances, it was shown that the eigenmodes are proportional to the resolvent response modes. Stability and resolvent analyses are formally related in Chapter 3 by a dyad expansion of the resolvent operator.

1.3.3 Non-normality

Subsequent to the work of Jovanović and Bamieh (2005), Marquet et al. (2009) and Brandt et al. (2011) investigated the distribution of energy and phase between the velocity components of the most amplified input/output in analyses about laminar, base flows in recirculation bubbles and the flat plate boundary layer, respectively. These studies distinguished between component-type non-normalities, which distribute energy in different velocity components, and convective non-normalities, which separate the spatial support of forcing upstream of the response. The roots of these non-normalities are the mean shear and mean flow advection terms, respectively, in the linearized NSE. These terms also result in the Orr mechanism (Orr, 1907), which reorients upstream-leaning forcing modes with the mean shear such that the response modes are leaning downstream (Farrell, 1987). Chomaz (2005) quantified non-normality via the inner product between the most amplified input and output. A response dominated by non-normality results in a smaller inner product and this has an impact on the amplification mechanisms identified by the resolvent.

1.3.4 Instabilities and Pseudoresonance

The response of a system to harmonic input (forcing) can be classified as resonant or pseudoresonant. The latter occurs due to nonmodal effects associated with the sensitivity of the spectrum to perturbation (Trefethen et al., 1993). The former is generally an instability mechanism and corresponds to excitation in the vicinity of an eigenvalue. Instability mechanisms can be further classified as convective or absolute depending on the nature of the base or mean flow. A convective instability is one where perturbations grow downstream as they are swept away by the flow while an absolute instability is one where perturbations grow upstream and downstream of where they originated (Huerre and Monkewitz, 1985; Schmid and Henningson, 2001). The presence of reverse flow tends to result in a region of absolute instability (Rowley et al., 2002; Saponitsky et al., 2005; Juniper, 2012). The nature of the instability, thus, has a bearing on the strength of the component-type non-normality mentioned earlier and this is explored in Chapter 3.

1.3.5 Data-Driven Methods

Resolvent analysis also has relationships with data-driven methods such as Dynamic Mode Decomposition (DMD). Introduced by Schmid (2010), DMD extracts modes which are distinguished by their frequency content. Gómez et al. (2014)

demonstrated the similarity between resolvent modes and DMD modes corresponding to the same temporal frequency in turbulent pipe flow. DMD is also related to the Koopman operator (Rowley et al., 2009), or an infinite-dimensional linear operator associated with the full nonlinear system. This is significant since Sharma et al. (2016a) later noted that both DMD and resolvent modes may approximate the ‘true’ Koopman modes of the system.

Other connections between resolvent analysis and data-driven methods, in particular spectral proper orthogonal decomposition (SPOD), have been expounded upon by Towne et al. (2018). Originally introduced by Lumley (1970), SPOD results in modes which are orthogonal in space and time. They can be interpreted, therefore, as optimally averaged DMD modes which are obtained from an ensemble DMD problem. Connections were also drawn between SPOD and resolvent analysis. When the input forcing to the resolvent operator can be approximated as white-noise, resolvent modes are identical to SPOD modes. It has been noted by, e.g., Zare et al. (2017) that white-in-time stochastic forcing is insufficient to explain turbulent flow statistics and, in instances where the forcing is correlated, differences arise between the SPOD and resolvent modes. Schmidt et al. (2017b) demonstrated this phenomenon for low temporal frequency modes in a turbulent jet.

While SPOD is not considered here, its connection to DMD and resolvent analysis is useful as it suggests there are cases when the modes predicted by resolvent analysis will not match DMD modes computed directly from the data. The conditions under which these circumstances are likely to arise will be addressed. Since this has an impact on flow reconstruction as noted by Towne et al. (2018), the nature of the nonlinear forcing needs examination to correctly identify the structures in the flow when the forcing cannot be treated as white-noise (McKeon et al., 2013; McKeon, 2017).

1.4 Data-Assimilation

It has been noted in several studies (Nisugi et al., 2004; Suzuki et al., 2009a; Suzuki et al., 2009b; Suzuki, 2012; Foures et al., 2014) that despite recent advances in computational fluid dynamics (CFD) and experiments, both techniques have several disadvantages. Despite capturing the "true" physics of the flow, for example, experiments are corrupted by noise, limited by field of view, and have insufficient resolution to capture small scales. CFD, on the other hand, requires modeling assumptions about boundary conditions and sub-grid scale models unless there is sufficient

computational power to resolve all scales in turbulence. Data-assimilation (Lewis et al., 2006) is a technique whereby experimental measurements can be merged with computational fluid dynamics (CFD) to improve prediction of real-world flows (Hayase, 2015). The underlying principle is to complement CFD, which lack full fidelity, with experimental measurements, which typically lack full-field information, so that the simulation reflects the dynamics observed in the laboratory. The assimilated or hybrid flow is able to recover quantities in the experiment which would otherwise be inaccessible or difficult to measure such as pressure, vorticity, and Reynolds stresses, by reducing noise and improving resolution. It is also possible to extrapolate the flow beyond the experimental view by solving the equations on a larger domain.

Data-assimilation can be traced back to meteorology (Le Dimet and Talagrand, 1986) and is of particular interest to the experimental fluid mechanics community since it may be used to complete experimental observations by enforcing dynamical constraints (Heitz et al., 2010). One of the first hybrid simulations conducted by Nisugi et al. (2004) used offline, sequential assimilation for flow behind a square cylinder. By measuring the discrepancy between experimental and numerical pressure measurements at finite time intervals to drive the momentum equations, the simulation was altered to match the experiment. Sequential assimilation was greatly expanded by Suzuki et al. (2009a) and Suzuki et al. (2009b) when particle-tracking velocimetry (PTV) data of an airfoil at high angle of attack was fed into a two-dimensional direct numerical simulation (DNS). The resulting hybrid flows contained less noise and recovered the unsteady pressure fields. They also offered insight into the statistics of the mean flow and the three-dimensional instabilities which attenuate vorticity strength.

1.4.1 Variational Methods

Data-assimilation has also been extended to the NSE using a variational approach (Papadakis and Mémin, 2008; Gronskis et al., 2013; Foures et al., 2014) where, similar to optimal control, the objective is to minimize a cost function. This generally involves penalizing the distance between experimental and numerical velocity fields subject to governing equations. Ensemble Kalman filter or ensemble-based variational approaches (Colburn et al., 2011; Suzuki, 2012; Kato et al., 2015; Silva and Colonius, 2017) rely on the Kalman filter and its ensemble variant (Kalman, 1960; Evensen, 1994), which are appropriate when the data-assimilation problem is viewed from a stochastic perspective. The roots of variational data-assimilation

can be traced back to optimal control theory, which has been applied to various flow control problems (see Kim and Bewley, 2007, for an overview). Bewley et al. (2001) studied the control side of the problem by investigating various control strategies applied to turbulent channel flow simulated using a DNS. Data-assimilation, on the other hand, functions more closely to an estimator which reads in inputs from various sensors and fits them to an underlying model. The idea is to read in a sparse number of measurements and use the model to produce an estimated state which is more highly resolved in space and time.

Foures et al. (2014) used a variational method to minimize the discrepancy between the mean velocity fields of a DNS and an incompressible RANS simulation for flow around a circular cylinder at a Reynolds number of $Re = 150$. An improved estimation technique for mean flows has potential applications in mean flow modification studies. A large body of work has attempted to investigate this problem which has its roots in the experiments of Strykowski and Sreenivasan (1990), who showed experimentally that for low Reynolds numbers, a small control cylinder inserted in the wake behind a larger cylinder can completely suppress vortex shedding. Numerical studies including Giannetti and Luchini (2007) and Marquet et al. (2008) looked at the sensitivity of the cylinder instability to base flow modification and steady forcing near critical Reynolds number Re_{crit} . Meliga et al. (2012) and Mettot et al. (2014) expanded this framework to higher Reynolds numbers and determined how a small control cylinder could impact the frequency of vortex shedding as predicted by the most unstable global mode of the mean flow. Data-assimilation could expand control techniques to wall actuators such as an oscillating ribbon or synthetic jet which are difficult to model computationally due to ambiguous boundary conditions at the wall. It is possible, for example, to determine the mean flow from an experiment and recover a more highly resolved mean flow by tuning the boundary condition at the wall so that the simulated mean flow matches the experimental one.

A study conducted by Gronskis et al. (2013) illustrates this point quite well. They employed adjoint data-assimilation to generate initial and inflow conditions for a DNS of flow around a cylinder at a Reynolds number of $Re = 172$. The resulting simulation reflected the flow physics from large scale PIV measurements but contained far lower noise levels. Data-assimilation has also been demonstrated by Mons et al. (2016) to be applicable for perturbed fluid problems. They compared variational, ensemble Kalman filter-based, and ensemble-based variational data-assimilation techniques to reconstruct the flow around a cylinder subject to co-

herent gusts. They found that the variational data-assimilation approach produced the best results since the adjoint method can effectively capture the first-order sensitivity of the cost functional which penalizes the distance between experimental and numerical velocity fields.

1.4.2 Incorporating Pressure and Extension to Experiments

In this thesis, the algorithm of Foures et al. (2014) is modified to include mean pressure measurements since, without pressure data, only the solenoidal component of the forcing to the mean momentum equations can be recovered. The irrotational component is lumped into the mean pressure gradient term, which prevents recovery of the mean pressure field. It is possible to solve a Poisson equation for mean pressure using the RANS equations, e.g., Oudheusden (2013), but this relies on computing two gradients of the Reynolds stresses which suffer from noise contamination. One can ask whether limited mean pressure measurements can account for this problem and where in the flow do they have the greatest impact. Recent studies, e.g., Kang and Xu (2012), Mons et al. (2017), and Manohar et al. (2017), have investigated optimal placement of sensors for flow reconstruction and this is addressed in the thesis. The framework of Foures et al. (2014) is also adapted to mean flows obtained from experimental data at significantly higher Reynolds numbers. The mean profiles are obtained from time-averaged particle image velocimetry (PIV) data from a free-surface water tunnel. The necessary experimental parameters such as field of view or vector resolution for successful mean flow reconstruction are addressed.

1.5 Flow Reconstruction

A growing number of studies have investigated flow reconstruction using resolvent modes as a basis. Analysis of the linear operator via the singular value decomposition only does not yield the complex weights, or expansion coefficients, of the modes. This has to be calculated by projecting the resolvent forcing modes onto the nonlinear forcing. One approach adopted by Moarref et al. (2014) was to formulate a convex optimization problem which optimally reproduced the energy spectra of a turbulent channel flow. The study led to close agreement with the DNS spectra using only 12 modes per wall-parallel wavenumber pair and temporal frequency although the method tended to overpredict the streamwise energy and underpredict the wall-normal energy. Jeun et al. (2016) studied the input-output behavior of high-speed isothermal turbulent jets to improve the understanding of aeroacoustics. They

solved for the expansion coefficients by projecting the Fourier-transformed nonlinear term onto the resolvent forcing modes. With 24 modes, a substantial amount of the acoustic energy predicted by a large-eddy simulation (LES) could be recovered.

Resolvent modes have also been used to find a low-dimensional representation of exact coherent states of the NSE by Sharma et al. (2016b) and Rosenberg and McKeon (2018). In the latter study, the wall-normal velocity and vorticity fields could be decomposed into their Orr-Sommerfeld (OS) and Squire (SQ) contributions. This reduced the number of singular modes per wavenumber pair needed to represent the velocity fields. A Helmholtz decomposition of the nonlinear term, furthermore, highlighted the role of the solenoidal forcing which could be isolated to determine its contribution to the OS and SQ modes. It was then possible to obtain an efficient representation of the wall-normal velocity which was not possible using previous techniques.

A different approach outlined in Towne et al. (2015) and Towne (2016) is to reduce the nonlinear forcing experienced by wavepackets in a Mach 0.9 turbulent jet using empirical resolvent-mode decomposition. The objective was to identify and characterize the missing dynamics which were responsible for the failure of linear wavepacket models to predict acoustic radiation.

Other methods have bypassed consideration of the nonlinear term by taking advantage of a flow's low-rank behavior. In such circumstances, the resolvent modes provide an efficient basis for the fluctuations in the flow. Gómez et al. (2016a) and Beneddine et al. (2016), for example, used resolvent analysis of a lid-driven cavity and backward facing step, respectively, to determine the shapes of the velocity fluctuations at various temporal frequencies. Since the dominant singular value of the resolvent operator was sufficiently greater than all the others, the first resolvent mode could account for most of the fluctuation energy. It should be emphasized that this method works if the basis is good and the resolvent operator is low-rank. There are several examples in this thesis where this is not the case.

A single pointwise unsteady measurement was necessary to calibrate the complex amplitudes of the resolvent modes. This concept has also been applied by Gómez et al. (2016b) to estimate forces on a square cylinder, Beneddine et al., 2017 to experimental data of a jet, and Thomareis (2017) to DNS data of airfoils. Each of these studies noted that the robustness of the reconstruction depended on the location of the unsteady measurement. They concluded that the unsteady measurement should be within energetic regions of the flow. The quality of the reconstruction could be

improved by using multiple measurements.

1.6 Approach and Outline of Thesis

Figure 1.5 is a schematic which outlines the thesis and the procedure for reconstructing statistically stationary flows with as few measurements as possible. The first step is to collect experimental data consisting of a rudimentary, in terms of spatial resolution and field of view, mean profile and a single probe point which contains time-resolved information. Chapter 2 describes the experimental methods used to obtain the data and provides the mathematical background for the various modal decomposition techniques used as well as the data-assimilation framework. Chapter 3 is an original perspective on stability and resolvent analyses for base and mean flows with an emphasis on the latter. The analyses are formally related through a dyad expansion and the real part of an eigenvalue, which is difficult to interpret when the NSE are linearized around the mean flow, is shown to be important as it influences the degree to which a disturbance is amplified. It also has a bearing on whether or not the resolvent operator is low-rank since an eigenvalue must be sufficiently separated from the rest of the spectrum in order to dominate over the contribution of other eigenvalues in the dyad expansion. Non-normality plays a role in amplification and is investigated through the lens of the pseudospectrum (e.g. Trefethen et al., 1993; Reddy et al., 1993; Trefethen et al., 1999; Schmid and Henningson, 2001; Schmid, 2007; Schmid and Brandt, 2014) of the LNS operator.

When an eigenvalue is marginally stable, or very close to the imaginary axis, it drowns out the effect of other eigenvalues over a large range of temporal frequencies. It is shown in Chapter 4 that the most amplified structure for a cylinder at a Reynolds number of $Re = 100$ over this range of frequencies is a stretched or compressed version of the shedding mode. The dominance of this mode leaves a significant footprint on the mean profile whose geometry scales with the shedding frequency for Reynolds numbers $60 < Re < 320$. Similar to Dergham et al. (2013) for a backward facing step, various branches of singular values are identified for the cylinder including one for the shedding mode and another for free-stream modes which are less amplified near the shedding frequency. At temporal harmonics, however, the resolvent predicts a structure which is completely different from its DMD counterpart, suggesting that the structure of the nonlinear forcing has a significant influence on the structure amplified by the resolvent.

Resolvent analysis requires explicit knowledge of the mean profile yet it can be tricky to obtain from experiments, particularly for flows around bodies where it is difficult to obtain measurements near the surface. In Chapter 5, flow around a circular cylinder at a Reynolds number of $Re = 100$ is chosen to investigate measurements for data-assimilation. The results of Mantič-Lugo et al. (2014) are the motivation for choosing a Reynolds number lower than Foures et al., 2014, where $Re = 150$, since it is possible to obtain a self-consistent model that predicts the amplitude of the most unstable eigenmode in the flow. There are connections between this approach and the effect of data-assimilation which are drawn out in Chapter 5, in which the main objective is to identify the domain where measurements are needed for successful data-assimilation of the mean flow. The quality of the reconstruction is assessed by comparing the mean forcings identified by the optimization to their counterparts from DNS. Resolvent analysis assists by identifying the energetic structures which interact to produce the necessary Reynolds stresses to sustain the mean profile. Additionally, it identifies information about the scaling of the mean profile and how the dominant amplification mechanisms are sustained by nonlinear interactions, both of which guide how to choose and weight measurements.

In Chapter 6, the data-assimilation algorithm is adapted for experimental data at higher Reynolds numbers. The complexity of flows is gradually increased from symmetric flows, i.e., airfoils at zero angle of attack, to airfoils at angles of attack where the flow is stalled. The minimum resolution of data is assessed for an idealized airfoil. Once this information is known, the objective is to reconstruct the flows around a NACA 0018 airfoil in Chapter 7. The parameters related to the PIV data such as resolution, domain, etc. are selected based on the results from Chapters 5 and 6. Chapters 3 and 4, where the fundamentals of resolvent analysis are discussed, guide the choice of temporal frequencies where the resolvent operator identifies the correct coherent structures in the flow. Nonlinear interactions of amplified modes may need to be considered at other temporal frequencies. Finally, the probe is used to calibrate the resolvent modes and reconstruct all fluctuating quantities as well as the mean pressure in Chapter 7. The thesis concludes in Chapter 8 and some avenues for future work are proposed.

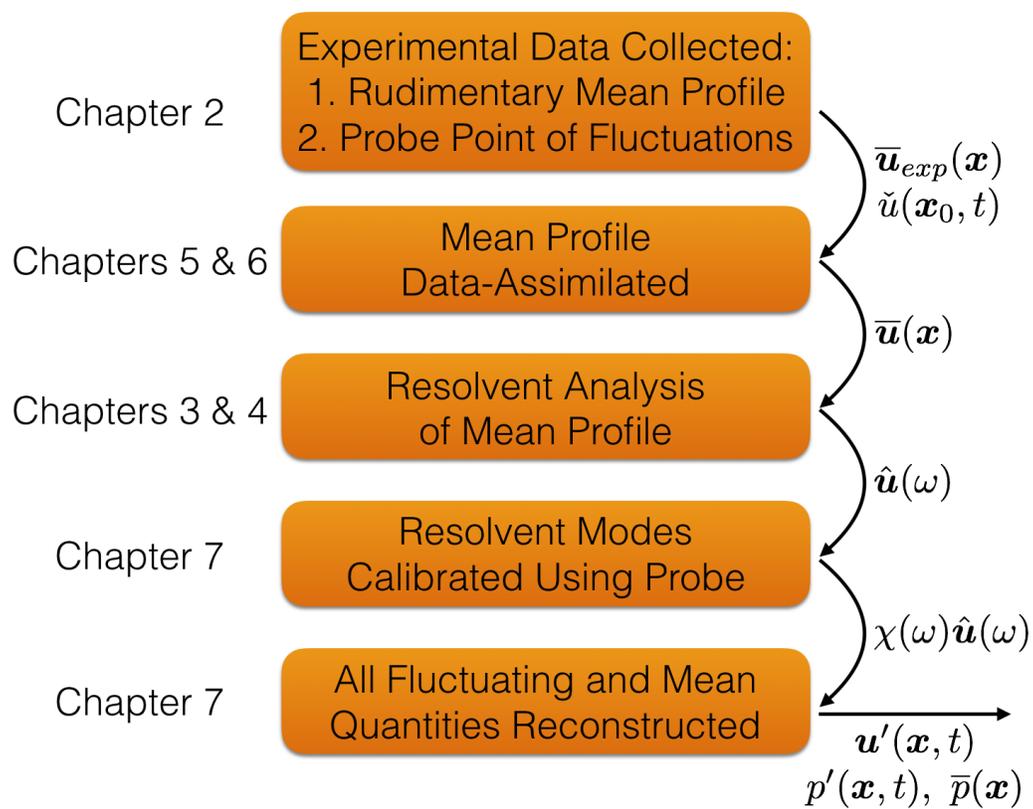


Figure 1.5: Schematic of flow reconstruction approach. The left-hand side indicates the chapters where the steps are discussed.

Chapter 2

METHODS

This chapter describes the numerical and experimental methods used to study the flows of interest. The two- and three-dimensional flows around a circular cylinder are solved using direct numerical simulation (DNS). The flows around an idealized airfoil and a NACA 0018 airfoil are acquired experimentally using particle image velocimetry (PIV). The mean profiles for turbulent channel flow are obtained from an eddy viscosity model. The modal decomposition techniques which are applied to these flows include dynamic mode decomposition (DMD), stability analysis, and resolvent analysis. Finally, the data-assimilation technique, which recovers a more highly resolved mean profile from incomplete measurements, is described.

2.1 Governing Equations

The flows are governed by the incompressible Navier-Stokes Equations (NSE) which are non-dimensionalized by the characteristic length and velocity scales, L and U :

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + Re^{-1} \nabla^2 \mathbf{u} \quad (2.1a)$$

$$\nabla \cdot \mathbf{u} = 0. \quad (2.1b)$$

The states $\mathbf{u}(\mathbf{x}, t)$ and $p(\mathbf{x}, t)$ are the spatially- and temporally-varying velocity and pressure fields, respectively (explicit statement of the spatial and temporal dependences will be dropped hereon for conciseness), while Re is the Reynolds number based on L and U . Steady-state solutions to the NSE are referred to as *base flows* and satisfy

$$\mathbf{U}_0 \cdot \nabla \mathbf{U}_0 = -\nabla P_0 + Re^{-1} \nabla^2 \mathbf{U}_0 \quad (2.2a)$$

$$\nabla \cdot \mathbf{U}_0 = 0, \quad (2.2b)$$

where the states \mathbf{U}_0 and P_0 are the base flow velocity and pressure, respectively.

2.2 Numerical Methods for Cylinder Flow

2.2.1 2D Cylinder

The NSE (Equation 2.1) are non-dimensionalized by the cylinder diameter D and inlet velocity U_∞ . For the base flow calculation, a uniform inlet velocity condition is prescribed while no-slip Dirichlet boundary conditions are applied to the

cylinder surface, symmetric conditions to the upper and lower boundaries, and advective conditions to the outlet. The nonlinear equations for U_0 are solved using a Newton method on a finite-element mesh generated by FreeFem++ (see Hecht, 2012). Taylor-Hood finite elements (P1b, P1b, P1 for U_0 , V_0 , and P_0 respectively) are used for spatial discretizations. The computational domain Ω spans $-30 \leq x/D \leq 60, -25 \leq y/D \leq 25$ with the cylinder centered at the origin and the mesh is made up of 104,214 triangles resulting in 365,358 degrees of freedom for velocity and pressure.

The unsteady flow is obtained with a DNS for $47 < Re < 190$ using the same boundary conditions and mesh. A second-order semi-implicit time discretization is employed with a non-dimensional time step $\Delta t = 0.02$. Beyond Re_{crit} , the simulated flow settles into regular vortex shedding at a fixed amplitude A_s and temporal frequency ω_s where the subscript s denotes shedding. The mean flow $\bar{\mathbf{u}}$ is computed by time-averaging the DNS state vector over 25 complete shedding cycles.

2.2.2 3D Cylinder

Three-dimensional simulations are performed using the spectral element-Fourier solver of Blackburn and Sherwin (2004) for $190 \leq Re \leq 320$. The code uses a spectral-element method for the spatial differencing in the cylinder plane and a Fourier decomposition for the spatial differencing in the spanwise direction. The domain has a blockage ratio of 4% with the upstream and side boundaries located $12D$ from the center of the cylinder and the outflow positioned $25D$ downstream of the cylinder. The spanwise extent is $12D$ and 48 Fourier planes are used. The mesh consists of 196 spectral elements and seventh-order Lagrangian interpolating polynomials. The flow is averaged in the spanwise direction once the flow has reached an asymptotic state. It is then averaged in time over about 60 shedding cycles since the flow is quasi-periodic.

2.2.3 Results and Validation

The DNS results are validated with respect to the literature. The shedding frequency ω_s , which is determined from the peak of the power spectrum, has been plotted in Figure 2.1 alongside experimental results from Williamson (1992). The agreement in the two-dimensional regime ($Re < 190$) is excellent although there are minor discrepancies in the three-dimensional regime. The simulations successfully distinguish between the two modes which characterize the three-dimensional transition of the flow (Williamson, 1988). Mode A results in a discontinuous drop of the shed-

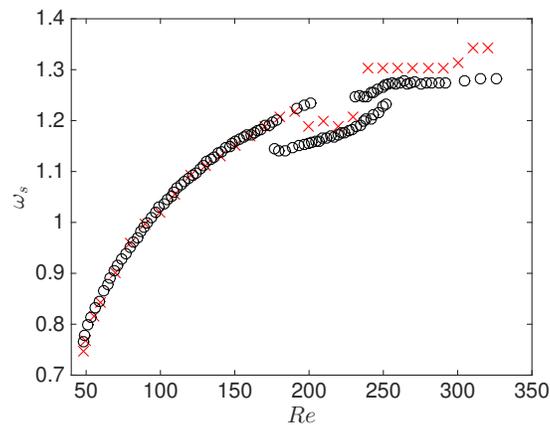


Figure 2.1: Shedding frequency ω_s as a function of Reynolds number Re : red crosses denote the frequency obtained from a Fourier-transform of the 2D/3D DNS at a point in the wake and black circles are from Williamson (1992).

ding frequency for $190 \leq Re < 230$ before it reverts back to larger values when Mode B takes over for $Re \geq 240$.

The mean profile and second-order statistics for $Re = 150$ are presented in Figure 2.2 to be compared against the results from Foures et al., 2014. Figure 2.2(a) shows contours of the mean streamwise velocity \bar{u} including the $\bar{u} = 0$ contour in green. The length of the recirculation zone agrees with Leontini et al. (2010). The mean pressure \bar{p} is plotted in Figure 2.2(b) and agrees with the mean pressure in Foures et al. (2014). The x - and y - components of the mean forcing \mathbf{f} are plotted in Figures 2.2(c) and (d), respectively, where \mathbf{f} corresponds to the divergence of the Reynolds stress tensor \mathbf{R} , defined as

$$\mathbf{f} = -\nabla \cdot \mathbf{R}, \quad \text{with } R_{ij} = \overline{u'_i u'_j}. \quad (2.3)$$

u' and v' denote the streamwise and transverse fluctuating velocity components, respectively. Both f_x and f_y match with their counterparts in Foures et al. (2014).

2.2.4 Mesh Size and Resolution

The effects of mesh size and resolution are assessed to guarantee that they do not influence the simulation results. The two-dimensional simulations are carried out on a smaller mesh of size $-10 < x < 25 \cup -7 < y < 7$. The shedding frequency and length of the recirculation bubble agree to within 3% of the values obtained from the larger mesh. The resolution of the larger mesh is augmented by a factor of 1.5 by increasing the vertex density along each boundary in the domain. The

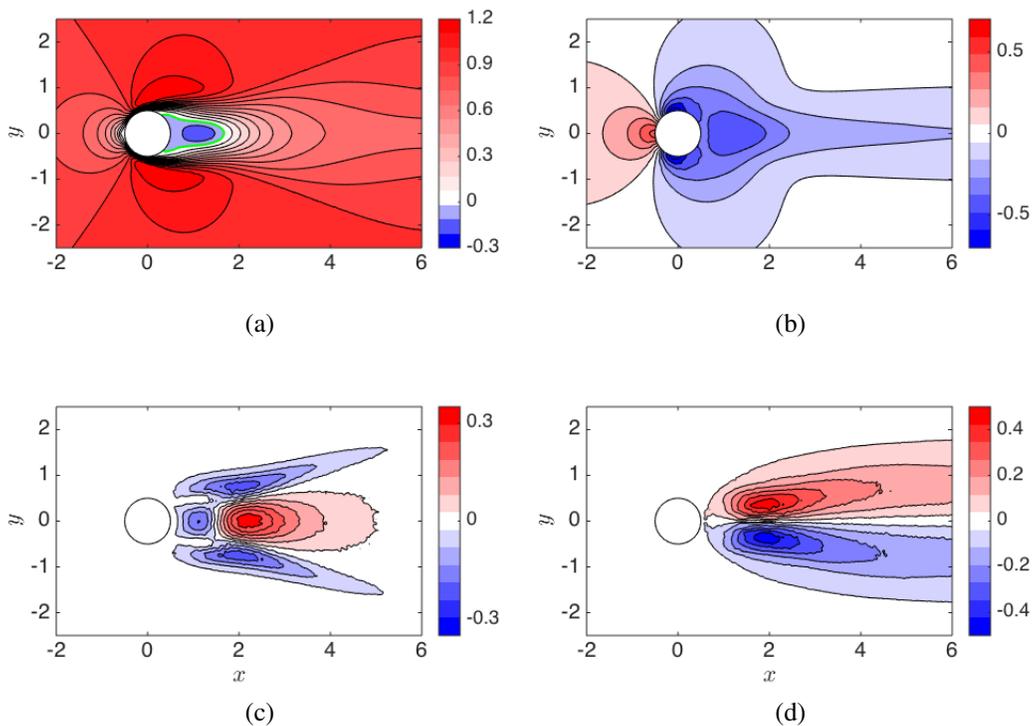


Figure 2.2: The time-averaged 2D DNS results for $Re = 150$: mean streamwise velocity component \bar{u} in (a) and mean pressure \bar{p} in (b). The two components of the mean forcing are plotted in the second row: f_x in (c) and f_y in (d).

impact on the shedding frequency and length of the recirculation bubble is negligible. The resolution effects of the three-dimensional simulations are examined by doubling the number of Fourier planes used in the spanwise direction, elongating the domain in the spanwise direction, and increasing the order of Lagrangian interpolating polynomials. The impact on the measured statistics for all three changes is found to be minor as long as the spanwise domain length is at least $8D$ and the number of Fourier planes is at least 32.

2.3 Experimental Methods

2.3.1 Idealized Airfoil

Experiments are performed on an idealized airfoil (see Figure 2.3) with a chord-length of 15.86 cm, a width of 3.43 cm, and a spanwise extent of 50.8 cm. The airfoil is symmetric about the chord and consists of a cylindrical leading edge followed by two plane surfaces connected at the trailing edge, which has a thickness of 0.15 cm. The diameter of the cylinder is equal to the width of the airfoil so that the

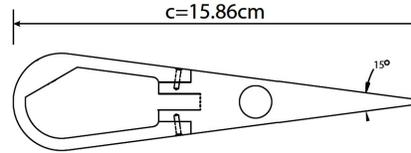


Figure 2.3: Schematic of the idealized airfoil adapted from Gonzalez et al. (2010).

junction between the cylinder and planar surfaces is half a diameter aft of the leading edge. At a zero angle of attack and a chord-based Reynolds number greater than approximately $Re_{chord} = 10,000$, the flow separates at the transition point between the cylinder and the plane surfaces and the time-averaged flow reveals recirculation bubbles that form on both sides of the airfoil (Wallace and McKeon, 2012). The critical Reynolds number, which coincides with the onset of von-Karman vortex shedding, is measured experimentally to be approximately $Re_{crit} = 2000$. Furthermore, the shear layers which form around the separation bubbles are convectively unstable, giving rise to the formation of Kelvin-Helmholtz vortices. These two instabilities ensure dynamics which are reminiscent of the behavior of separated flows at far higher Reynolds numbers (Prasad and Williamson, 1997). Consequently, this configuration is an attractive choice for studying the capability of data-assimilation to capture the flow dynamics around aerodynamic geometries.

Experiments are conducted in a free-surface water facility (see Wallace and McKeon, 2012). The test section measures 1.6 m in length, 0.46 m in width, and 0.5 m in height. The airfoil is mounted vertically so that its span is parallel to the test section height. The flow is conditioned by a perforated plate, a honey-comb mesh, three turbulence-reducing screens, and a 4-to-1 fifth-order-polynomial contraction (Gharib, 1983). The free-stream velocity is 8.1 cm/s and the free-stream turbulence intensity is less than 0.1% at the centerline. The water temperature is 23°C, which results in a chord-based Reynolds number of $Re_{chord} = 13,500$.

A LaVision time-resolved 2D-PIV setup is used consisting of two Photron Fastcam APX-RS high-speed cameras with 50 mm focal length Nikon lenses and 1:1.2 aperture. The cameras are synchronized with a high-speed controller and sample the flow at a frame rate of 83 Hz. The camera resolution is 1024×1024 pixels and the cameras are calibrated at 5.25 px/mm. The snapshot frequency is selected to guarantee a particle displacement between 5 and 7 pixels between any two consecutive snapshots. The seeding particles are hollow glass spheres (reference 110P8 with an average diameter of $11.7 \mu\text{m}$ and a specific gravity of 1.1) and the seeding

density is about 0.1 particles per square pixel. The particles are illuminated by a 2 mm-thick laser sheet provided by a Photonics DM20-527 solid-state laser. In an effort to avoid large uncertainty near the illuminated profile due to surface reflection, the image intensity is calibrated using white-image subtraction (normalization of the image intensity using the average light distribution) and background-image subtraction. The white and background images are taken before each run and averaged over 100 snapshots. The camera view, shown in Figure 2.4, encompasses the flow from 6.5 cm upstream of the leading edge to 10.5 cm downstream of the trailing edge with a 15% overlap in order to include a large area of mean flow and fluctuation measurements. Finally, 10240 instantaneous flow fields are captured over 5 runs (2048 snapshots per run), which represents approximately 35 complete vortex shedding cycles.

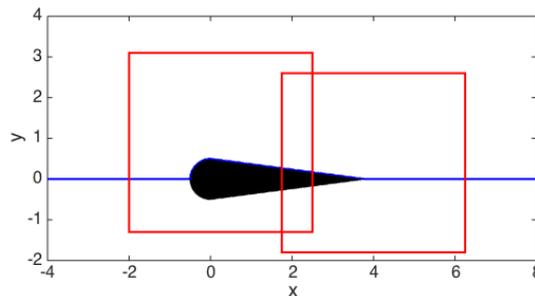


Figure 2.4: Experimental setup of the flow around an idealized airfoil showing the spatial coverage of the flow by the cameras. The dimensions are normalized by the diameter of the cylinder. Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

2.3.2 NACA 0018 Airfoil

PIV data are also collected for a NACA 0018 airfoil with a chord length of 10 cm and a spanwise extent of 48 cm resulting in an aspect ratio $AR = 4.8$. Gerakopoulos et al. (2010) studied this flow at Reynolds numbers between $8 \times 10^4 \leq Re \leq 2 \times 10^5$ and determined that the stall angle α_s is between 10° and 14° , increasing with Reynolds number. On the suction side, there is a laminar separation bubble which decreases in size with Reynolds number. Flow separation may result in 3D stall cells and significant out-of-plane velocity. Weihs and Katz (1983) found that the number of stall cells is $n_s = AR/2.28$, which would indicate approximately two stall cells appearing for this experimental setup. The airfoil is mounted vertically in the tunnel at nine different angles of attack ranging from $\alpha = -10^\circ$ to $\alpha = 10^\circ$ in 2.5° increments. The results for $\alpha = 0, \pm 10^\circ$ only are reported in this thesis.

In addition, the airfoil is connected to an ATI Mini 40 IP68 force transducer at its quarter-chord point to measure the forces and torques on the airfoil over time. The force data were collected over 100 seconds at a sampling frequency of 5000 Hz. These data are not reported in this thesis but will be used in the future to assess whether the pressure fields from the reconstruction procedure yield accurate predictions of the forces on the airfoil.

The PIV setup is slightly different from the idealized airfoil case as it consists of two Phantom Miro 320 cameras with 50 mm focal length Nikon lenses and 1:1.8 aperture. They have an overlap of 18% in the streamwise direction only. The cameras sample the flow at a frequency of 125 Hz for $Re_{chord} = 10,250$ and 250 Hz for $Re_{chord} = 20,700$. Each run consists of 3,500 snapshots; two runs (7,000 snapshots) are collected for the lower Reynolds number case while three runs (10,500 snapshots) are collected for the higher Reynolds number case. The camera resolution is 1920×1200 pixels and they are calibrated at 8.2 px/mm. The laser sheet is provided by a YLF dual cavity solid-state laser and is centered at a height of 220 mm which is in between the two stall cells. The out-of-plane velocity component is less than 4% of the free-stream velocity as documented by Dunne (2016).

2.3.3 Vector Post-Processing

The computation of velocity vectors is performed using the software package DaVis provided by LaVision. A standard cross-correlation technique via Fast Fourier Transformation (see Adrian, 1991) is applied to each sequential image with a window-size reduced from 32×32 px² to 16×16 px² over three passes, a 50% overlap, and a 2:1 elliptic weight (see Kompenhans et al., 2007). Finally, the data are post-processed and single missing vectors are interpolated using an average of all the non-zero neighborhood vectors. A median filter, as described by Westerweel and Scarano (2005), is used for outlier detection. The snapshots are then averaged to obtain the mean velocity field and Reynolds stresses.

2.4 Modal Analysis Methods

Three modal decomposition techniques are applied (see Taira et al., 2017, for more details) to the numerical and experimental datasets. The first is stability analysis, which assumes small perturbations to the linearized NSE. An eigenanalysis of the linear Navier-Stokes (LNS) operator yields the spatial modes which are most prone to instability. The second is resolvent analysis which analyzes the LNS operator from an input-output perspective. The third is DMD (Schmid, 2010; Schmid, 2011;

Rowley et al., 2009), a method proposed by Schmid (2010) to extract spatial modes, each of which has a complex eigenvalue, that are orthogonal in time. The real part of the eigenvalue describes the rate at which the mode grows or decays while the imaginary part is the frequency at which it oscillates.

The development considers global and spatially periodic modes in the context of flows around bodies and turbulent channel flow, respectively. The mean profile is used as an input to resolvent analysis and it can be obtained experimentally from PIV. To compute global modes, however, requires a large domain to adapt to the upstream and downstream boundary conditions. The resolution of PIV near the surface of the body, furthermore, may not be sufficient to accurately capture the resolvent modes. For now, it is assumed that the temporally-averaged (mean) velocity profiles are known throughout the domain from numerical simulation or experiment. In the turbulent channel flow case, the mean profiles are obtained via an eddy viscosity model (e.g. Reynolds and Tiederman, 1967). The following derivations concern the general case in which there exists invariance only in time. In other words, the analysis is performed in the frequency domain, such that mode shapes may be functions of all three spatial dimensions.

2.4.1 Stability Analysis

In temporal stability analysis, the velocity field is decomposed into a base flow plus a perturbation

$$\mathbf{u} = \mathbf{U}_0 + \epsilon \mathbf{u}', \quad (2.4)$$

where $\epsilon \ll 1$. At $O(\epsilon^0)$, this simply reduces to Equation 2.2, which governs the base flow. At $O(\epsilon^1)$, the equations governing the perturbation are obtained:

$$\partial_t \mathbf{u}' + \mathbf{U}_0 \cdot \nabla \mathbf{u}' + \mathbf{u}' \cdot \nabla \mathbf{U}_0 + \nabla p' - Re^{-1} \nabla^2 \mathbf{u}' = 0 \quad (2.5a)$$

$$\nabla \cdot \mathbf{u}' = 0. \quad (2.5b)$$

Nonlinear terms are neglected since they are $O(\epsilon^2)$. Assuming perturbations of the form $\mathbf{u}' = \tilde{\mathbf{u}} e^{\lambda t}$ yields

$$\lambda \tilde{\mathbf{u}} = -\mathbf{U}_0 \cdot \nabla \tilde{\mathbf{u}} - \tilde{\mathbf{u}} \cdot \nabla \mathbf{U}_0 - \nabla \tilde{p} + Re^{-1} \nabla^2 \tilde{\mathbf{u}} \quad (2.6a)$$

$$\nabla \cdot \tilde{\mathbf{u}} = 0, \quad (2.6b)$$

where a tilde will be used to denote stability analysis. The real part of the eigenvalue λ is the growth/decay rate while the imaginary part is the temporal frequency.

Equation 2.6 is recast into operator form to obtain

$$\lambda \mathbf{B} \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{p} \end{pmatrix} = \mathbf{A} \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{p} \end{pmatrix}, \quad (2.7)$$

where \mathbf{A} is the LNS operator with respect to the base flow,

$$\mathbf{A} = \begin{pmatrix} -\mathbf{U}_0 \cdot \nabla() - () \cdot \nabla \mathbf{U}_0 + Re^{-1} \nabla^2() & -\nabla() \\ \nabla \cdot () & 0 \end{pmatrix}, \quad (2.8)$$

and

$$\mathbf{B} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}. \quad (2.9)$$

Equation 2.6 is the base flow stability analysis case in Table 2.1. The flow is linearly stable when the real part of all eigenvalues of operator \mathbf{A} are negative and unstable if the real part of at least one eigenvalue is positive. If the real part of the least stable eigenvalue is zero, the flow is said to be marginally stable.

The adjoint NSE have been derived by, e.g., Luchini and Bottaro (2014), and the linearized operator \mathbf{A}^* for the adjoint variables $\tilde{\mathbf{u}}^\dagger$ and \tilde{p}^\dagger is

$$\mathbf{A}^* = \begin{pmatrix} \mathbf{U}_0 \cdot \nabla() - () \cdot (\nabla \mathbf{U}_0)^* + Re^{-1} \nabla^2() & \nabla() \\ \nabla \cdot () & 0 \end{pmatrix}. \quad (2.10)$$

The operator \mathbf{A}^* satisfies

$$\langle \tilde{\mathbf{u}}, \mathbf{A} \tilde{\mathbf{u}}^\dagger \rangle = \langle \mathbf{A}^* \tilde{\mathbf{u}}, \tilde{\mathbf{u}}^\dagger \rangle, \quad (2.11)$$

where $\langle \cdot, \cdot \rangle$ is the scalar product associated with the energy in the whole domain. Throughout the thesis, the continuous adjoint formulation is used, signifying that the adjoint operator \mathbf{A}^* is obtained before discretizing the equations, e.g., Chandler et al. (2012). An alternative approach is to use the discrete adjoint formulation, where the adjoint is determined after discretizing the equations. The benefit of the latter approach is the guarantee that the eigenvalues of both the forward and adjoint operator are equal. This is not the case for the continuous adjoint formulation although the eigenvalues gradually converge as the mesh density is increased. The downside, as discussed by Chandler et al. (2012), is that the global modes may contain numerical artifacts. It is also more difficult to derive the discrete adjoint equations. The discrepancy between the forward and adjoint eigenvalues is small (on the order of 0.01%) for the flows studied in this thesis.

For a general operator, \mathcal{T} , that is normal, i.e. $\mathcal{T} \mathcal{T}^* = \mathcal{T}^* \mathcal{T}$, the eigenvectors of \mathcal{T} corresponding to distinct eigenvalues are orthogonal although the eigenvalues may

be complex. Self-adjoint operators ($\mathcal{T} = \mathcal{T}^*$), on the other hand, have orthogonal eigenvectors and real eigenvalues. In general, the LNS operator \mathbf{A} is neither self-adjoint nor normal which accounts for the differences between Equations 2.8 and Equations 2.10 and their influence is discussed in Chapter 3.

The NSE can also be linearized about the mean flow by Reynolds-decomposing the velocity field into a temporal mean (denoted by an overline) and a fluctuating component (denoted by a prime). The mean flow obeys the following set of equations:

$$\bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} + \nabla \bar{p} - Re^{-1} \nabla^2 \bar{\mathbf{u}} = -\overline{\mathbf{u}' \cdot \nabla \mathbf{u}'} \quad (2.12a)$$

$$\nabla \cdot \bar{\mathbf{u}} = 0. \quad (2.12b)$$

Subtracting the mean momentum equations (Equation 2.12) from the NSE (Equation 2.1) results in the equations governing the fluctuations:

$$\partial_t \mathbf{u}' + \bar{\mathbf{u}} \cdot \nabla \mathbf{u}' + \mathbf{u}' \cdot \nabla \bar{\mathbf{u}} + \nabla p' - Re^{-1} \nabla^2 \mathbf{u}' = -\mathbf{u}' \cdot \nabla \mathbf{u}' + \overline{\mathbf{u}' \cdot \nabla \mathbf{u}'} = \mathbf{f}' \quad (2.13a)$$

$$\nabla \cdot \mathbf{u}' = 0. \quad (2.13b)$$

Equation 2.13a has been written such that all linear terms appear on the left-hand side. They can be recast as the linear operator

$$\mathbf{L} = \begin{pmatrix} -\bar{\mathbf{u}} \cdot \nabla() - () \cdot \nabla \bar{\mathbf{u}} + Re^{-1} \nabla^2() & -\nabla() \\ \nabla \cdot () & 0 \end{pmatrix}, \quad (2.14)$$

which is the LNS operator about the mean flow. Mean stability analysis corresponds to an eigenanalysis of the LNS operator linearized about the mean, i.e.,

$$\lambda \mathbf{B} \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{p} \end{pmatrix} = \mathbf{L} \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{p} \end{pmatrix}. \quad (2.15)$$

The interpretation of this analysis, which is summarized in Table 2.1, is not straightforward since the mean is not a solution of the NSE; its significance will be discussed in Chapter 3. Similar to the base flow case, it is possible to derive the adjoint LNS operator

$$\mathbf{L}^* = \begin{pmatrix} \bar{\mathbf{u}} \cdot \nabla() - () \cdot (\nabla \bar{\mathbf{u}})^* + Re^{-1} \nabla^2() & \nabla() \\ \nabla \cdot () & 0 \end{pmatrix}, \quad (2.16)$$

which satisfies the inner product

$$\langle \tilde{\mathbf{u}}, \mathbf{L} \tilde{\mathbf{u}}^\dagger \rangle = \langle \mathbf{L}^* \tilde{\mathbf{u}}, \tilde{\mathbf{u}}^\dagger \rangle. \quad (2.17)$$

Like its base flow counterpart, L is generally non-normal, the implications of which are discussed in Chapter 3.

The linear operators for stability analysis are formed in FreeFem++. Homogeneous boundary conditions are enforced at the inlet so that the perturbations vanish at infinity. The other boundary conditions for the domain are identical to those applied in the base flow calculation. The eigenvalues are computed using a shift-and-invert strategy, the details of which are discussed in Nayar and Ortega (1993). The generalized eigenvalue problem is then solved with the Implicitly Restarted Arnoldi method using the ARPACK library developed by Lehoucq and Sorensen (1996).

	Base Flow	Mean Flow
Stability Analysis	$\lambda \mathbf{B} \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{p} \end{pmatrix} = \mathbf{A} \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{p} \end{pmatrix}$	$\lambda \mathbf{B} \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{p} \end{pmatrix} = \mathbf{L} \begin{pmatrix} \tilde{\mathbf{u}} \\ \tilde{p} \end{pmatrix}$
Resolvent Analysis	$i\omega \mathbf{B} \begin{pmatrix} \hat{\mathbf{u}} \\ \hat{p} \end{pmatrix} = \mathbf{A} \begin{pmatrix} \hat{\mathbf{u}} \\ \hat{p} \end{pmatrix} + \mathbf{C} \hat{\mathbf{f}}$	$i\omega \mathbf{B} \begin{pmatrix} \hat{\mathbf{u}} \\ \hat{p} \end{pmatrix} = \mathbf{L} \begin{pmatrix} \hat{\mathbf{u}} \\ \hat{p} \end{pmatrix} + \mathbf{C} \hat{\mathbf{f}}$

Table 2.1: Operator form of the equations for stability and resolvent analyses. Variables with a tilde correspond to stability analysis while a caret indicates resolvent analysis.

2.4.2 Resolvent Analysis

Unlike stability analysis, the role of nonlinearity is retained in resolvent analysis by treating it as a forcing to the linear dynamics of the NSE. The nonlinear terms on the right-hand side of Equation 2.13a are lumped together as a forcing \mathbf{f}' without loss of generality. Endogenous nonlinear terms or exogenous forcing can be treated equally well via \mathbf{f}' , although the interpretation of the resulting system is different. For a harmonic forcing and response at temporal frequency ω , i.e.,

$$\mathbf{f}' = \hat{\mathbf{f}} e^{i\omega t}, \quad \mathbf{u}' = \hat{\mathbf{u}} e^{i\omega t}, \quad (2.18)$$

Equation 2.13a can be rewritten as

$$\hat{\mathbf{u}} = \mathcal{H}(\omega) \hat{\mathbf{f}}, \quad (2.19)$$

where the caret denotes that the fluctuation is associated with a resolvent analysis. The resolvent operator $\mathcal{H}(\omega)$ is given by

$$\mathcal{H}(\omega) = \mathbf{C}^T (i\omega \mathbf{B} - \mathbf{L})^{-1} \mathbf{C}, \quad (2.20)$$

where

$$\mathbf{C} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad (2.21)$$

restricts the forcing to velocity space since it arises from $\mathbf{u} \cdot \nabla \mathbf{u}$. The mean profile $\bar{\mathbf{u}}$ must be known *a priori* to form $\mathcal{H}(\omega)$ and Equation 2.20 is presented in Table 2.1 as mean flow resolvent analysis. It should be noted that the sense of the imaginary and real parts of ω is reversed in the resolvent analysis relative to the definition customary to the stability literature. The real part of ω is the frequency associated with a mode while the imaginary part is set to zero as only neutral disturbances are considered.

$\mathcal{H}(\omega)$ can be decomposed via a singular value decomposition (SVD), e.g., McKeon and Sharma (2010):

$$\mathcal{H}(\omega) = \mathbf{\Psi}(\omega)\mathbf{\Sigma}(\omega)\mathbf{\Phi}^*(\omega), \quad (2.22)$$

where $\mathbf{\Psi}$ and $\mathbf{\Phi}$ are the left and right singular vectors corresponding to the response and forcing modes, often called resolvent modes (see McKeon and Sharma, 2010), respectively. Both sets of singular vectors are guaranteed to be orthonormal bases and are ranked according to their gain, or singular value, contained in the diagonal matrix $\mathbf{\Sigma}$. The resolvent operator can thus be written as the sum of outer products of the left and right singular vectors

$$\mathcal{H}(\omega) = \sum_{j=1}^{\infty} \hat{\boldsymbol{\psi}}_j(\omega) \sigma_j(\omega) \hat{\boldsymbol{\phi}}_j^*(\omega). \quad (2.23)$$

$\mathcal{H}(\omega)$ is (approximately) low rank if

$$\sum_{j=1}^p \sigma_j^2 \approx \sum_{j=1}^{\infty} \sigma_j^2, \quad (2.24)$$

where $\sigma_p \gg \sigma_{p+1}$ and p is small (Moarref et al., 2013; McKeon, 2017). If the leading singular value is significantly greater than all others ($\sigma_1 \gg \sigma_2$) then the rank-1 approximation can be invoked and the resolvent is approximated by the outer product of the leading optimal response and forcing modes:

$$\mathcal{H}(\omega) \approx \sigma_1 \hat{\boldsymbol{\psi}}_1 \hat{\boldsymbol{\phi}}_1^*. \quad (2.25)$$

The physical interpretation of the resolvent response modes is the response to forcing that results in a neutrally stable response, i.e., with the real component of frequency equal to zero. The singular value gives the input-output gain, here associated with the energy norm.

The resolvent may also be formed for the base flow (see Table 2.1)

$$\mathcal{H}(\omega) = \mathbf{C}^T (i\omega \mathbf{B} - \mathbf{A})^{-1} \mathbf{C}, \quad (2.26)$$

where the forcing is now treated as an extrinsic disturbance (see Jovanović and Bamieh, 2005). This is different from the mean flow case where the forcing is intrinsic and stems from the nonlinear term.

The singular values of the resolvent operator are computed in a manner outlined by Sipp and Marquet (2013); a brief summary of the procedure is presented here. The singular value problem is reformulated as the following eigenvalue problem:

$$\mathcal{H}(\omega)^* \mathcal{H}(\omega) \hat{\phi}_i = \sigma_i^2 \hat{\phi}_i, \quad (2.27)$$

where $\hat{\phi}_i$ is the i th right singular vector corresponding to the singular value σ_i of $\mathcal{H}(\omega)$. The largest eigenvalues of the Hermitian operator $\mathcal{H}(\omega)^* \mathcal{H}(\omega)$ are computed using the ARPACK library and the parallel MUMPS solver developed by Amestoy et al. (2001). The response modes are then computed from Equation 2.19.

2.4.3 Link to Pseudo-Spectrum

Analyzing the resolvent corresponds to considering the spectrum of the perturbed LNS operator:

$$\Lambda_\epsilon(\mathbf{A}) = \{z \in \mathbb{C} : z \in \Lambda(\mathbf{A} + \mathbf{E}) \text{ where } \|\mathbf{E}\| \leq \epsilon\}, \quad (2.28)$$

where Λ_ϵ is the pseudospectrum of \mathbf{A} under a perturbation magnitude $\epsilon > 0$ and $\|\cdot\|$ is the (operator) 2-norm (Trefethen et al., 1993; Reddy et al., 1993; Taira et al., 2017). An equivalent definition is given by

$$\Lambda_\epsilon(\mathbf{A}) = \left\{ z \in \mathbb{C} : \|(z\mathbf{I} - \mathbf{A})^{-1}\| \geq \epsilon^{-1} \right\} \cup \Lambda(\mathbf{A}), \quad (2.29)$$

where $\Lambda = \Lambda_0$ is the spectrum of \mathbf{A} . Throughout the thesis, Λ denotes the set of eigenvalues and $\mathbf{\Lambda}$ is the diagonal matrix of eigenvalues. If \mathbf{A} is normal, Λ_ϵ can be interpreted as the set of points away from Λ by only less than or equal to ϵ on the complex plane (see Taira et al., 2017). If \mathbf{A} is non-normal, this distance may be greater than ϵ , signifying that an eigenvalue is sensitive to perturbation of the LNS operator..

For a given z , the resolvent norm $\|(z\mathbf{I} - \mathbf{A})^{-1}\|$ is equal to the largest value of ϵ^{-1} such that z is contained within Λ_ϵ . The resolvent norm is (by definition) the maximum singular value of the resolvent operator, and quantifies the system's sensitivity

to temporal forcing. The neutrally stable response of the system to harmonic forcing is characterized by the value of the resolvent norm along the imaginary axis. For a stable, normal operator, the largest response occurs at the frequency corresponding to the imaginary part of the least stable eigenvalue λ_{ls} and the resolvent norm is $1/\text{Real}(\lambda_{ls})$ since $\text{Real}(\lambda_{ls})$ is the minimum distance between the eigenvalue and the imaginary axis (Chomaz, 2005). If the operator is marginally stable ($\text{Real}(\lambda_{ls}) \approx 0$), then $\epsilon \approx 0$ and the response is dominated by the corresponding eigenmode. The resolvent is not defined when the flow is unstable or $\text{Real}(\lambda_{ls}) = 0$. For a stable, non-normal operator, the frequency and gain of the largest response is less predictable since it is necessary to find the smallest value of ϵ for which the pseudospectrum crosses the imaginary axis.

2.4.4 Flows with Homogeneous Directions

In the case of turbulent channel flow, it is possible to Fourier-transform in the streamwise and spanwise directions resulting in the spatial wavenumbers k_x and k_z , respectively, since these directions are homogeneous. The incompressible NSE are non-dimensionalized by the channel half-height h and the friction velocity $u_\tau = \sqrt{\tau_w/\rho}$ (where τ_w is the wall shear stress, ρ is the density) resulting in

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + Re_\tau^{-1} \nabla^2 \mathbf{u} \quad (2.30a)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2.30b)$$

where $Re_\tau = hu_\tau/\nu$ and ν is the kinematic viscosity. The wall-normal domain extends from $y/h = -1$ to $y/h = 1$ with no-slip and no-penetration conditions imposed at the wall. The fluctuations are expressed as Fourier modes in the streamwise/spanwise directions and in time,

$$\hat{\mathbf{u}}(k_x, k_z, \omega; y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \mathbf{u}'(x, y, z, t) e^{-i(k_x x + k_z z - \omega t)} dx dz dt. \quad (2.31)$$

Upon elimination of the pressure term, the governing equations can be expressed in terms of the fluctuating vertical velocity \hat{v} and normal vorticity $\hat{\eta} = ik_z \hat{u} - ik_x \hat{w}$,

$$-i\omega \begin{pmatrix} \hat{v} \\ \hat{\eta} \end{pmatrix} + \begin{pmatrix} k^2 - \mathcal{D}^2 & 0 \\ 0 & 1 \end{pmatrix}^{-1} \begin{pmatrix} \mathcal{L}_{OS} & 0 \\ ik_z \bar{u}' & \mathcal{L}_{SQ} \end{pmatrix} \begin{pmatrix} \hat{v} \\ \hat{\eta} \end{pmatrix} = \mathbf{J} \hat{\mathbf{f}}, \quad (2.32)$$

where the Orr-Sommerfeld (OS) and Squire (SQ) operators are given by

$$\mathcal{L}_{OS} = ik_x \bar{u} (k^2 - \mathcal{D}^2) + ik_x \bar{u}'' + \frac{1}{Re_\tau} (k^2 - \mathcal{D}^2)^2, \quad (2.33)$$

$$\mathcal{L}_{SQ} = ik_x \bar{u} + \frac{1}{Re_\tau} (k^2 - \mathcal{D}^2), \quad (2.34)$$

and

$$\mathbf{J} = \begin{pmatrix} k^2 - \mathcal{D}^2 & 0 \\ 0 & 1 \end{pmatrix}^{-1} \begin{pmatrix} -ik_x \mathcal{D} & -k^2 & -ik_z \mathcal{D} \\ ik_z & 0 & -ik_x \end{pmatrix}, \quad (2.35)$$

$$\hat{\mathbf{f}} = \begin{pmatrix} \hat{f}_u \\ \hat{f}_v \\ \hat{f}_w \end{pmatrix} = -\langle \mathbf{u}' \cdot \nabla \mathbf{u}' \rangle_{\mathbf{k}}. \quad (2.36)$$

Here $\mathcal{D} = \frac{\partial}{\partial y}$, $k^2 = k_x^2 + k_z^2$, and $\langle \cdot \rangle_{\mathbf{k}}$ denotes the Fourier component associated with the wavenumber vector $\mathbf{k} = (k_x, k_z, \omega)$.

The wall-normal operators are discretized numerically with Chebyshev collocation points using the suite developed by Weideman and Reddy (2000). Equation 2.32 is recast into the following input/output form

$$\begin{pmatrix} \hat{u} \\ \hat{v} \\ \hat{w} \end{pmatrix} = \mathcal{H}(k_x, k_z, \omega) \begin{pmatrix} \hat{f}_u \\ \hat{f}_v \\ \hat{f}_w \end{pmatrix}, \quad (2.37)$$

where the resolvent operator \mathcal{H} is given by

$$\mathcal{H}(k_x, k_z, \omega) = \mathbf{K}(-i\omega + \mathbf{L})^{-1} \mathbf{J}, \quad (2.38)$$

and

$$\mathbf{L} = \mathbf{G}^{-1} \mathcal{L}, \quad (2.39)$$

$$\mathbf{G} = \begin{pmatrix} k^2 - \mathcal{D}^2 & 0 \\ 0 & 1 \end{pmatrix}, \quad (2.40)$$

$$\mathcal{L} = \begin{pmatrix} \mathcal{L}_{Os} & 0 \\ ik_z \bar{\mathbf{u}}' & \mathcal{L}_{SQ} \end{pmatrix}, \quad (2.41)$$

$$\mathbf{K} = \frac{1}{k^2} \begin{pmatrix} ik_x \mathcal{D} & -ik_z \\ k^2 & 0 \\ ik_z \mathcal{D} & ik_x \end{pmatrix}. \quad (2.42)$$

As before, the resolvent operator can be decomposed via the SVD as

$$\mathcal{H}(k_x, k_z, \omega) = \mathbf{\Psi}(k_x, k_z, \omega) \mathbf{\Sigma}(k_x, k_z, \omega) \mathbf{\Phi}^*(k_x, k_z, \omega). \quad (2.43)$$

The Orr-Sommerfeld/Squire formulation is used in Chapter 3 with respect to the discussion of amplification mechanisms in turbulent channel flow. The selection of the wavenumber vector \mathbf{k} is based on knowledge of highly amplified or known turbulent structures to highlight the role of non-normality in the resolvent operator.

2.4.5 Dynamic Mode Decomposition

In DMD, snapshots of data, which are collected at a fixed time interval Δt and contain m states, are arranged into the columns of matrices \mathbf{X} and $\mathbf{X}^\#$, such that

$$\mathbf{X} = [\mathbf{x}(t_1) \ \mathbf{x}(t_2) \ \cdots \ \mathbf{x}(t_{n-1})] \text{ and } \mathbf{X}^\# = [\mathbf{x}(t_2) \ \mathbf{x}(t_3) \ \cdots \ \mathbf{x}(t_n)], \quad (2.44)$$

where n is the number of snapshots. The objective is to approximate the linear operator \mathbf{T} which maps the snapshot matrix in time, e.g.

$$\mathbf{X}^\# = \mathbf{T} \mathbf{X}. \quad (2.45)$$

The DMD eigenvalues and modes are the eigenvalues and eigenvectors of \mathbf{T} (Tu et al., 2014). Equation 2.45 is an undetermined system of equations which is recast into an optimization problem to find \mathbf{T} . The solution can be expressed as

$$\mathbf{T} = \mathbf{X}^\# (\mathbf{X})^+, \quad (2.46)$$

where $(\cdot)^+$ denotes the pseudo-inverse. The calculation of the pseudo-inverse is achieved by using the singular value decomposition (SVD) of \mathbf{X} resulting in $\mathbf{X} = \mathbf{U} \mathbf{\Sigma} \mathbf{W}^*$ where $(\cdot)^*$ denotes the conjugate transpose. $\mathbf{\Sigma}$ is a matrix of size $m \times n$ containing the singular values of \mathbf{X} while \mathbf{U} and \mathbf{W} are unitary matrices ($\mathbf{U} \mathbf{U}^* = \mathbf{I}$) of size $m \times m$ and $n \times n$, respectively.

Equation 2.46 can be rewritten as

$$\mathbf{T} = \mathbf{X}^\# \mathbf{W} \mathbf{\Sigma}^{-1} \mathbf{U}^*. \quad (2.47)$$

Left and right-multiplication of Equation 2.47 by \mathbf{U}^* and \mathbf{U} , respectively, results in

$$\mathbf{U}^* \mathbf{T} \mathbf{U} = \mathbf{U}^* \mathbf{X}^\# \mathbf{W} \mathbf{\Sigma}^{-1}, \quad (2.48)$$

signifying that the columns of \mathbf{T} have been projected onto the singular vectors of \mathbf{X} . The left-hand side of Equation 2.48 is the companion matrix \mathbf{S} which is similar to \mathbf{T} . The eigenvalues of \mathbf{S} , consequently, approximate those of \mathbf{T} and the corresponding eigenvector (DMD mode) \mathbf{v}_i is given by $\mathbf{v}_i = \mathbf{U} \boldsymbol{\gamma}_i$, where $\boldsymbol{\gamma}_i$ is an eigenvector of \mathbf{S} . The DMD eigenvalues are computed from $\lambda_i = \log(\mu_i) / \Delta t$, where μ_i is an eigenvalue of \mathbf{S} .

2.5 Data-Assimilation of the Mean Flow

The difficulties associated with obtaining the mean velocity profile from experimental data (e.g. limited field of view, noise, etc.) can be overcome by data-assimilation. The primary elements of the data-assimilation framework are similar

to those outlined in Foures et al. (2014), and so only an abridged version is presented here. Modifications to the procedure for experimental data at higher Reynolds numbers developed in this work are discussed in Section 2.6.

2.5.1 Computational Model

The measured mean flow is obtained from 2D PIV data or numerical simulation and is assumed to satisfy the 2D incompressible RANS equations given by

$$\bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} + \nabla \bar{p} - Re^{-1} \nabla^2 \bar{\mathbf{u}} = \mathbf{f} \quad (2.49a)$$

$$\nabla \cdot \bar{\mathbf{u}} = 0. \quad (2.49b)$$

The forcing term \mathbf{f} on the right-hand side of Equation 2.49a is treated as an unknown momentum forcing which is data-driven whereas for ordinary RANS simulations it is solved for using a turbulence model. Out-of-plane velocities are not captured by the experiment, which means enforcing Equation 2.49b may not strictly be true at every point in the domain. Non-zero three-dimensional mean flow effects are compensated for by the momentum forcing term.

When the algorithm is formulated on velocity-only measurements, only a partial recovery of the pressure is possible by data-assimilation of the mean velocity field. Upon further inspection of Equation 2.49a, it is noted that the forcing term \mathbf{f} can be decomposed into the following

$$\mathbf{f} = \nabla \xi + \mathbf{f}_s, \quad (2.50)$$

where $\nabla \xi$ and \mathbf{f}_s represent the irrotational part and solenoidal part of \mathbf{f} , respectively (Foures et al., 2014). As will be shown below, the recovered forcing is divergence-free without pressure measurements; therefore, only the solenoidal part can be captured. Boundary conditions need to be specified for this decomposition to be unique. This involves enforcing $\nabla \xi \cdot \mathbf{n} = 0$ and $\mathbf{f}_s \cdot \mathbf{n} = 0$ on the airfoil walls where \mathbf{n} is the outward normal (Foures et al., 2014). The model given by Equation 2.49 can now be recast as follows

$$\bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} + \nabla \bar{p}' - Re^{-1} \nabla^2 \bar{\mathbf{u}} = \mathbf{f}_s \quad (2.51a)$$

$$\nabla \cdot \bar{\mathbf{u}} = 0, \quad (2.51b)$$

where $\bar{p}' = \bar{p} - \xi$. Despite the limitation that the pressure field cannot be reconstructed in its entirety, it is still possible to reconstruct the mean velocity field using

experimental data to correctly model f_s . The portion of f which is not captured is lumped in with the pressure term. The full f can be captured by incorporating pressure measurements into the data-assimilation algorithm which will be demonstrated in due course.

2.5.2 Fitting Criterion

The goal of the data-assimilation algorithm is to determine the solenoidal forcing f_s such that the estimate $\bar{\mathbf{u}}$ matches the available mean velocity field measurements. In order to determine this forcing, it is first necessary to calculate the discrepancy between the measured flow $\bar{\mathbf{u}}_{exp}$ and the estimate $\bar{\mathbf{u}}$. The discrepancy velocity field is computed using

$$\Delta\bar{\mathbf{u}} = \bar{\mathbf{u}}_{exp} - \bar{\mathbf{u}}, \quad (2.52)$$

where $\Delta\bar{\mathbf{u}}$ is a two-dimensional vector containing both streamwise and transverse discrepancy velocity measurements for each point in the domain. This computation is not trivial since the experimental and numerical data are frequently defined on different meshes. A detailed explanation for how to compute this field is given in the next section. For now, it is assumed that the discrepancy field has the same spatial resolution as the numerical field and that it is possible to perform the above operation using techniques which do not involve interpolation.

The L_2 -norm of the discrepancy field yields a scalar function called the fitting criterion, which quantifies the distance between the current estimate and the measured mean velocity. It is calculated using the following operation:

$$\mathcal{E}_A(\bar{\mathbf{u}}) = \frac{1}{2} \sum_{j=1}^N \|\Delta\bar{\mathbf{u}}_j\|^2, \quad (2.53)$$

where N denotes the number of points on the mesh and the index j denotes the j th point on the mesh. The subscript A signifies that only mean velocity measurements are available (i.e. no mean pressure measurements are available).

2.5.3 Cost Functional and Adjoint Equations

Similar to 3D-Var (see Lewis et al. (2006)), an objective functional using a variational formulation yields an iterative optimization scheme that minimizes the fitting criterion. The final result is an optimal match between the data-assimilated mean $\bar{\mathbf{u}}$ and the measured data $\bar{\mathbf{u}}_{exp}$. An augmented Lagrangian is formed consisting of the

objective functional (the fitting criterion) and the constraints (the model equations) which are enforced in a weak form by Lagrange multipliers or adjoint variables. For readers familiar with optimal control, the augmented Lagrangian can also be thought of as a least-squares cost functional which penalizes the difference between the computed state, which in this case is $\bar{\mathbf{u}}$, and the observed state, $\bar{\mathbf{u}}_{exp}$. The cost functional is written as follows:

$$\mathcal{L}(f_s, \bar{\mathbf{u}}, \bar{\mathbf{u}}^\dagger, \bar{p}', \bar{p}^\dagger) = \mathcal{E}(\bar{\mathbf{u}}) - \left\langle \bar{\mathbf{u}}^\dagger, \bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} + \nabla \bar{p}' - \frac{1}{Re} \nabla^2 \bar{\mathbf{u}} - f_s \right\rangle - \left\langle \bar{p}^\dagger, \nabla \cdot \bar{\mathbf{u}} \right\rangle. \quad (2.54)$$

The scalar product denoted by $\langle \cdot, \cdot \rangle$ is associated with the Euclidian norm for vector and scalar fields on the domain Ω .

In order to minimize the functional, first-order variational derivatives are taken with respect to each independent variable and set equal to zero. Differentiation with respect to the adjoint variables $\bar{\mathbf{u}}^\dagger$ and \bar{p}^\dagger yields the direct equations given by Equation 2.51 while differentiation with respect to the direct variables produces the adjoint equations given below:

$$-\bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}}^\dagger + \bar{\mathbf{u}}^\dagger \cdot \nabla \bar{\mathbf{u}}^T - \nabla \bar{p}^\dagger - \frac{1}{Re} \nabla^2 \bar{\mathbf{u}}^\dagger = \frac{\delta \mathcal{E}}{\delta \bar{\mathbf{u}}} \quad (2.55a)$$

$$\nabla \cdot \bar{\mathbf{u}}^\dagger = 0. \quad (2.55b)$$

The derivations of these equations and the boundary conditions are detailed in Appendix A. The right-hand term of Equation 2.55a is discussed in Section 2.6 as it pertains to the smoothing procedure needed to recover adjoint solutions on the higher-dimensional subspace of the simulation.

Finally, the variational derivative with respect to the forcing f_s yields

$$\nabla_{f_s} \mathcal{E} = \bar{\mathbf{u}}^\dagger, \quad (2.56)$$

which is the steepest descent direction towards the optimality condition. It can also be interpreted as the sensitivity of the fitting criterion to the forcing vector. An initial guess is necessary to begin the optimization procedure and compute the first descent direction. Since the forcing must be divergence-free and the Reynolds stresses must vanish on the airfoil surface, a natural initial guess is $f_s = 0$. The solution to Equation 2.55 then provides $\bar{\mathbf{u}}^\dagger$, which is the direction in which the guess to the forcing is updated. This can be stated mathematically by the following equation

$$f_{n+1} = f_n + \beta_n \bar{\mathbf{u}}_n^\dagger, \quad (2.57)$$

where $\mathbf{f}_{n=0}$ denotes the initial guess. This method is known as the simple gradient descent method which, while effective, has shown to be prone to zigzagging since the new search direction is always orthogonal to the previous search direction (see Bewley et al., 2001). To avoid this problem, a conjugate-gradient approach using the Polak-Ribière formula (Polak and Ribiere, 1969) is adopted as well as a line-search algorithm to determine β_n for each new descent direction.

As seen from Equation 2.57, the forcing is a linear combination of adjoint velocity fields which, from Equation 2.55b, must be divergence-free. This reinforces the validity of the initial guess and justifies the earlier statement that only the solenoidal component of the forcing can be captured by the data-assimilation algorithm.

2.5.4 Mean Pressure Reconstruction

Up until this point, Section 2.5 has recapped the algorithm first developed in Foures et al. (2014). The extension proposed here is a novel contribution to the framework although it was originally suggested as a topic of future work in the original paper. To capture the full forcing and reconstruct the mean pressure, \mathcal{E} is modified in Equations 2.53 and 2.54 to include mean pressure measurements such that $\mathcal{E} = \mathcal{E}_B(\bar{\mathbf{u}}, \bar{p})$, where the subscript B represents the consideration of mean velocity and pressure measurements. This modifies the variational derivative of Equation 2.54 with respect to \bar{p} and so Equation 2.55b becomes

$$\nabla \cdot \bar{\mathbf{u}}^\dagger = \frac{\delta \mathcal{E}}{\delta \bar{p}}, \quad (2.58)$$

implying that the adjoint velocity is no longer constrained to be divergence-free as it was in Foures et al. (2014). The modified cost function is

$$\mathcal{E}_B(\bar{\mathbf{u}}, \bar{p}) = \frac{1}{2} \sum_{j=1}^N \|\Delta \bar{\mathbf{u}}_j + \Delta \bar{p}_j\|^2, \quad (2.59)$$

and

$$\Delta \bar{p} = \bar{p}_{exp} - \bar{p}. \quad (2.60)$$

Even if no mean pressure measurements are available, the resolvent response modes can be utilized to correct the mean pressure. The forcing to the mean momentum equation can be rewritten as a sum of resolvent modes:

$$\bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} + \nabla \bar{p} + Re^{-1} \nabla^2 \bar{\mathbf{u}} = \sum_{\omega} \sum_{j=1}^N 2Real(\chi_j^2 \hat{\boldsymbol{\psi}}_{\omega_j} \cdot \nabla \hat{\boldsymbol{\psi}}_{-\omega_j}), \quad (2.61)$$

where χ_j is the amplitude of the j th resolvent response mode (see Section 2.5.5). The left-hand side of (2.61) is a sum of triadic interactions where a mode with frequency ω interacts with its conjugate at frequency $-\omega$ to produce the Reynolds stresses required to sustain the mean profile. The right-hand side of (2.61) consists of both a rotational and irrotational component unlike the velocity-only data-assimilated forcing, which is divergence-free.

Subtracting Equation 2.51a from Equation 2.61 yields an equation for ξ :

$$\nabla^2 \xi = \sum_{j=1}^N 2A^2 \text{Real}(\hat{\psi}_{\omega_j} \cdot \nabla \hat{\psi}_{-\omega_j}) - \mathbf{f}_s. \quad (2.62)$$

Taking the divergence of Equation 2.62, one eliminates \mathbf{f}_s and obtains a Poisson equation for ξ :

$$\nabla^2 \xi = \nabla \cdot \left[\sum_{j=1}^N 2A^2 \text{Real}(\hat{\psi}_{\omega_j} \cdot \nabla \hat{\psi}_{-\omega_j}) \right], \quad (2.63)$$

where, as mentioned earlier, it is necessary to set $\partial \xi / \partial n = 0$ at the surface of the body to define a unique solution.

2.5.5 Amplitude Calibration

The weight of a response mode is determined by computing the projection of the nonlinear forcing onto the resolvent forcing mode:

$$\hat{\mathbf{u}} = \sum_j \hat{\psi}_j(\mathbf{x}, \omega) \sigma_j(\omega) \langle \hat{\mathbf{f}}(\mathbf{x}, \omega), \hat{\phi}_j(\mathbf{x}, \omega) \rangle = \sum_j \hat{\psi}_j(\mathbf{x}, \omega) \sigma_j(\omega) \bar{\chi}_j(\omega), \quad (2.64)$$

where $\bar{\chi}_j$ is the complex weight associated with the j th response mode (McKeon et al., 2013). If $\sigma_1 \gg \sigma_2$, the weights of the first resolvent response modes ($j = 1$) can be approximated using the method proposed by Gómez et al. (2016a) and Beneddine et al. (2016). To begin, it is necessary to have time-resolved knowledge of the fluctuations at a single point, \mathbf{x}_0 . The location and flow quantity (e.g. streamwise velocity, pressure, etc.) is likely to be flow specific and should include regions in which the resolvent response modes are energetic (Gómez et al., 2016a; Beneddine et al., 2016). For now, it is assumed the streamwise velocity is measured and the signal is Fourier-transformed in time yielding $\check{u}^\omega(\mathbf{x}_0)$. This complex-valued function of ω predicts the amplitude of the corresponding resolvent mode through the following expression

$$\chi_1 = \hat{u}(\omega, \mathbf{x}_0) / \check{u}^\omega(\mathbf{x}_0), \quad (2.65)$$

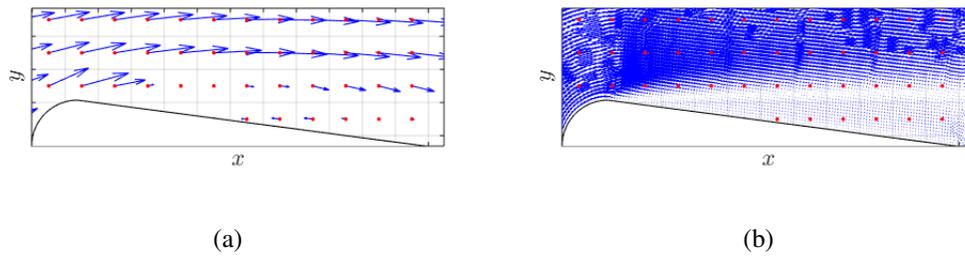


Figure 2.5: Subdivision of the domain into cells delineated by grey lines. The center of each cell is indicated by a red point which corresponds to the spatial location of a PIV measurement shown by a blue vector. Only a limited section of the total field of view is shown. In (a) one out of every five PIV vectors in both the streamwise and transverse directions appear for clarity while (b) displays the density of vectors in the CFD mesh. Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

where $\chi_1(\omega) = \sigma_1(\omega)\overline{\chi}_1(\omega)$ and $\hat{u}(\omega, \mathbf{x}_0)$ is the value of the unscaled resolvent response mode at \mathbf{x}_0 . The reader is referred to e.g. Towne et al. (2018) for other ways to obtain the amplitudes from data.

2.6 Smoothing Procedure

It was assumed in the previous section that the resolutions of the experimental and numerical velocity fields are identical when computing the discrepancy velocity field. This is generally not the case since the experimental velocity field is defined on a coarse, uniform Cartesian grid while the numerical velocity field is continuous and approximated with quadratic basis functions. To obtain a discrete representation of the numerical velocity field, the values of the velocity defined on the vertices and midpoints of the finite element triangles can be outputted from the finite element code. A smoothing procedure described in Symon et al. (2017) is used to cope with the difference in resolution when forcing the adjoint RANS equations. For convenience, the points where experimental measurements exist are referred to as grid points while their numerical counterparts are referred to as mesh points.

The smoothing procedure begins by dividing the domain into cells as seen in Figure 2.5 where the red dot denotes the center of each cell and the blue vector the associated PIV measurement. Mesh points are sorted into the cells whose boundaries are delineated by the grey lines in Figure 2.5. This means that for a given cell, there is a unique experimental measurement and multiple numerical measurements. The

fitting criterion in Equation 2.53 can now be recast in terms of a new scalar product defined in the experimental measurement space:

$$\mathcal{E}(\bar{\mathbf{u}}) = \frac{1}{2} \langle \Delta\bar{\mathbf{u}}, \Delta\bar{\mathbf{u}} \rangle_M, \quad (2.66)$$

where the subscript M represents the number of measures in the inner product defined by

$$\langle \Delta\bar{\mathbf{u}}_1, \Delta\bar{\mathbf{u}}_2 \rangle_M = \sum_{j=1}^M \Delta\bar{\mathbf{u}}_{1,j} \Delta\bar{\mathbf{u}}_{2,j}. \quad (2.67)$$

The numerical measurements are projected down to the measurement space using a smoothing function or operator \mathcal{P} :

$$\Delta\bar{\mathbf{u}} = \bar{\mathbf{u}}_{exp} - \mathcal{P}\bar{\mathbf{u}}. \quad (2.68)$$

For this configuration, \mathcal{P} is chosen to be an area average of the numerical measurements in a given cell. For the j th cell, therefore, \mathcal{P} can be written as

$$\mathcal{P}_j \bar{\mathbf{u}} = \int_{\Omega} \bar{\mathbf{u}}(x) \varphi_j(x) d\Omega, \quad (2.69)$$

where

$$\varphi_j(x) = \begin{cases} \frac{1}{\int_{\Omega_j} d\Omega} & \text{in } \Omega_j \\ 0 & \text{outside of } \Omega_j, \end{cases} \quad (2.70)$$

and $\int_{\Omega_j} d\Omega$ is the area of a cell.

The forcing term of Equation 2.55a can be derived by substituting Equation 2.68 into Equation 2.66 and taking the variational derivative with respect to $\bar{\mathbf{u}}$. After simplifying the expansion, one obtains

$$\left\langle \frac{\delta \mathcal{E}}{\delta \bar{\mathbf{u}}}, \delta \bar{\mathbf{u}} \right\rangle = \langle \bar{\mathbf{u}}_{exp} - \mathcal{P}\bar{\mathbf{u}}, -\mathcal{P}\delta \bar{\mathbf{u}} \rangle_M = \langle \mathcal{P}^\dagger (\mathcal{P}\bar{\mathbf{u}} - \bar{\mathbf{u}}_{exp}), \delta \bar{\mathbf{u}} \rangle, \quad (2.71)$$

and so $\frac{\delta \mathcal{E}}{\delta \bar{\mathbf{u}}}$, therefore, reads

$$\frac{\delta \mathcal{E}}{\delta \bar{\mathbf{u}}} = \mathcal{P}^\dagger (\mathcal{P}\bar{\mathbf{u}} - \bar{\mathbf{u}}_{exp}). \quad (2.72)$$

The adjoint operator of \mathcal{P} satisfies

$$\langle \mathcal{P}^\dagger \Delta\bar{\mathbf{u}}, \delta \bar{\mathbf{u}} \rangle = \langle \Delta\bar{\mathbf{u}}, \mathcal{P}\delta \bar{\mathbf{u}} \rangle_M \text{ for all } \Delta\bar{\mathbf{u}} \text{ and } \delta \bar{\mathbf{u}}. \quad (2.73)$$

For the \mathcal{P} chosen in this study, the adjoint operator acting on $\Delta\bar{\mathbf{u}}$ can be shown to be

$$\mathcal{P}^\dagger \Delta\bar{\mathbf{u}} = \sum_{j=1}^M \Delta\bar{\mathbf{u}}_j \varphi_j(x), \quad (2.74)$$

and so the final expression for the forcing term of Equation 2.55a reduces to the following expression:

$$\frac{\delta \mathcal{E}}{\delta \bar{\mathbf{u}}} = \sum_{j=1}^M (\mathcal{P}_j \bar{\mathbf{u}} - \bar{\mathbf{u}}_{j,exp}) \varphi_j(x). \quad (2.75)$$

This calculation need not include every PIV vector in the experimental data set. It is possible to artificially lower the resolution by omitting one out of two vectors in both the x - and y -directions, for example, which would increase the area of each cell by a factor of four. It is also possible to change the weight of a cell to increase its relative importance or turn it off completely to artificially decrease the PIV window. A second projection operator \mathcal{S} is introduced into Equation 2.68 to account for weighting or artificially reducing the number of experimental measurements:

$$\Delta \bar{\mathbf{u}} = \mathcal{S}(\bar{\mathbf{u}}_{exp} - \mathcal{P} \bar{\mathbf{u}}). \quad (2.76)$$

If \mathcal{S} is the identity, then all experimental data are used to compute the fitting criterion and drive the data-assimilation process. This is referred to as the full-field case in Chapter 6. The modifications to the previous derivations for $\frac{\delta \mathcal{E}}{\delta \bar{\mathbf{u}}}$ are minor even though it will change the measurement space in Equation 2.67 as the effect of \mathcal{S} can be incorporated into $\varphi_j(x)$ of the original smoothing function \mathcal{P} .

The smoothing procedure can be interpreted as a way to compute the pseudo-inverse of the projection operator \mathcal{P} . For this setup, the numerical data must be projected down onto the experimental data subspace by \mathcal{P} to compute the discrepancy field but the adjoint solution needs to be forced by a quantity defined on the same subspace as the original numerical data. The method outlined in the previous paragraph is one way to perform this pseudo-inverse operation. Alternatives, which are not implemented in this thesis, include enforcing global smoothness constraints or including regularization terms in Equation 2.54. Failure to include the smoothing procedure results in forcing the adjoint equations by Dirac delta functions located at the center of each cell and the resulting forcing fields are no longer smooth. The algorithm, furthermore, will only attempt to match the velocities at discrete points where the PIV vectors are located, and this can lead to spurious recirculation bubbles or other unphysical flow structures appearing in the assimilated flow field between measurement points.

2.7 Adjoint Looping: Implementation of Data-Assimilation Algorithm

This section describes the details for implementing data-assimilation for numerical and experimental data. The block diagram in Figure 2.6 illustrates the adjoint loop-

ing procedure used to iteratively update the forcing f_s . The process is terminated once the new steepest descent direction no longer results in a reduction of the fitting criterion. The data-assimilation begins with an initial guess $f_0 = 0$, which is used to solve the RANS equations given by Equation 2.51. Next, the discrepancy field is computed using Equation 2.76 once the relevant projection operators have been applied to \bar{u} and \bar{u}_{exp} , and these values are smoothed over all mesh points in a grid cell as explained earlier. Finally, the adjoint equations are solved to find the steepest descent direction that updates the forcing.

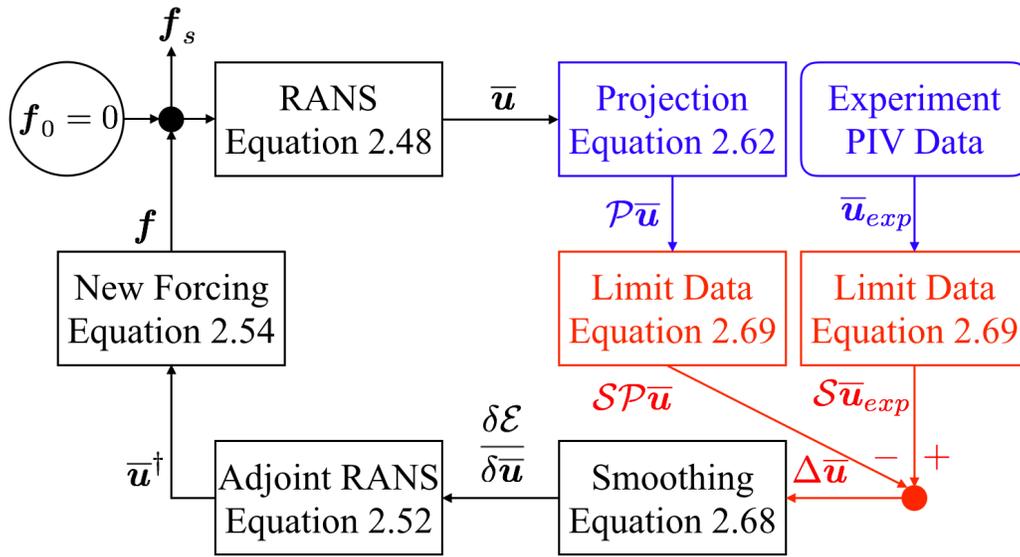


Figure 2.6: A block diagram representation of the adjoint looping procedure used to determine the optimal forcing f_s which will minimize the fitting criterion. The steps and quantities in black are in the subspace of the numerical data, while blue indicates the subspace of the experiment and brown represents the subspace of fewer measurements after applying the projection operator \mathcal{S} .

If $f_0 = 0$ is the initial guess, then a problem that arises in this procedure is finding a base flow solution at the experimental Reynolds number. This is due to the fact that the critical Reynolds number, or the Reynolds number at which the base flow undergoes a supercritical Hopf bifurcation and becomes unstable, tends to be rather low. The base flows at higher Reynolds numbers, as a result, tend to be unphysical and require exceedingly high numerical precision to compute (Sipp et al., 2010). A base flow is computed at a lower Reynolds number and the algorithm is run for this lower number until the fitting criterion is minimized. The Reynolds number is then increased incrementally until it reaches the experimental Reynolds number. It is

not a requirement to use a base flow solution as the starting point to the algorithm. One could use the results of a RANS simulation at the same Reynolds number, for example.

2.8 Numerical Details of Experimental Data-Assimilation

The idealized airfoil is mounted at an angle of attack of zero degrees. Since its profile is symmetric, the time-averaged flow on either side of the airfoil is identical so data are collected for one side only and reflected over the centerline. Both the numerical and experimental data are non-dimensionalized to yield a unit input velocity and a unit cylindrical diameter of the idealized airfoil. The direct and adjoint equations are solved using FreeFem++ on a computational domain Ω which spans $-20 \leq x \leq 35$, $-12.5 \leq y \leq 12.5$ with the cylinder at the leading edge of the airfoil being centered at the origin as it is in the PIV domain shown in Figure 2.4. The mesh density, which is controlled by specifying the number of divisions along a boundary, increases from the outer boundaries towards the airfoil surface. The equations are spatially discretized using quadratic basis functions for the velocity and linear basis functions for the pressure, resulting in approximately 580,000 degrees of freedom for both velocity and pressure. FreeFem++ solves the RANS and adjoint RANS equations in weak form which can be obtained by dotting the equations with a test function and integrating by parts to remove second-derivative terms. As mentioned earlier, the direct equations are solved using a Newton-Raphson method while the adjoint equations are linear in $\bar{\mathbf{u}}^\dagger$ and can be solved efficiently without an iterative method.

For the NACA 0018 airfoil at an angle of attack, the laser sheet is obstructed by the airfoil's shadow and so data are not available for the pressure side of the airfoil. Since it is necessary to obtain the mean profile on both sides of the airfoil, data are also collected for the airfoil mounted at -10° . The time-averaged profiles for both the positive and negative angle of attack are later stitched together and are used as an input for the data-assimilation algorithm. The numerical and experimental data are non-dimensionalized to yield a unit input velocity and a unit chord length. The computational domain is $\Omega \in -6.5 \leq x \leq 15.5 \cup -5 \leq y \leq 5$ with the leading edge of the airfoil being centered at the origin. The equations are spatially discretized using quadratic basis functions for the velocity and linear basis functions for the pressure resulting in approximately 115,000 degrees of freedom for the 0° case and 331,000 degrees of freedom for the 10° case.

2.9 Evaluation of Agreement Between Experiment and Data-Assimilation

To quantify how well the data-assimilation recovers the mean flow, the L_2 -difference is computed between the PIV velocity field $\bar{\mathbf{u}}_{exp}$ and the assimilated flow $\bar{\mathbf{u}}$ interpolated onto the PIV grid, which will be referred to as $\bar{\mathbf{u}}_{int}$. This difference is then normalized by the L_2 -norm of the PIV velocity field and is referred to as the experimental mismatch. It can be written as

$$Q = \frac{\|\bar{\mathbf{u}}_{exp} - \bar{\mathbf{u}}_{int}\|}{\|\bar{\mathbf{u}}_{exp}\|}. \quad (2.77)$$

Q is computed on the domain $x \in [-1.5, 6] \cup y \in [0, 2]$. This definition is constant across all cases studied even if the experimental data are artificially limited.

Another quantity of interest is the decrease of the cost function for each case investigated. A bad cost function is one which leads to a large decrease in the fitting criterion without reducing the experimental mismatch. To determine the success of a given cost function, its final value \mathcal{E}_{end} is normalized by the initial value computed using the base flow \mathcal{E}_0 . This ratio can be compared to the corresponding decrease of the experimental mismatch

$$Q_r = \frac{Q_{end}}{Q_0}, \quad \mathcal{E}_r = \sqrt{\frac{\mathcal{E}_{end}}{\mathcal{E}_0}}, \quad (2.78)$$

where Q_0 is computed using the base flow. The square root of the cost function ratio is necessary since \mathcal{E}_r is based on $\|\cdot\|^2$ while Q_r is based on $\|\cdot\|$. If $\mathcal{E}_r \ll Q_r$, the cost function is not effectively reducing the mean velocity field discrepancy.

2.10 Summary

The flow configurations are summarized in Table 2.2, which includes the Reynolds number ranges, the method used to obtain the data, and the chapters where they are analyzed.

	Parameters	Methods	Chapters
2D Cylinder	$47 < Re < 190$	DNS	3, 4, 5
3D Cylinder	$190 \leq Re \leq 320$	DNS	3, 5
Channel Flow	$Re_\tau = 2000$	Eddy Viscosity	3
Idealized Airfoil	$Re_{chord} = 13,500$	PIV	6
NACA 0018 Airfoil	$Re_{chord} = 10,250$ $Re_{chord} = 20,700$ $-10^\circ \leq \alpha \leq 10^\circ$	PIV	6, 7

Table 2.2: The various flows considered including the parameters, method used to obtain the data, and the chapters where they are discussed.

Chapter 3

AMPLIFICATION MECHANISMS IN STABILITY AND RESOLVENT ANALYSIS

Symon, S., Rosenberg, K., Dawson, S. T. M., and McKeon, B. J. (2018). “Non-normality and classification of amplification mechanisms in stability and resolvent analysis”. In: *Phys. Rev. Fluids* 3.053902.

The focus of this chapter is to understand the types of amplification mechanisms identified by the resolvent and predict when the operator exhibits low-rank behavior. A lot of attention is devoted to the interplay of terms in the LNS operator which introduce non-normality and whose competing influence affects the mode shapes and size of amplification. They also influence the pseudo-spectrum of the LNS operator which characterizes the sensitivity of eigenvalues to perturbations of the operator. The beginning of this chapter considers 2-by-2 matrix models to illustrate the roles of various terms on the mode shapes and the pseudo-spectrum. Their roles are then examined in the context of oscillator flows for cylinder flow at $Re = 100$ and noise-amplifier flows for turbulent channel flow at $Re_\tau = 2000$. Features of the modes can be predicted by properties of these mean profiles.

3.1 Resonance and Pseudoresonance

The origin of the amplification mechanisms characterized by the resolvent norm can be identified by expanding the resolvent through an eigenvalue decomposition of the LNS operator,

$$\mathcal{H}(\omega) = \mathbf{C}^T (i\omega \mathbf{B} - \mathbf{V} \mathbf{\Lambda} \mathbf{B} \mathbf{V}^{-1})^{-1} \mathbf{C}. \quad (3.1)$$

Here \mathbf{V} represents the matrix of eigenvectors of the LNS operator for either the base or mean flow profile and $\mathbf{\Lambda}$ the diagonal matrix of eigenvalues. These can be used to find an upper and lower bound for the resolvent norm (see Schmid and Henningson, 2001):

$$\|i\omega \mathbf{I} - \mathbf{\Lambda}\|^{-1} \leq \|\mathcal{H}(\omega)\| \leq \underbrace{\|\mathbf{V}\| \|\mathbf{V}^{-1}\|}_{\text{pseudoresonance}} \underbrace{\|i\omega \mathbf{I} - \mathbf{\Lambda}\|^{-1}}_{\text{resonance}}. \quad (3.2)$$

Considering first the far-righthand term in Equation 3.2, it is clear that forcing in the vicinity of an eigenvalue, i.e. $\omega = \lambda$, is likely to lead to amplification due to

resonance and this is predictable from an eigenanalysis. Large amplification also arises in the event of pseudoresonance when the condition number $\kappa = \|\mathbf{V}\|\|\mathbf{V}^{-1}\|$ is large due to non-orthogonality of the eigenvectors, a consequence of the non-normal nature of \mathbf{L} and hence $\mathcal{H}(\omega)$. In the formulation of Equation 3.2, the resolvent therefore contains both the amplification mechanisms associated with the eigenvalue spectrum accessed via eigenanalysis (normal mode linear stability analysis) and the pseudoresonant amplification that is possible when the eigenvectors are not orthogonal to each other.

3.2 Amplification Mechanisms Highlighted by a Model Operator

A model operator \mathbf{M} inspired by Gebhardt and Grossmann (1994) is proposed to represent the LNS for a two-dimensional velocity field associated with a one-dimensional base or mean flow

$$\mathbf{M} = \begin{pmatrix} m_1 & d \\ 0 & m_2 \end{pmatrix}. \quad (3.3)$$

By selecting a one-dimensional operator, the spatial (streamwise and spanwise) dependence of the base/mean flow and therefore the modes is neglected. Nevertheless, the impact of the various types of term in the LNS operator on the resolvent modes can be modeled. Here $Real(m_j) < 0$ is analogous to the stabilizing role of viscosity through the $Re^{-1}\nabla^2(\cdot)$ term and $Imag(m_j)$ represents mean flow advection through the $-\bar{\mathbf{u}} \cdot \nabla(\cdot)$ term. d is real and is analogous to mean shear $(\cdot) \cdot \nabla\bar{\mathbf{u}}$, which here is equal to the gradient in the 2-direction of mean flow in the 1-direction. Thus in this simple 2-by-2 example, d models the lift-up mechanism (Landahl, 1980) by coupling forcing in the n_2 -direction (second component of the vector) with a response in the n_1 -direction (first component of the vector).

3.2.1 Eigenvectors, singular functions, and non-normality of the operator

The resolvent of \mathbf{M} is

$$\mathcal{H}(\omega) = \begin{pmatrix} -1/(m_1 - i\omega) & d/[(m_1 - i\omega)(m_2 - i\omega)] \\ 0 & -1/(m_2 - i\omega) \end{pmatrix}. \quad (3.4)$$

In order to isolate the effect of non-normality introduced via the off-diagonal term in Equation 3.5, the eigenvalues are assumed to be real and the immediate development is limited to stationary disturbances ($\omega = 0$) to eliminate the remaining imaginary terms, such that

$$\mathcal{H}(\omega = 0) = \begin{pmatrix} -1/(m_1) & d/[(m_1 m_2)] \\ 0 & -1/(m_2) \end{pmatrix}. \quad (3.5)$$

For the limiting case of $d = 0$, i.e. zero mean shear, \mathbf{M} and its resolvent are self-adjoint and therefore normal. If the least stable eigenvalue has real part close to zero which would occur, say, if $m_1 \rightarrow 0$, the singular value decomposition can be simplified to

$$\lim_{m_1 \rightarrow 0} \text{SVD}(\mathcal{H}(\omega = 0)) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad (3.6)$$

where $\sigma_1/\sigma_2 \rightarrow \infty$, i.e. the resolvent is low-rank. Thus the response can be well predicted from the leading singular vectors $\hat{\psi}_1$ and $\hat{\phi}_1$. For a normal operator, these are identical to each other, $\hat{\psi}_1 = \hat{\phi}_1 = [1 \ 0]^T$, and identical to the corresponding eigenvectors. The inner product $|\hat{\phi}_1^* \hat{\psi}_1|$ quantifies the componentwise correspondence between the forcing and response modes which, in this limit, is equal to unity. A schematic of the variation in the n_1 and n_2 directions of the forcing and response mode, i.e. the singular vectors, in this case is shown in Figure 3.1(a-b).

In the limit $d \rightarrow \infty$

$$\lim_{d \rightarrow \infty} \text{SVD}(\mathcal{H}(\omega = 0)) = \begin{pmatrix} 1 - \epsilon_2 & -\epsilon_1 \\ \epsilon_1 & 1 - \epsilon_2 \end{pmatrix} \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} \begin{pmatrix} \epsilon_1 & \epsilon_2 - 1 \\ 1 - \epsilon_2 & \epsilon_1 \end{pmatrix}, \quad (3.7)$$

where $\sigma_1/\sigma_2 \rightarrow \infty$ and $\epsilon_1, \epsilon_2 \rightarrow 0$. The constants ϵ_1 and ϵ_2 are real and positive. The resolvent operator is still low-rank in this limit, but $\hat{\psi}_1$ and $\hat{\phi}_1$ are now orthogonal to each other and thus the inner product $|\hat{\phi}_1^* \hat{\psi}_1| \rightarrow 0$. The perturbation energy in the optimal forcing mode is concentrated in the second component of the vector while the perturbation energy in the optimal response mode is concentrated in the first component, as sketched in Figure 3.1(c-d).

The analogous eigenvalue decomposition of (non-normal) \mathbf{M} is

$$\lim_{d \rightarrow \infty} \text{EIG}(\mathbf{M}) = \begin{pmatrix} 1 & 1 \\ 0 & \epsilon_3 \end{pmatrix} \begin{pmatrix} m_1 & 0 \\ 0 & m_2 \end{pmatrix} \begin{pmatrix} 1 & -1/\epsilon_3 \\ 0 & 1/\epsilon_3 \end{pmatrix}, \quad (3.8)$$

where $\epsilon_3 \rightarrow 0$ is a positive, real constant. Unlike the resolvent response modes which are orthogonal to one another, the eigenvectors are non-orthogonal, such that $\kappa \rightarrow \infty$. Thus, in this case the stability and resolvent modes are different. Since the eigenvectors are nearly parallel, they both project equally well onto the optimal resolvent response mode. The same can be said for the projection of the adjoint modes, which are also nearly parallel, onto the optimal resolvent forcing mode. The optimal response and forcing modes, therefore, are linear combinations

of multiple stability and adjoint modes, respectively. Furthermore, it is not clear from this decomposition that the resolvent operator is low-rank.

As outlined in Figure 3.1(a-d), the model operator given in Equation 3.3 can be used to elucidate that the presence of mean shear, which introduces non-normality into the LNS operator, is to concentrate energy in different velocity components of the resolvent forcing and response modes. The model LNS operator does not need a base or mean flow with any special characteristics or spatial dependence other than an off-diagonal term to reveal this. Nevertheless, the same mechanism will be present in the case of more complex variation of shear, as is common in real flows. In the normal limit where there is no mean shear, the forcing and response are in the same velocity component. In the pseudoresonant limit, the forcing acts in the n_2 direction while the response is in the n_1 direction.

3.2.2 Eigenvectors, singular functions, and self-adjointness of the operator

The objective is to isolate the effect (or lack thereof) of self-adjointness on the characteristics of the resolvent modes. To this end, the mean shear term is set to zero, i.e. $d = 0$ in Equation 3.3, guaranteeing that the LNS and resolvent operators are normal:

$$\mathbf{M} = \begin{pmatrix} m_1 & 0 \\ 0 & m_2 \end{pmatrix}, \quad (3.9)$$

$$\mathcal{H}(\omega) = \begin{pmatrix} 1/(m_1 - i\omega) & 0 \\ 0 & 1/(m_2 - i\omega) \end{pmatrix}. \quad (3.10)$$

The discussion is limited to stationary disturbances ($\omega = 0$) once again as the presence of $i\omega$ in the resolvent operator will always guarantee that it is not self-adjoint.

If the eigenvalues are real, \mathbf{M} and $\mathcal{H}(\omega = 0)$ are self-adjoint and the SVD of the resolvent is

$$\lim_{\text{Imag}(m_j) \rightarrow 0} \text{SVD}(\mathcal{H}(\omega = 0)) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} m_1 & 0 \\ 0 & m_2 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (3.11)$$

Equation 3.12 is similar to the normal limit case of Equation 3.6 and illustrates that the resolvent forcing and response modes are identical.

The simplest way to demonstrate what happens when the resolvent is not self-adjoint is to assume that the eigenvalues are purely imaginary. In this case, the

SVD of the resolvent is

$$\lim_{\text{Real}(m_j) \rightarrow 0} \text{SVD}(\mathcal{H}(\omega = 0)) = \begin{pmatrix} i & 0 \\ 0 & i \end{pmatrix} \begin{pmatrix} \|m_1\| & 0 \\ 0 & \|m_2\| \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}. \quad (3.12)$$

Since the singular values of the resolvent are required to be real yet the eigenvalues are imaginary, the resolvent response modes must be 90° out of phase from the resolvent forcing modes as portrayed in Figure 3.1(e-f).

3.2.3 Mean flow advection and the Orr mechanism

In the case of an operator that is not self-adjoint and has non-zero mean shear, i.e. $d \neq 0$, the spatio-temporal manifestation of the Orr mechanism is observed. This mechanism results in disturbance amplification by the reorientation of an input tilted against the mean shear into a response which is aligned with the mean shear, with the maximum amplification in the transient case occurring when the disturbance is vertical.

In the continuously forced formulation of the present problem, the Orr mechanism is attributable to the space dependence of $\bar{\mathbf{u}}$ in the mean flow advection term. Assuming a parallel flow ($\bar{\mathbf{u}} = \bar{u}(y)$) with non-zero mean shear ($\partial \bar{u} / \partial y \neq 0$) and traveling at the local mean velocity for a given y_c , the flow above y_c moves towards the downstream direction while the flow below moves upstream in a relative sense. There is, therefore, a phase difference of π across y_c which manifests itself in the response modes (e.g. McKeon, 2017). The decrease in phase results in the response modes leaning downstream, aligned with the mean shear. When considering the adjoint LNS operator \mathbf{L}^* , the mean flow advection acts in the upstream direction due to the sign change of $\bar{\mathbf{u}} \cdot \nabla(\cdot)$. The corresponding phase jump across y_c is now in the opposite sense and this results in the forcing modes leaning upstream. Artificially removing mean flow advection from the resolvent operator suppresses the Orr mechanism. Figure 3.1(g-h) is a cartoon illustrating how the Orr mechanism results in the tilting of the resolvent forcing and response modes.

3.2.4 Resolvent (approximate) wavemaker

For a spatially-varying base or mean flow, further statements can be made concerning the effect of the non-self-adjoint nature of the resolvent operator on the resolvent mode shapes, namely a difference in the spatial support of forcing and response modes, with the latter being downstream of the former. For a spatially-developing base or mean flow, normal mechanisms can be categorized in terms of either convective or absolute instability depending on the characteristics of the pro-

file. The non-self-adjoint nature of \mathbf{L} changes the influence of the convective terms in the adjoint operator, per Equations 2.10 and 2.16, or, for the model operator of Equation 3.3,

$$\mathbf{M}^* = \begin{pmatrix} m_1^* & 0 \\ d & m_2^* \end{pmatrix}. \quad (3.13)$$

The direction of mean flow advection is reversed since the adjoint of the derivative operator introduces a negative sign implying that adjoint perturbations are transported upstream. Direct or forward perturbations, on the other hand, are transported downstream. For absolutely unstable flows, for which perturbations grow both upstream and downstream of the source, the advection term may no longer separate the spatial support of the forcing/adjoint and response/forward modes leading to regions of overlap at resonant frequencies. This region, known as the wavemaker, is traditionally computed from the eigenmodes, and is associated with non-zero values of \mathcal{W} , where

$$\mathcal{W}(\mathbf{x}_0) = \|\tilde{\mathbf{u}}(\mathbf{x}_0)\| \|\tilde{\mathbf{u}}^\dagger(\mathbf{x}_0)\|, \quad (3.14)$$

and \mathbf{x}_0 denotes a position in space (see derivation of Giannetti and Luchini, 2007).

In cases where amplification is due to normal mechanisms, the resolvent modes can be used to find the wavemaker as long as they are normalized appropriately. The wavemaker approximates regions of the flow which are absolutely unstable or self-sustaining since perturbations are prevented from convecting due to reverse flow (Juniper, 2012). Huerre and Monkewitz (1985) have shown that when a mean profile of hyperbolic tangent form exhibits greater than 13.6% reverse flow with respect to the free stream, the flow is absolutely unstable. The streamwise extent of absolute instability and the wavemaker is finite since flow reversal is confined to a certain portion of the flow. This information is encoded within the advection term $\bar{\mathbf{u}} \cdot \nabla()$ through the sign of $\bar{\mathbf{u}}$. The overlap of the resolvent forcing and response modes was identified as a qualitative proxy for sensitivity for base flows by Brandt et al. (2011). Henceforth, the wavemaker of Equation 3.14 is denoted as the “true” wavemaker and the approximation using resolvent modes as the “resolvent” wavemaker. If the flow is convectively unstable, there is no region of reverse flow and so $\bar{\mathbf{u}}$ is always positive. In this case, the optimal response or stability mode will be downstream of the optimal forcing or adjoint mode. The inner product $|\hat{\phi}_1^* \hat{\psi}_1|$ decreases when the modes are separated in space. Chomaz (2005) noted that this is due to the convective-type non-normality introduced from advection of the base flow.

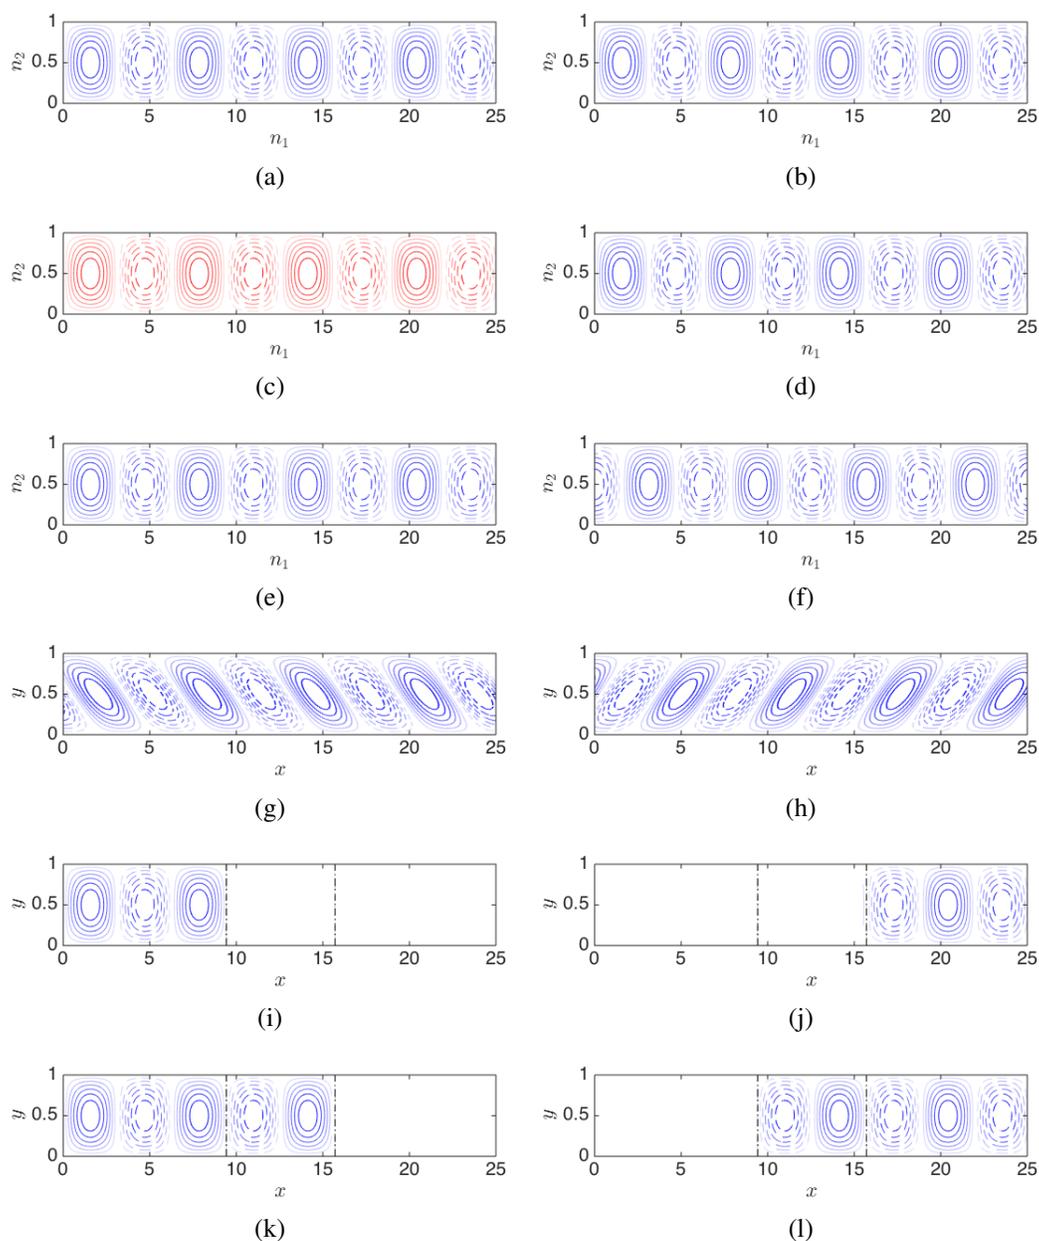


Figure 3.1: Cartoon of forcing (left column) and response (right column) modes for various amplification mechanisms. Panels (a) and (b) are the forcing and response for a completely normal operator. Panels (c) and (d) are the forcing and response in Equation 3.3, for a component-type non-normality. Panels (e) and (f) depict a $\pi/2$ phase difference between the forcing and response mode. Panels (g) and (h) illustrate the action of the Orr mechanism where the forcing mode leans upstream against the mean shear and the response mode leans downstream and is aligned with the mean shear. Panels (i) and (j) denote a convective-type non-normality where the forcing is upstream of the response. This is a convective instability while panels (k) and (l) include a region of overlap between the modes where the flow is absolutely unstable. Positive/negative isocontours are denoted by solid/dotted lines and blue/red colors indicate streamwise/transverse components, in the n_1 or x and n_2 or y directions, respectively.

3.3 Lower Rank Approximations of the Resolvent Operator

The conditions under which analysis of the resolvent is likely to identify eigenmodes as the most amplified disturbance are formalized in this section. In these scenarios, it is also feasible to differentiate the contributions from resonance and non-normality to the resolvent norm.

3.3.1 Dyad expansion of the Resolvent Operator

As is more customary for the eigenvalue problem (Luchini and Bottaro, 2014; Schmid and Brandt, 2014) a dyad expansion of the resolvent \mathcal{R} for a generic, non-singular linear operator \mathbf{Q} can be performed,

$$\mathcal{R} = (z\mathbf{I} - \mathbf{Q})^{-1} = \sum_{j=1}^n \frac{1}{z - \lambda_j} \tilde{\mathbf{g}}_j \tilde{\mathbf{h}}_j^*, \quad (3.15)$$

where $\tilde{\mathbf{g}}_j$ and $\tilde{\mathbf{h}}_j$ are the j th left and right eigenvectors of \mathbf{Q} , respectively. Since the objective of resolvent analysis is often the identification of the most amplified neutral disturbance, $i\omega$ is substituted for z and the eigenvectors of the LNS operator for $\tilde{\mathbf{g}}_j$ and $\tilde{\mathbf{h}}_j$ into Equation 3.15 to give

$$\mathcal{H}(\omega) = \sum_{j=1}^n \frac{1}{i\omega - \lambda_j} \tilde{\mathbf{u}}_j \tilde{\mathbf{v}}_j^*. \quad (3.16)$$

Thus if the real part of an eigenvalue λ_p is sufficiently close to zero and the forcing frequency ω is identical to the imaginary part, then its contribution to the series dominates over the contributions from all other eigenvalues. The resolvent, furthermore, can be approximated by the forward and adjoint eigenvectors corresponding to that frequency weighted by the inverse distance between the eigenvalue and the imaginary axis:

$$\mathcal{H}(\omega) \approx \frac{1}{i\omega - \lambda_p} \tilde{\mathbf{u}}_p \tilde{\mathbf{v}}_p^*. \quad (3.17)$$

Equation 3.17 represents a rank-1 approximation of the resolvent operator using eigenvectors. In the context of base flows, the resolvent is singular at the critical Reynolds number since the real part of the least stable eigenvalue is identically zero when it crosses the imaginary axis. In the case of mean flows, which tend to be marginally stable (e.g. Reynolds and Tiederman, 1967; Barkley, 2006; Turton et al., 2015), Equation 3.17 is applicable for eigenvalues near the imaginary axis. It is important to note that an eigenvalue does not have to be marginally stable, but it must be the dominant contribution to the series in Equation 3.16. It is possible to

obtain a rank-1 approximation of the resolvent even when an eigenvalue is highly damped.

This rank-1 approximation fails when there is not sufficient separation of eigenvalues at the frequency of interest as will be seen for cylinder flow. If there are several eigenvalues in the vicinity of the imaginary axis at a frequency ω , then the resolvent operator can no longer be approximated by just one outer product in Equation 3.17. Equating the two low-rank approximations of the resolvent operator in terms of eigenvectors (Equation 3.17) and resolvent modes (Equation 2.25) implies the following:

$$\sigma_1 \hat{\psi}_1 \hat{\phi}_1^* \approx \frac{1}{i\omega - \lambda_r} \tilde{\mathbf{u}}_r \tilde{\mathbf{v}}_r^* \implies \hat{\psi}_1 \propto \tilde{\mathbf{u}}_r, \quad \hat{\phi}_1 \propto \tilde{\mathbf{v}}_r, \quad (3.18)$$

since

$$\hat{\psi}_1 \approx \frac{1}{\sigma_1(i\omega - \lambda_r)} \tilde{\mathbf{u}}_r \tilde{\mathbf{v}}_r^* \hat{\phi}_1 = C \tilde{\mathbf{u}}_r, \quad (3.19)$$

where C is a complex constant. The leading resolvent response and forcing modes are proportional to the forward and adjoint eigenmodes, respectively, and this holds for any base or mean flow as long as only one eigenvalue leads to amplification. The similarity between the resolvent forcing and adjoint stability modes draws out how the resolvent operator contains sensitivity information, as described by, e.g., Qadri and Schmid, 2017. The development is less amenable to pseudoresonant mechanisms where the proximity of an eigenvalue to the imaginary axis does not necessarily govern the behavior of the resolvent.

3.3.2 The Relationship Between Spectral Radius and Spectral Norm for Approximately Low-Rank Operators

For a nonsingular linear operator \mathbf{Q} , an explicit relationship between the spectrum of \mathbf{Q} and the spectral norm of its resolvent $\mathcal{R}(z) = (z\mathbf{I} - \mathbf{Q})^{-1}$ is desired. Of particular interest are cases where $\mathcal{R}(z)$ is approximately low-rank (i.e., a small number of leading singular values are much larger than the others).

The spectral radius Θ of an operator \mathbf{Q} can be defined through the eigendecomposition $\mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$

$$\Theta(\mathbf{Q}) = \max_{\lambda_j \in \mathbf{\Lambda}} (|\lambda_j|). \quad (3.20)$$

The spectral radius of the corresponding resolvent operator is

$$\Theta(\mathcal{R}(z)) = \max_{\lambda_j \in \mathbf{\Lambda}} (|z - \lambda_j|^{-1}) = \left[\min_{\lambda_j \in \mathbf{\Lambda}} (|z - \lambda_j|) \right]^{-1}. \quad (3.21)$$

Note that, with this definition, Equation 3.2 may be expressed as

$$\Theta(\mathcal{R}) \leq \sigma_1 \leq \kappa\Theta(\mathcal{R}).$$

For non-normal operators with large condition numbers, the upper and lower bounds span a large range, and thus do not give much insight into the size of the resolvent norm, σ_1 . To estimate the resolvent norm in terms of the spectral radius, the following relationship from Gelfand (1941) can be used:

$$\Theta(\mathcal{R}) = \lim_{n \rightarrow \infty} \|\mathcal{R}^n\|^{1/n}. \quad (3.22)$$

Suppose now that the largest singular value of resolvent operator (for a given z) is much larger than the rest, such that

$$\mathcal{R}(z) = \Psi \Sigma \Phi^* \approx \Psi \Sigma_1 \Phi^*, \quad (3.23)$$

where Σ_1 is Σ with all but the first singular value set to zero. Suppose in addition that this truncation is also accurate for powers of \mathcal{R} , i.e.,

$$\mathcal{R}^n \approx (\Psi \Sigma_1 \Phi^*)^n. \quad (3.24)$$

Defining the quantity

$$r_{ij} = \frac{\hat{\phi}_i^* \hat{\psi}_j}{\hat{\phi}_j^* \hat{\psi}_j}, \quad (3.25)$$

implies that

$$(\Phi^* \Psi \Sigma_1)^n = \sigma_1^n (\hat{\phi}_1^* \hat{\psi}_1)^n \begin{pmatrix} r_{11} & 0 & \cdots & 0 \\ r_{21} & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \end{pmatrix}. \quad (3.26)$$

The norm of powers of the resolvent may be estimated as

$$\|\mathcal{R}^n\| = \|\Phi^* \mathcal{R}^n \Phi\| \approx \|(\Phi^* \Psi \Sigma_1)^n\| = \sigma_1^n |\hat{\phi}_1^* \hat{\psi}_1|^{n-1} \|\mathbf{r}\|, \quad (3.27)$$

where $\mathbf{r} = [r_{11} \ r_{21} \ \cdots]^T$. Consequently, assuming that Equation 3.24 holds, Equation 3.22 results in the estimate

$$\Theta(\mathcal{R}) \approx \sigma_1 |\hat{\phi}_1^* \hat{\psi}_1|. \quad (3.28)$$

In other words, the resolvent norm is estimated to be larger than the lower bound in Equation 3.2 by a factor of $|\hat{\phi}_1^* \hat{\psi}_1|^{-1}$. This analysis relied on the rather restrictive assumption that only the leading singular value was large. If there is a pair of large singular values, as is often the case in channel flows (owing to spatial symmetry

across the mid-plane of the channel) then the argument may be generalized as follows. Suppose that σ_1 and σ_2 are of comparable size, and that all other singular values are negligibly small. If it is assumed that $|\hat{\phi}_1^* \hat{\psi}_2|, |\hat{\phi}_2^* \hat{\psi}_1| \approx 0$, then

$$\begin{aligned} (\Phi^* \Psi \Sigma)^n &\approx \begin{pmatrix} \sigma_1 \hat{\phi}_1^* \hat{\psi}_1 & \sigma_2 \hat{\phi}_1^* \hat{\psi}_2 & 0 & \cdots & 0 \\ \sigma_1 \hat{\phi}_2^* \hat{\psi}_1 & \sigma_2 \hat{\phi}_2^* \hat{\psi}_2 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \end{pmatrix}^n \\ &\approx \sigma_1^n |\hat{\phi}_1^* \hat{\psi}_1|^{n-1} \begin{pmatrix} r_{11} & 0 & \cdots & 0 \\ r_{21} & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \end{pmatrix} + \sigma_2^n |\hat{\phi}_2^* \hat{\psi}_2|^{n-1} \begin{pmatrix} 0 & r_{12} & 0 & \cdots & 0 \\ 0 & r_{22} & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \end{pmatrix}, \end{aligned}$$

which, following the same approach as before, gives

$$\Theta(\mathcal{R}) \approx \max\{\sigma_1 |\hat{\phi}_1^* \hat{\psi}_1|, \sigma_2 |\hat{\phi}_2^* \hat{\psi}_2|\}. \quad (3.29)$$

Thus the inverse of $|\hat{\phi}_1^* \hat{\psi}_1|$ can be interpreted as the contribution of non-normality to the resolvent norm. The product $\sigma_1 |i\omega - \lambda|$ can also be identified as a quantification of non-normality since $|i\omega - \lambda|^{-1}$ represents the resonance contribution to the resolvent norm. However, since highly amplified modes may occur at non-resonant frequencies, the contribution from $|i\omega - \lambda|$ is typically overestimated as it is likely for a pseudoeigenvalue to reside much closer to the imaginary axis than the nearest eigenvalue of the unperturbed spectrum. These two predictions tend to agree in cases where amplification can be attributed to a single eigenvalue and mean stability analysis is valid.

3.4 Resonance curve and resolvent norm

To demonstrate the amplification characteristics of the resolvent for normal and non-normal operators, the pseudo-spectra are sketched in Figure 3.2 for simple example operators,

$$\mathbf{M}_1 = \begin{pmatrix} -1.5 + 1.1i & 0 \\ 0 & -1.9 - 2.2i \end{pmatrix}, \quad (3.30)$$

$$\mathbf{M}_2 = \begin{pmatrix} -1.5 + 1.1i & 5 \\ 0 & -1.9 - 2.2i \end{pmatrix}, \quad (3.31)$$

where \mathbf{M}_1 is a normal operator containing only the eigenvalues of the non-normal operator \mathbf{M}_2 . Level curves of ϵ for operators \mathbf{M}_1 and \mathbf{M}_2 satisfy

$$\Lambda_\epsilon(\mathbf{M}_1) = \{z \in \mathbb{C} : \|(z\mathbf{I} - \mathbf{M}_1)^{-1}\| \geq \epsilon^{-1}\}, \quad (3.32)$$

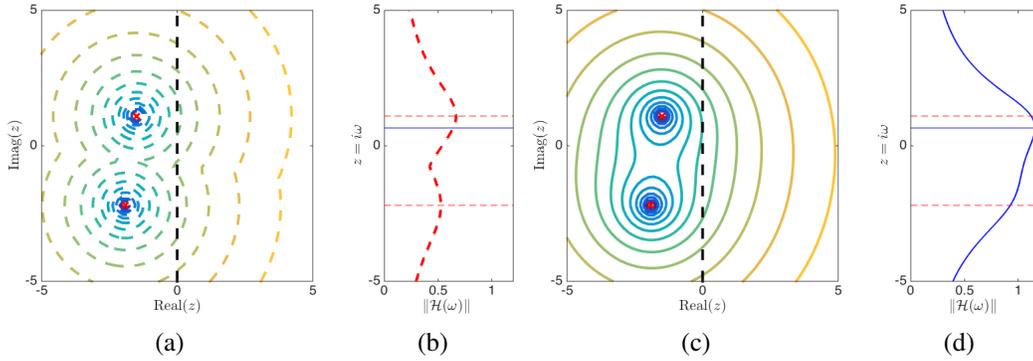


Figure 3.2: Comparison of the pseudospectra and resolvent norm for the operators given in Equations 3.30 and 3.31, which have the same eigenvalues. (a, b) normal operator \mathbf{M}_1 , (c, d) non-normal operator \mathbf{M}_2 . The eigenvalues, i.e., the eigenspectrum, are marked by red crosses and color contours outline the bounds of the perturbed spectrum for constant perturbation magnitudes in (a, c). The dashed contours in (a) reflect that the pseudospectra are circles centered on the eigenvalues, which is not true for the non-normal operator in (c). The resolvent norm in each case (b, d) reflects the value of these contours along the imaginary axis. Red, dashed horizontal lines indicate the resonant frequencies of the operator, i.e., the frequencies corresponding to the eigenvalues, while the blue, solid horizontal line represents the most highly amplified frequency in the non-normal case.

and

$$\Lambda_\epsilon(\mathbf{M}_2) = \{z \in \mathbb{C} : \|(z\mathbf{I} - \mathbf{M}_2)^{-1}\| \geq \epsilon^{-1}\}, \quad (3.33)$$

respectively.

For a normal operator such as \mathbf{M}_1 , $\kappa = 1$ and the level curves of ϵ are proportional to the distance from the closest eigenvalue. The resolvent norm for a particular ω is inversely proportional to the distance from $i\omega$ to the nearest eigenvalue. This shall be referred to as the resonance curve. The spectrum and pseudospectra of \mathbf{M}_1 are shown in Figure 3.2(a); there are two stable eigenvalues denoted by red crosses, and the pseudospectra consist of circular contours centered on the two eigenvalues. Since both eigenvalues are significantly damped, amplification due to resonance is not possible and the magnitude of the resolvent norm is less than one (Figure 3.2(b)). Moreover, the eigenvalue and singular value decompositions of \mathbf{M}_1 yield parallel basis functions, and the singular values are simply the magnitude of the eigenvalues.

Operator \mathbf{M}_2 , however, is non-normal due to the non-zero off-diagonal term and, with reference to Equation 3.2, $\kappa > 1$. The shifts of the eigenvalues of the per-

turbed operator are not proportional to ϵ , as indicated by the pseudospectrum isocontours in Figure 3.2(c), and the resolvent norm of Figure 3.2(d) is appreciably larger than that of the normal case in Figure 3.2(a), with values exceeding one. Furthermore, the maximum value of the resolvent norm occurs at a non-resonant frequency, $\omega = 0.66$. Amplification is possible for a linearly stable operator due to pseudoresonance under forcing at any frequency for which $\|\mathcal{H}(\omega)\| > 1$. It is important to add that even when the primary contribution to amplification is a normal mechanism, non-normality can still contribute to exacerbate the response. For example, the amplification at the frequencies of the eigenvalues in Figure 3.2 is higher for operator \mathbf{M}_2 than it is for the purely normal operator \mathbf{S} . That being the case, the right-hand side of Equation 3.2 may be large due to one or both terms in the product.

The projections of the first resolvent mode $\hat{\psi}_1$ onto the eigenvectors of operators \mathbf{M}_1 and \mathbf{M}_2 are plotted in Figure 3.3 for various ω . For \mathbf{M}_1 , the values are either one or zero meaning that $\hat{\psi}_1$ is one of the operator's eigenvectors, even at non-resonant frequencies. The eigenvector it chooses is simply whichever eigenvalue is closest to $z = i\omega$ for a given ω . The projections for operator \mathbf{M}_2 are far more interesting. There is a nontrivial projection of $\hat{\psi}_1$ onto both eigenvectors for every ω , hence why the resolvent norm in Figure 3.2 is higher for operator \mathbf{M}_2 at every frequency. The maximum $|\hat{\psi}_1^* \tilde{u}_1|$ and $|\hat{\psi}_1^* \tilde{u}_2|$ occur at λ_1 and λ_2 , respectively, which is to be expected since the forcing is at the frequency of these eigenvalues. The projections at the most amplified frequency ($\omega = 0.66$) are both high, which can be seen by their product. The frequency where the product peaks, however, does not match the most amplified frequency since $Real(\lambda_1) < Real(\lambda_2)$, hence the resolvent's preference for choosing a frequency closer to λ_1 .

Since both eigenvectors are needed to capture the behavior of operator \mathbf{M}_2 but only one resolvent mode is needed, it can be concluded that resolvent modes are a more efficient basis for capturing the dominant input-output behavior of the operator. It can also be inferred that the projection of $\hat{\psi}_2$ onto the eigenvectors is relatively small since it is orthogonal to $\hat{\psi}_1$. It is no surprise, therefore, that the resolvent operator tends to be low-rank at pseudo-resonant frequencies since the optimal resolvent response mode projects onto many eigenvectors. This forces suboptimal resolvent modes to be nearly orthogonal to many eigenvectors of the LNS operator and so their contribution to the input-output behavior is negligible.

Having examined the implications of the structure of the operator on amplification

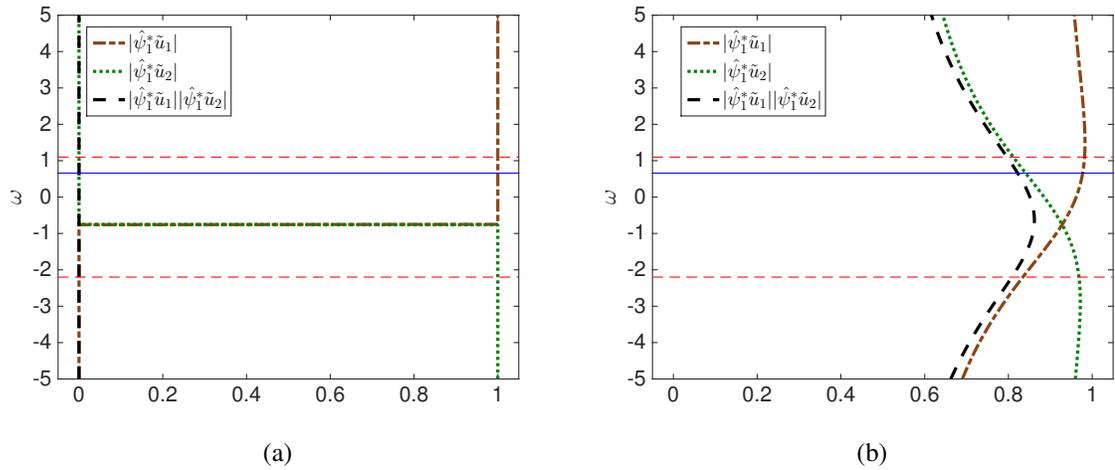


Figure 3.3: Projection of the first resolvent mode $\hat{\psi}_1$ onto the eigenvectors $\tilde{\mathbf{u}}_1$ (brown dashed-dotted line) and $\tilde{\mathbf{u}}_2$ (green dotted line). The product of these projections is the black, long dashed line. Red, dashed horizontal lines indicate the resonant frequencies of the operators and the blue, solid horizontal line represents the most highly amplified frequency in the non-normal case.

and forcing and response modes, the remainder of this chapter is devoted to analysis of two real example flows. Low Reynolds number cylinder flow is used to investigate the choice of base or mean flow as the linear stability threshold is crossed. A canonical wall turbulence configuration is employed to identify the influence of the various terms in the resolvent on the resulting singular value decomposition.

3.5 Application to Cylinder Flow

Global resolvent analysis is applied to the base and mean velocity profiles for cylinder flows under the critical Reynolds number $Re_c \leq 47$ (Provansal et al., 1987; Sreenivasan et al., 1987; Noack and Eckelmann, 1994), as well as mean flows of the 2D laminar vortex shedding regime where $Re \leq 189$ (Barkley and Henderson, 1996). Cylinder flow is a particularly suitable choice to investigate trends associated with the wavemaker since it exhibits a region of absolute instability.

3.5.1 Base flow velocity profile

A resolvent analysis is performed on the cylinder base flow for various Reynolds numbers over a range of ω . Contours of the pseudospectrum for $Re = 47$ are overlaid onto the spectrum, which is in agreement with Sipp and Lebedev (2007) to within the sensitivity to the mesh geometry, in Figure 3.4(a). The variation of the resolvent norm along the imaginary axis, i.e. σ_1 , is plotted alongside the second

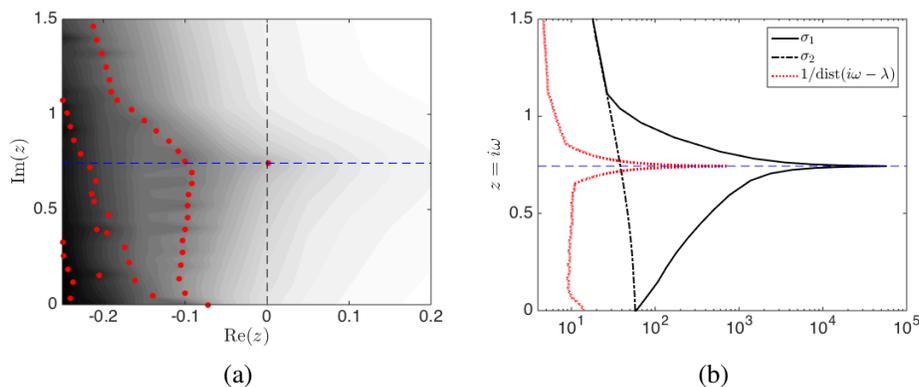


Figure 3.4: (a) Spectrum (red dots) and pseudospectrum (filled contours, ϵ increasing as colors change from dark to light) of the LNS operator for the cylinder base flow at $Re = 47$. (b) The resolvent norm, σ_1 (solid line), i.e. the value of ϵ^{-1} along the imaginary axis, second largest singular value σ_2 (dash-dotted line) and inverse distance from the imaginary axis to the nearest eigenvalue (dotted red line).

singular value of the resolvent, σ_2 , and the resonance curve in Figure 3.4(b). There is only one frequency ω_{max} where the first singular value is several orders of magnitude larger than all the others (only two are shown for clarity). The least stable eigenvalue, whose imaginary part is ω_{max} , dominates the behavior of the resolvent norm and its influence spans from $0 < \omega < 1.1$. For the cylinder base flow at $Re = 47$, $\omega_{max} = 0.742$ corresponds to the true ω_s at the onset of vortex shedding. The amplification is significantly lower for all other frequencies, including harmonics, and the resolvent is not low-rank since the eigenvalues are not separated for $\omega > 1.1$.

If the resolvent norm had been computed for Re_c exactly, it would be infinite at ω_{max} since the resolvent is undefined when there is an eigenvalue on the imaginary axis. The finite value quoted in Figure 3.4(b), therefore, is very sensitive to small changes of the Reynolds number or the mesh resolution as this may affect the real part of the least stable eigenvalue. It is misguided to ascribe significant importance to the precise value quoted in Figure 3.4(b). What is important is the separation of singular values, which implies a strong linear (in this case stability) mechanism, and a single peak. This reinforces why stability analysis about the base flow can predict Re_c in contrast to stability analysis of wall-bounded shear flows, which fails to predict the transition to instability by not taking into account non-normality and transient growth. In the case of the cylinder, only one structure at the globally most amplified frequency is prone to significant amplification at subcritical Reynolds

numbers and it is the first to become unstable. This is characteristic of an absolute instability mechanism in which frequency selection is not influenced by background noise.

The stability modes and resolvent modes are nearly identical as seen in Figure 3.5. The effect of lift-up is weak since the energy is fairly evenly distributed in the u - and v -components of both the forcing and response modes in Figure 3.5. Mean flow advection, on the other hand, plays a significant role in the spatial support of the forcing and response modes which are located upstream and downstream of the cylinder. Chomaz (2005) made an analogous observation for the forward and adjoint eigenmodes and attributed this to convective non-normality. Since the resolvent operator is low-rank, computing $|\hat{\phi}_1^* \hat{\psi}_1|^{-1}$ is a good estimate of the non-normal amplification experienced by the flow. A value of $|\hat{\phi}_1^* \hat{\psi}_1|^{-1} = 79.4$ is obtained, which is in good agreement with $\sigma_1(\omega_{max})|i\omega_{max} - \lambda_{ls}| = 79.3$ (see Table 3.1). This accounts for the large gap between the peaks of the resolvent norm and resonance curve in Figure 3.4(b). Despite the fact that the resolvent norm and real part of the least stable eigenvalue are sensitive to the mesh, the agreement between the two measures of non-normality is robust.

The least stable global mode and its adjoint counterpart are computed for various Reynolds numbers near and below Re_c to illustrate the cylinder transition from convective to absolute instability. Figure 3.6 juxtaposes the v -component of the adjoint mode, forward mode, and wavemaker. The forward mode has unit magnitude while the adjoint has been normalized with respect to the forward mode such that their inner product is unity. A wavemaker first appears for $Re = 25$, the Reynolds number at which Monkewitz (1988) determined the cylinder wake is absolutely unstable. There is no wavemaker for lower Reynolds numbers due to the downstream location of the forward eigenmode which is a consequence of mean flow advection; the strength of the reverse flow is not sufficient to produce an overlap region. For the lowest two Reynolds numbers considered in Figure 3.6, the contour levels of the forward eigenmode immediately behind the cylinder are three orders of magnitude smaller than their higher Reynolds number counterparts. The downstream location where the contour levels are significant does not appear within the plotted domain. As the Reynolds number increases, the velocity deficit grows and the reverse flow directly behind the cylinder strengthens. The forward eigenmode gradually appears closer to the cylinder until there is a nontrivial overlap between it and its adjoint counterpart.

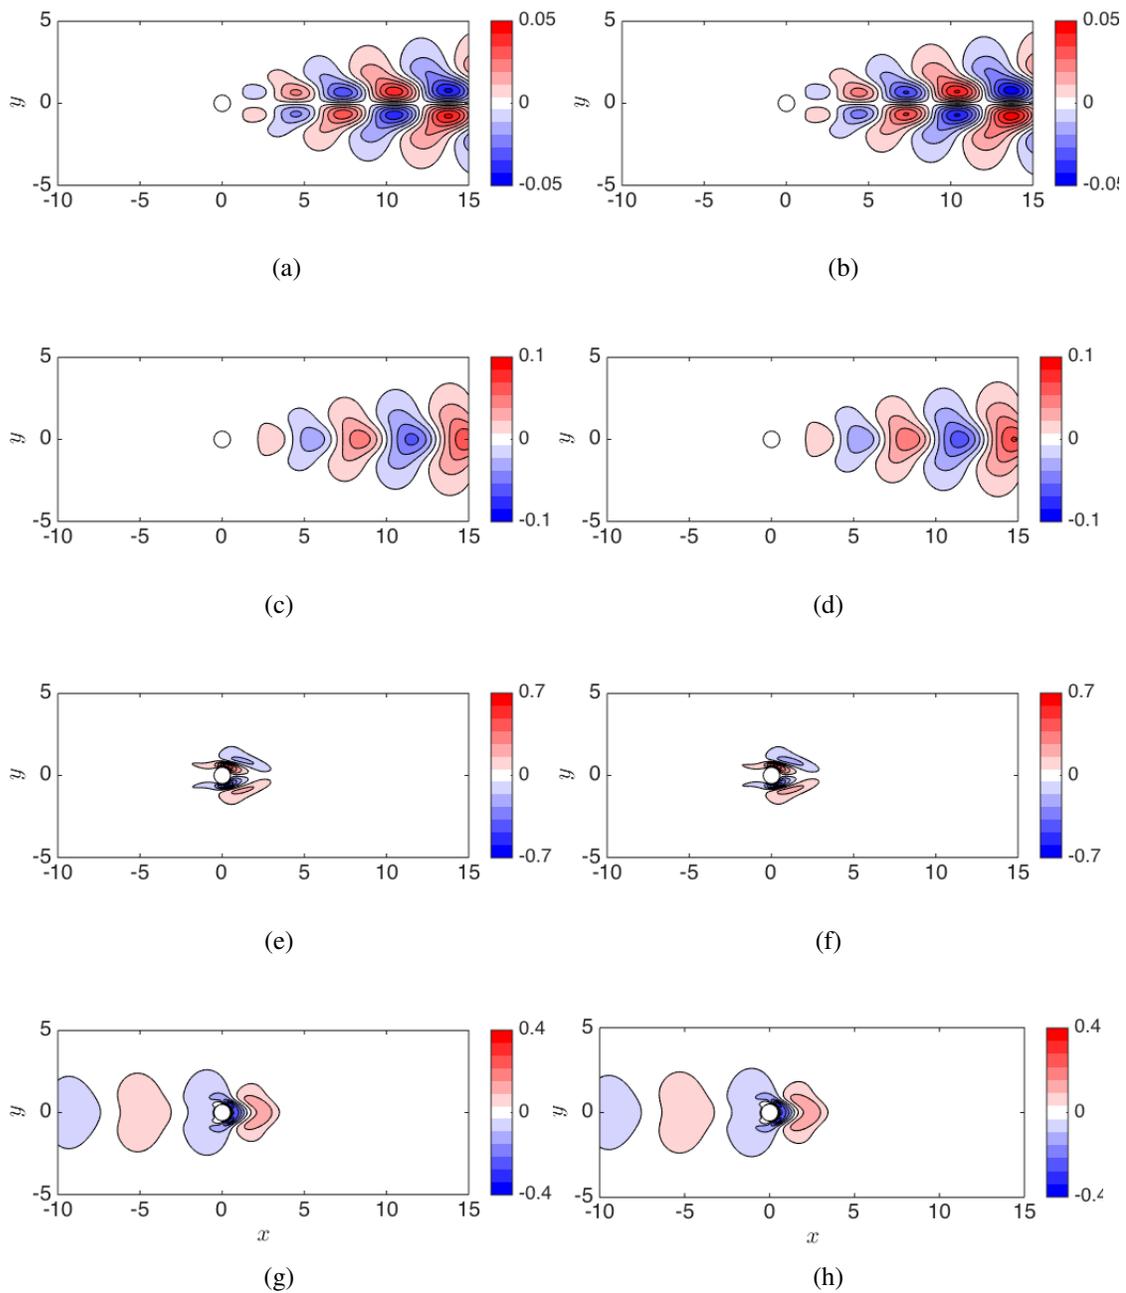


Figure 3.5: Comparison of stability modes (left) with resolvent modes (right) at the critical Reynolds number $Re_c = 47$ and a temporal frequency of $\omega = 0.742$. Panels (a, b) are the streamwise component of the forward or response mode, (c, d) are the transverse component of the forward or response mode, (e, f) are the streamwise component of the adjoint or forcing mode, and (g, h) are the transverse component of the adjoint or forcing mode. The eigenmodes and resolvent mode shapes are essentially indistinguishable for this flow.

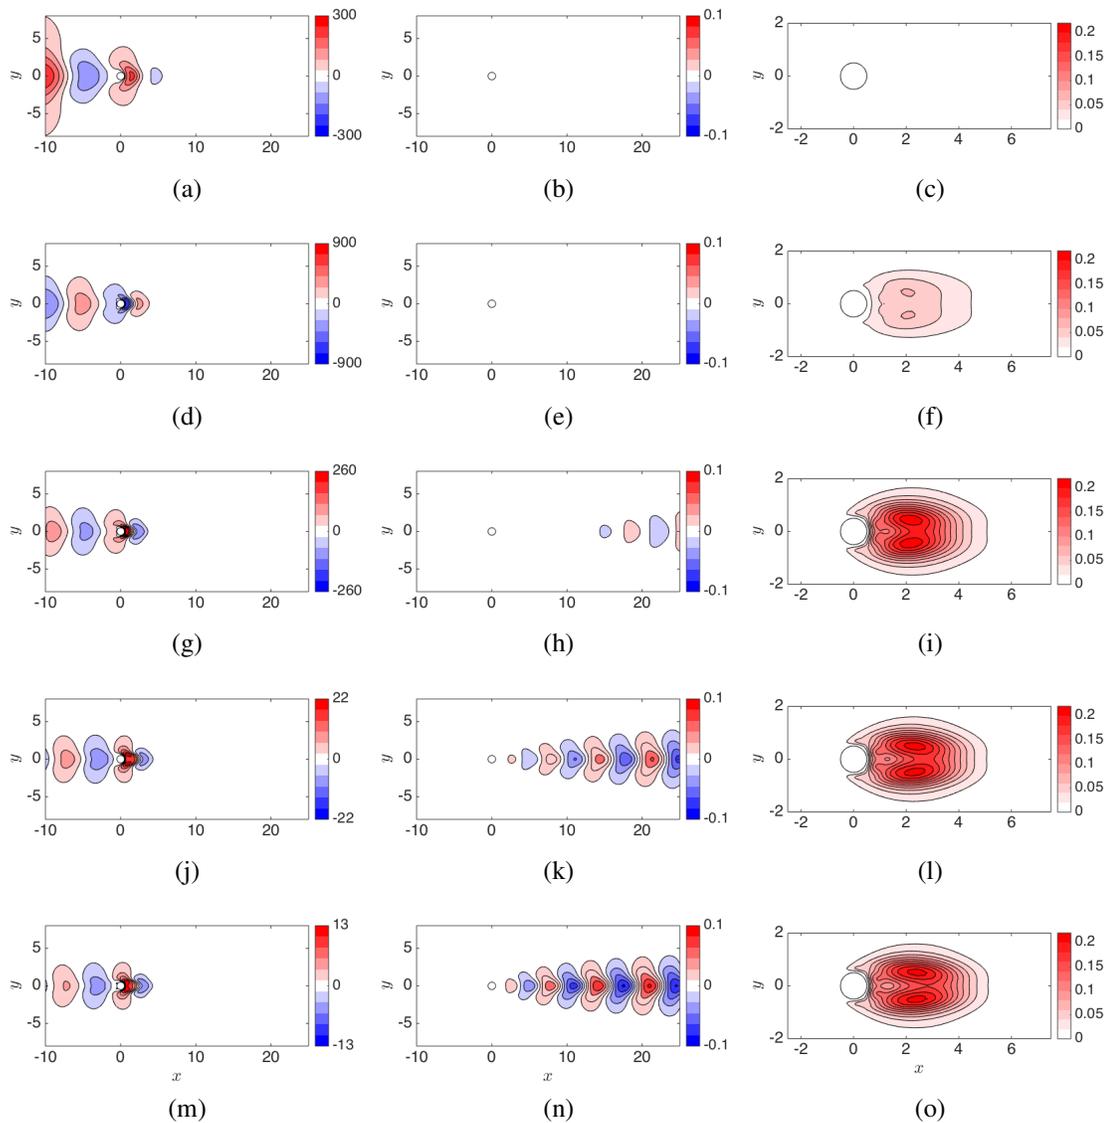


Figure 3.6: Contours of the transverse velocity for the leading adjoint modes \tilde{v}^\dagger (left) and forward modes \tilde{v} (middle) of the base flow. The Reynolds numbers ($Re = 15; 25; 35; 45; \text{ and } 50$) increase from top to bottom. The wavemaker \mathcal{W} (right) is computed using the forward and adjoint modes. Contour levels are not identical for the adjoint modes which are normalized based on the forward modes. Note that the streamwise velocity component has not been plotted even though the wavemaker depends on this quantity.

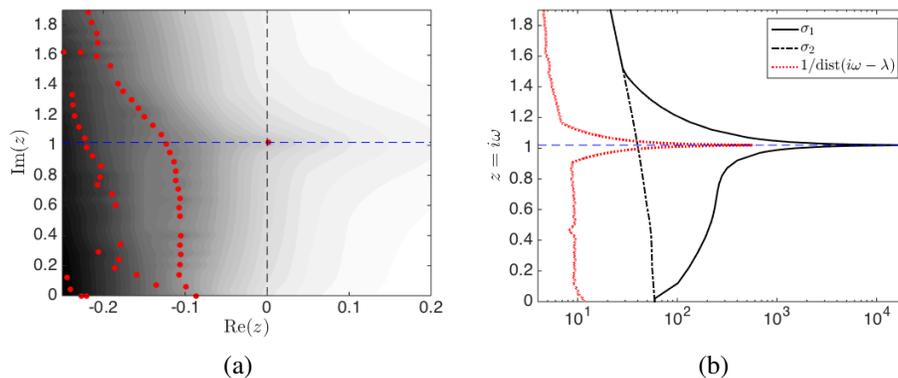


Figure 3.7: (a) Spectrum (red dots) and pseudospectrum (filled contours, ϵ increasing as colors change from dark to light) of the LNS operator for the cylinder mean flow at $Re = 100$. (b) The resolvent norm, σ_1 (solid line), i.e., the value of ϵ^{-1} along the imaginary axis, second largest singular value σ_2 (dash-dotted line) and inverse distance from the imaginary axis to the nearest eigenvalue (dotted red line).

Beyond the Re_{crit} , the region of the flow which is absolutely unstable is sufficiently long for the flow to become globally unstable. Perturbations grow exponentially in time until they are saturated by nonlinearities. The resulting velocity fluctuations are dominated by the vortex shedding. Once the flow has reached a limit cycle, the shedding frequency is different from that predicted by resolvent analysis of the base flow since the frequency of the least stable perturbations is altered during the saturation process. Additionally, the mean recirculation region behind the cylinder is shorter than its base flow counterpart in the streamwise direction.

3.5.2 Mean velocity profile

The focus of resolvent analysis typically shifts when using the mean velocity profile rather than the base flow. The goal becomes identification of the energetically important structures and their frequencies in the unsteady flow rather than prediction of the external forcing and structure which appears when the flow becomes unsteady.

Contours of the pseudospectrum corresponding to the mean flow in the 2D laminar shedding regime are overlaid with the spectrum of the mean flow at $Re = 100$ in Figure 3.7. The resolvent norm along the imaginary axis is also plotted alongside the second largest singular value and resonance curve. Similar to the base flow resolvent, there is only one frequency at which there is a resonant peak. Unlike the base flow case, the most amplified frequency at supercritical Reynolds numbers

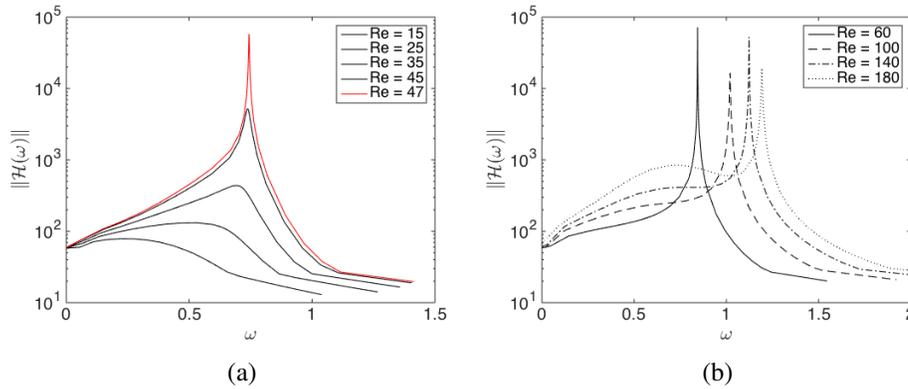


Figure 3.8: (a) The resolvent norm for the critical (solid red line), subcritical (solid black lines), and supercritical (dotted black lines) base flows. (b) The resolvent norm for supercritical mean flows.

correctly predicts the shedding frequency as seen in Figure 3.8. While the resolvent norm always peaks at a distinct frequency for all cases, the growth rate of the least stable eigenvalue of the base flow continues to grow while the frequency remains roughly constant. Figure 3.8 shows that the largest peak occurs at the stability limit, Re_{crit} . The maximum amplification, which here is proportional to the inverse distance between the eigenvalue and the imaginary axis, indicates the progression of the least stable pole across the complex plane and over the imaginary axis. The resolvent norm has not been plotted for supercritical base flows since the resolvent attempts to quantify the size of perturbation necessary for the spectrum to cross the neutral axis. For the base flow at the critical Reynolds number Re_{crit} and mean flows where $Re > Re_{crit}$, the size of this perturbation is very small, leading to very highly amplified disturbances.

The peak resolvent norm for the mean flows has no discernible pattern in Figure 3.8 since the real part of the eigenvalue is approximately zero. It is very sensitive, therefore, to the spatial resolution and temporal convergence of the mean flow in addition to the discretization of ω . Similar to the base flow case, the precise value of σ_1 is not of fundamental importance. Nevertheless, proportionality between the resolvent and stability mode shapes can be expected. Substituting $i\omega_s$ for λ into Equation 3.16, since the real part of the marginally stable mode is nearly zero, yields

$$\mathcal{H}(\omega_s) \approx \tilde{\mathbf{u}}_s \tilde{\mathbf{v}}_s^* \approx \hat{\boldsymbol{\psi}}_s \hat{\boldsymbol{\phi}}_s^*. \quad (3.34)$$

The sum in Equation 3.16 is dominated by the contribution from the marginally stable mode so the resolvent operator can be approximated by the outer product of the marginally stable mode and its adjoint counterpart or the optimal resolvent response and forcing modes at ω_s . Similar to the base flow case, the prediction from $|\hat{\phi}_1^* \hat{\psi}_1|^{-1} = 26.9$ agrees fairly well with the ratio of $\sigma_1(\omega_s) |i\omega_s - \lambda_s| = 28.9$ (see Table 3.1).

Rather than comparing the stability and resolvent mode shapes as shown in Figure 3.5, the contribution of the fluctuating feedback force to the wavemaker (see Meliga et al., 2016) for $Re = 100$ is computed using stability and resolvent modes, respectively, in Figure 3.9(a,b). Figure 3.9(a) is in good agreement with Meliga et al. (2016) and the agreement between Figures 3.9(a,b) is excellent, implying that the underlying modes are indeed proportional to each other. Streamlines from the mean flow are superimposed to observe how the wavemaker is related to the mean recirculation bubble, the size of which depends on Reynolds number. The length of the recirculation bubble scales with the streamwise extent of the wavemaker region for any Reynolds number for either the mean flow as seen in Figure 3.9(a,b) or base flow as seen in Figure 3.9(c,d). The wavemaker regions associated with the base profile are shown in Figures 3.9(c) and (d) for $Re = 47$ and $Re = 100$, respectively. Figures 3.9(a) and (d) compare the mean and base wavemakers at $Re = 100$, the main difference being that as the Reynolds number increases, both the mean recirculation bubble (Zielinska et al., 1997) and wavemaker region shrink. For the unstable base flows, increasing the Reynolds number will also increase the streamwise length of the recirculation bubble and wavemaker region as seen in Figures 3.9(c,d).

3.5.3 Resolvent modeling of (low Reynolds number) cylinder flow

In the case of flow around a circular cylinder, the resolvent identifies only stability mechanisms for both the base and mean flow cases. There are no frequencies at which amplification occurs due to pseudoresonance. This is consistent with the work of Abdessemed et al. (2009) who investigated direct transient growth analysis to study its role in the primary and secondary bifurcations of cylinder flow. Since only one mode becomes unstable, the effect of the nonlinearities is to saturate the growth mechanism and this alters the frequency of the structure (Barkley, 2006). While the base flow is unable to predict the nonlinear frequency in the saturated state, it can be correctly predicted from the mean flow. Thus in the case of stability mechanisms leading to unsteadiness (e.g. cylinder flow, Rayleigh-Bénard convec-

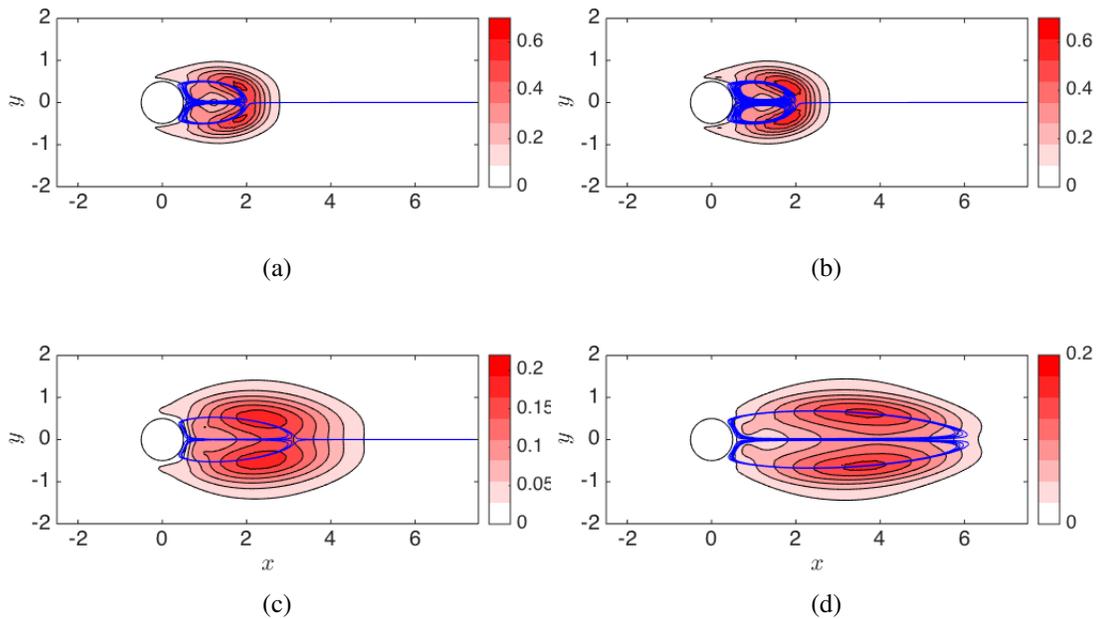


Figure 3.9: Wavemakers for mean flow at $Re = 100$ computed from stability modes (a) and resolvent modes (b). Wavemakers for the base flow at $Re = 47$ (c) and $Re = 100$ (d). Blue lines superimpose the mean flow streamlines which delineate the mean recirculation bubble.

tion (Turton et al., 2015)), mean stability analysis is successful at predicting the frequency of the unsteady flow.

3.6 Application to Wall-Bounded Turbulence

Unlike the cylinder flow, which is an oscillator with intrinsic dynamics that are insensitive to background noise, wall-bounded shear flows are an example of a noise-amplifier; as such, pseudoresonance plays a big role and leads to significant amplification at non-resonant frequencies. Due to its geometric simplicity, channel flow is chosen at $Re_\tau = 2000$ which has a parallel mean velocity profile $\bar{\mathbf{u}} = \bar{u}(y)$.

The off-diagonal term in \mathbf{L} is proportional to the mean shear \bar{u}' which is maximum at the wall. It remains large in the inner region before it begins to decline in the log region. Mean shear is the primary source of non-normality leading to significant amplification. Its spatial variation is important since it has been shown by McKeon and Sharma (2010) that a critical-layer mechanism tends to localize activity at the wall-normal location where the phase speed of the disturbance is equal to the local mean velocity. This is explored further by considering three particular wavenumber triplets that are representative of the near-wall cycle, a very large-scale motion

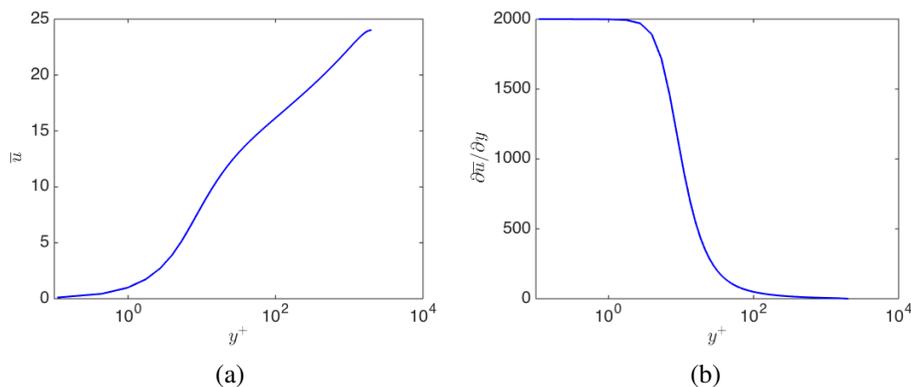


Figure 3.10: (a) Turbulent mean velocity profile for channel flow at $Re_\tau = 2000$ plotted in inner units alongside (b) the mean shear.

(VLSM), and a stationary disturbance. The roles of normal and non-normal mechanisms are studied by analyzing the mode shapes of the leading resolvent response modes, the pseudospectrum of the LNS operator, and the resolvent norm compared with the resonance curve.

3.6.1 Near-wall cycle

The first wavenumber combination considered is $(k_x, k_z, c^+) = (4\pi, 40\pi, 14)$ which is representative of the near-wall cycle (McKeon and Sharma, 2010). Here the wavespeed is given by $c^+ = \omega/k_x$. Figure 3.10 shows that the mean shear is very large at the wall-normal height where the wavespeed matches the local mean, resulting in the off-diagonal terms of the resolvent operator being large. This is similar to the model LNS operator in Equation 3.7 where the influence of non-normality concentrates energy in different velocity components for $\hat{\psi}_1$ and $\hat{\phi}_1$. The optimal resolvent forcing and response modes are plotted in Figure 3.11 to illustrate that the forcing is primarily concentrated in v and w while the response is mostly in u .

The strength of mean shear suggests that pseudoresonance is the primary driver of the near-wall cycle mode. The spectrum as well as contours of the pseudospectrum are plotted in Figure 3.12 for various ϵ . Figure 3.12 also includes the resolvent norm and the resonance curve. The ratio of the resolvent norm to the contribution from resonance is 19.6, which is of the same order of magnitude as the value predicted by $|\hat{\phi}_1^* \hat{\psi}_1|^{-1} = 4.81$ (see Table 3.1). Nevertheless, it is clear from this discrepancy that amplification cannot be attributed to one particular eigenvalue and that there is no eigenvector which is proportional to the resolvent mode. Using the expression

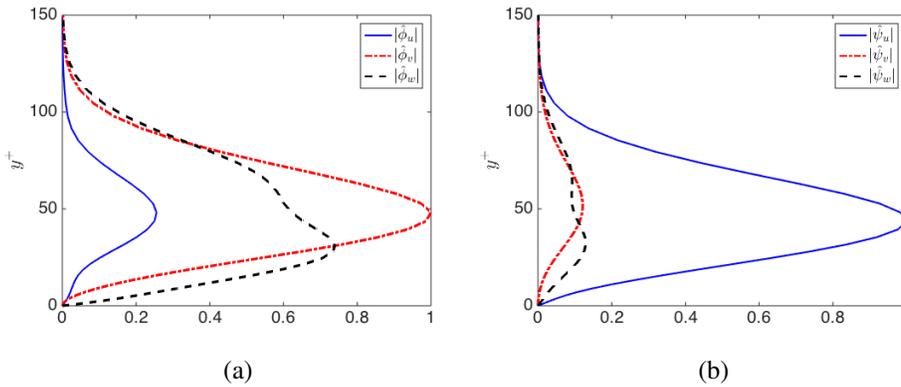


Figure 3.11: Velocity amplitudes for the optimal forcing mode $\hat{\phi}_1$ in (a) and optimal response mode $\hat{\psi}_1$ in (b) corresponding to the wavenumber triplet of $(k_x, k_z, c^+) = (4\pi, 40\pi, 14)$.

$\sigma_1|i\omega - \lambda|$ to quantify non-normality is problematic since there is no unique λ which is responsible for amplification.

The ω corresponding to $c^+ = 14$ is indicated by the horizontal, dashed blue line. At this frequency, the resolvent norm is significantly larger than the resonance term suggesting that amplification is due to non-normal mechanisms. This observation is confirmed by the spectrum where the least damped eigenvalues are clustered around higher frequencies and contribute to the resolvent norm for $\omega > 300$. It is also worth noting that the eigenvalues are significantly damped, which results in the leading singular values being on the order of unity. While these are not large, the rank-1 approximation is still valid since the first pair of singular values is approximately one order of magnitude larger than the others as seen in Figure 3.13.

3.6.2 VLMS

Further from the wall, the mean shear drops several orders of magnitude (see Figure 3.10) and the effect of viscosity decreases. As a result, amplification becomes a mix of both normal and non-normal effects. To reinforce this point, a wavenumber triplet which is representative of a VLMS is considered. The mode shapes for $(k_x, k_z, c^+) = (\pi/9, 2\pi/3, 22)$ are plotted in Figure 3.14. The forcing is dominated by the w -component while the response is dominated by the u -component. The v -component of the forcing, notably, is less significant than the near-wall cycle mode, implying that the role of lift-up is not as pronounced for this mode. The spectrum associated with the streamwise and spanwise wavenumbers of the VLMS is plotted along with the pseudospectrum of the LNS operator in Figure 3.15.

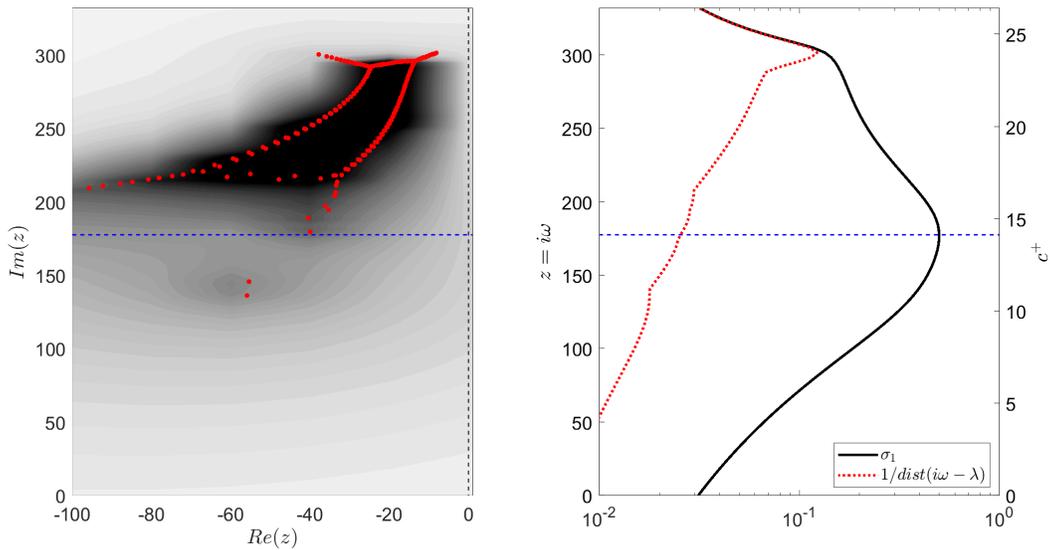


Figure 3.12: The eigenvalues of the operator $\mathbf{L}(k_x = 4\pi, k_z = 40\pi)$ in red circles overlaid with contours of the pseudospectrum (left). The resolvent norm is plotted in the solid black line along with the inverse distance from the imaginary axis to the nearest eigenvalue which is the red dotted line (right). The spatial wavenumbers correspond to the near-wall cycle, and the horizontal, dashed blue line represents the ω with the largest resolvent norm, which corresponds to a phase speed of $c^+ \approx 14$.

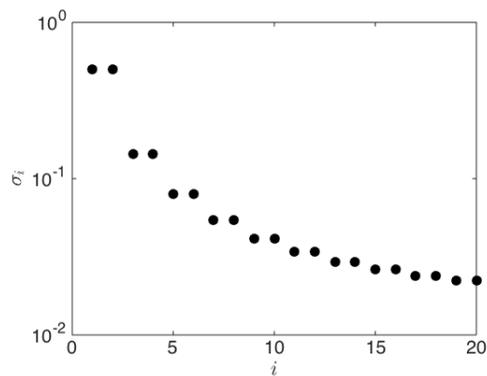


Figure 3.13: First 20 singular values of the resolvent operator for $(k_x, k_z, c^+) = (4\pi, 40\pi, 14)$.

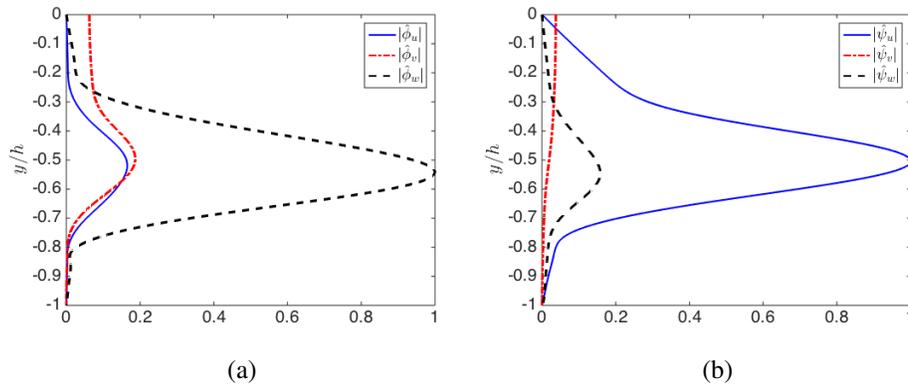


Figure 3.14: Velocity amplitudes for the optimal forcing mode $\hat{\phi}_1$ in (a) and optimal response mode $\hat{\psi}_1$ in (b) corresponding to the wavenumber triplet of $(k_x, k_z, c^+) = (\pi/9, 2\pi/3, 22)$.

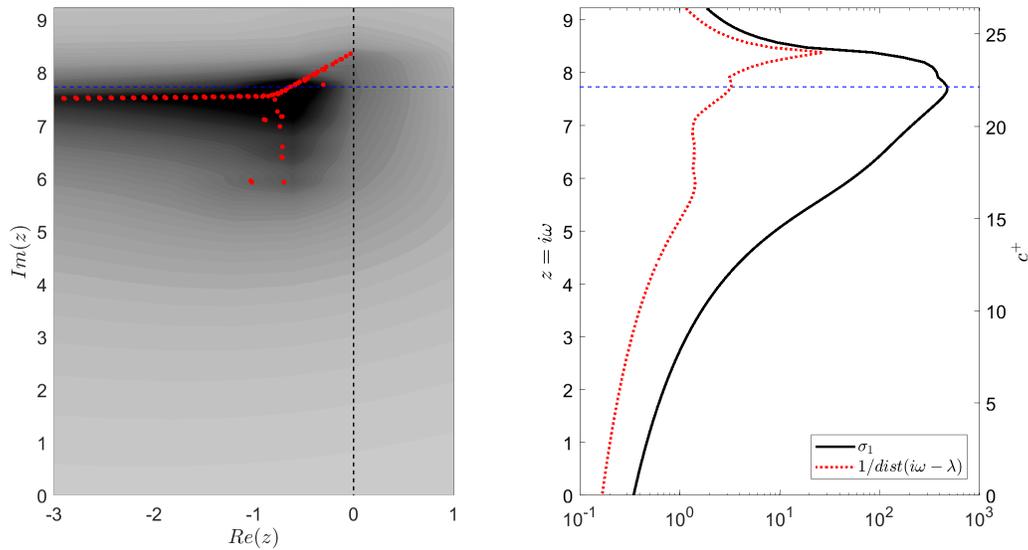


Figure 3.15: The eigenvalues of the operator $L(k_x = \pi/9, k_z = 2\pi/3)$ in red circles overlaid with contours of the pseudospectrum (left). The resolvent norm is plotted in the solid black line along with the inverse distance from the imaginary axis to the nearest eigenvalue which is the red dotted line (right). The spatial wavenumbers correspond to the VLSM mode, and the horizontal, dashed blue line represents the ω with the largest resolvent norm, which corresponds to a phase speed of $c^+ \approx 22$.

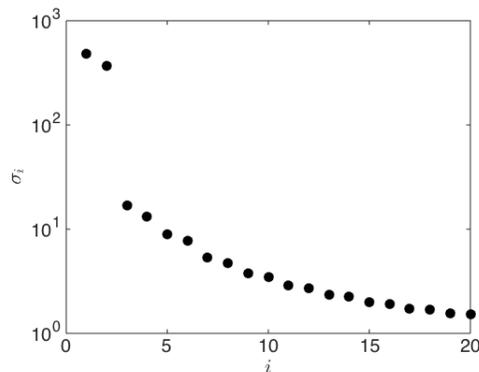


Figure 3.16: First 20 singular values of the resolvent operator for $(k_x, k_z, c^+) = (\pi/9, 2\pi/3, 22)$.

The results are drastically different from the near-wall cycle case as the resonant forcing of eigenvalues is greater than one so amplification is due to both terms on the right-hand side of Equation 3.2. Notably, the spectrum in Figure 3.15 resembles that of the base flow at $Re = 10,000$ for $k_x = 1$, $k_z = 0$ as there are three distinct branches. As observed by Reddy et al. (1993) and Schmid and Henningson (2001), the eigenvalues at the intersection of the branches are the most sensitive to perturbations and result in very large non-normal amplification. The product $\sigma_1 |i\omega - \lambda| = 146$ while $\|\hat{\phi}_1^* \hat{\psi}_1\|^{-1} = 34.9$ (see Table 3.1), suggesting that the nonorthogonality of many eigenfunctions leads to high pseudoresonance. Similar to the near-wall cycle case, there are no eigenvalues which exactly match the wave speed associated with the VLSM. The eigenvalue close to the dotted blue line in Figure 3.15, however, does seem to impact the resolvent norm which has an extra bump near its maximum value. This is also reflected in the red dotted line since the eigenvalue protrudes from the distinct Y-shape of the spectrum. The maximum singular values are on the order of 10^3 and the resolvent operator is low-rank as seen in Figure 3.16.

3.6.3 Stationary disturbances

Finally, stationary disturbances are considered as they tend to be the most amplified by the resolvent operator with singular values exceeding 10^4 . The specific wavenumber triplet selected for this study is $(k_x, k_z, \omega) = (0, 2\pi/3, 0)$. The roots behind such large amplification can be traced back to the model operator in Equation 3.7. In this example, the $\partial \bar{u} / \partial y \rightarrow \infty$ resulting in a low-rank system which concentrated all the forcing energy in the second velocity component and the re-

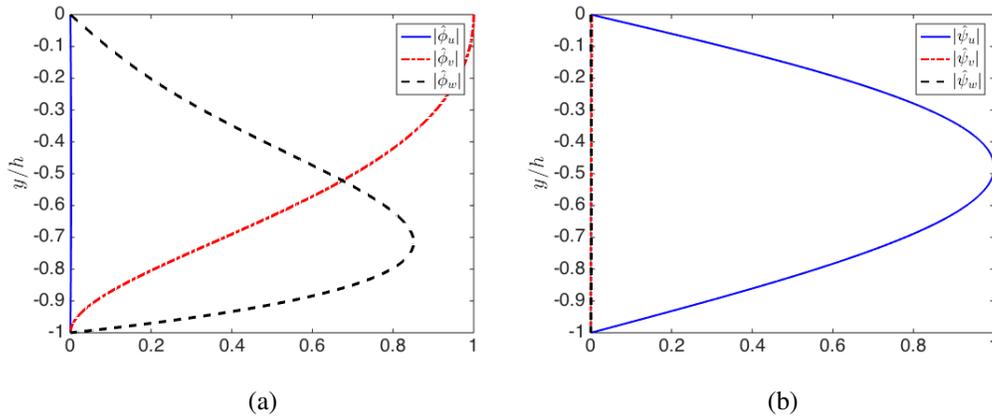


Figure 3.17: Velocity amplitudes for the optimal forcing mode $\hat{\phi}_1$ in (a) and optimal response mode $\hat{\psi}_1$ in (b) corresponding to the wavenumber triplet of $(k_x, k_z, \omega) = (0, 2\pi/3, 0)$.

sponse energy in the first velocity component. When $k_x = \omega = 0$, all of the diagonal terms of the resolvent become order ϵ small since imaginary terms are eliminated and \mathcal{D} scales with Re^{-1} . Thus, when the LNS operator is inverted, the determinant, which is the product of the diagonal terms, is very small. The energy for the forcing is almost totally in the wall-normal and spanwise components as seen in Figure 3.17 while the response is almost totally in the streamwise component.

Similar to the near-wall cycle and VLSM modes, the spectrum and contours of the pseudospectrum are presented in Figure 3.18 alongside the resolvent norm and resonance curve. All of the eigenvalues are real since the imaginary terms are eliminated from the resolvent operator when $k_x = 0$. Another implication of eliminating mean flow advection, as mentioned by Hack and Moin (2017), is that the Orr mechanism is absent from this mode.

Stationary disturbances are highly amplified and the singular values are plotted in Figure 3.19. The rank-1 approximation is quite applicable for this particular mode as the leading pair of singular values is on the order of 10^5 . Moreover, the contribution from non-normality is $|\hat{\phi}_1^* \hat{\psi}_1|^{-1} = 250$ which agrees quite well with $\sigma_1 |i\omega - \lambda| = 278$. Such agreement can be attributed to the eigenvalue closest to the imaginary axis which has an imaginary component that agrees with the most amplified frequency. Unlike the cylinder case where there exists a convective-type non-normality, the $k_x = 0$ modes are an example of the component-type non-normality and so $|\hat{\phi}_1^* \hat{\psi}_1|$ is small since the velocity for the forcing mode is almost completely

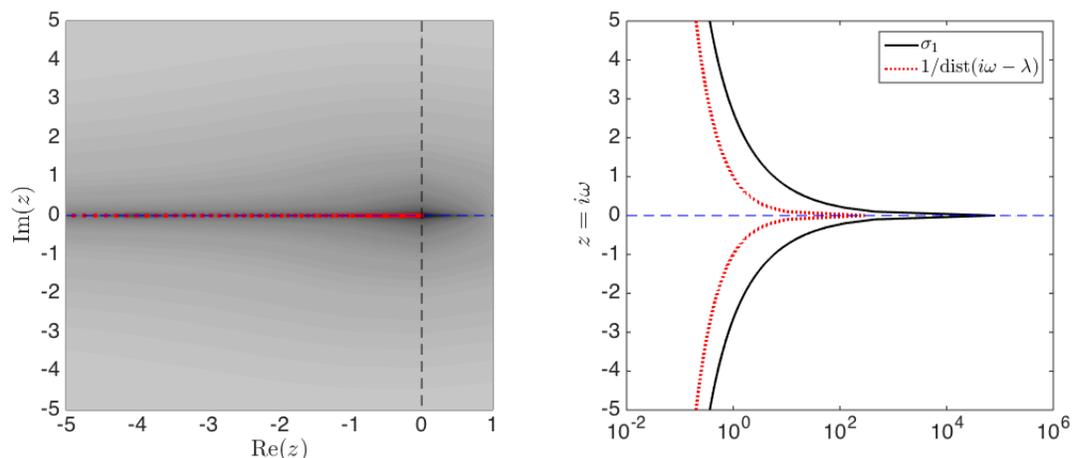


Figure 3.18: Spectrum and pseudospectrum of the stationary disturbance mode

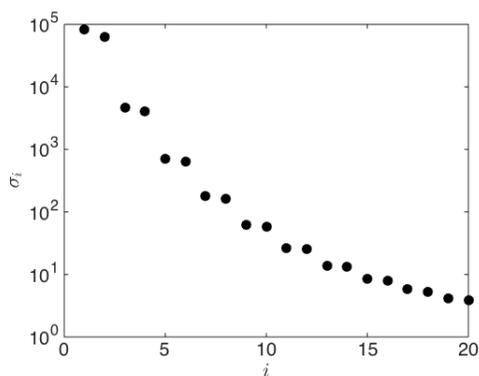


Figure 3.19: First 20 singular values of the resolvent operator for $(k_x, k_z, \omega) = (0, 2\pi/3, 0)$.

concentrated in the wall-normal plane while the velocity for the response mode is nearly all in the streamwise direction.

3.6.4 Influence of spatial wavenumber and wave speed

Based on the findings of this study and observations from Bourguignon (2012), it is possible to categorize amplification mechanisms in wall-bounded turbulence as either normal or non-normal depending on the wavenumber vector \mathbf{k} selected. When k_x is small, the influence of both mean flow advection and viscosity is diminished resulting in a non-normal system where there is high amplification. Low-order modes (those corresponding to the largest singular values), experience high amplification due to both normality and non-normality. Higher-order modes also experience amplification due to non-normality. At higher k_x , only low-order modes are

	\mathbf{k} or ω	σ_1	$ i\omega - \lambda ^{-1}$	$\sigma_1 i\omega - \lambda $	$ \hat{\phi}_1^* \hat{\psi}_1 ^{-1}$
Cyl. Base	$\omega = 0.743$	5.78e04	729	79.3	79.4
Cyl. Mean	$\omega = 1.02$	1.65e04	570	28.9	26.9
NWC	$\mathbf{k} = (4\pi, 40\pi, 14)$	0.502	2.56e-02	19.6	4.81
VLSM	$\mathbf{k} = (\pi/9, 2\pi/3, 22)$	479	3.28	146	34.9
Stationary	$\mathbf{k} = (0, 2\pi/3, 0)$	8.14e04	293	278	250

Table 3.1: Quantification of non-normality for the most amplified modes in cylinder and turbulent channel flows.

amplified as long as they are localized near the critical layer. Higher-order modes experience low amplification which is proportional to viscosity. The wall-normal height, furthermore, has implications on the type of amplification as it influences the choice of wave speed c , or ω , as well as the influence of mean shear $\partial\bar{u}/\partial y$. Closer to the wall, the mean shear is highest, while in the log region mean shear still plays an important role, resulting in preferential amplification of long streamwise structures.

3.7 Comparison of Flows

Cylinder flow is a case where resonance is the root of amplification leading to similarity of stability and resolvent modes. In both the base flow and mean flow cases, amplification occurs at a single frequency corresponding to the imaginary part of the least stable eigenvalue. The resolvent modes can be used to identify the wave-maker, which does not exist at very low Reynolds numbers when the flow is only convectively unstable. The cylinder exemplifies a convective-type non-normality where mean flow advection separates the spatial support of the forcing mode to be upstream of the response mode as long as $\bar{u} > 0$. Non-normality quantified by $|\hat{\phi}_1^* \hat{\psi}_1|^{-1}$ agrees well with $\sigma_1|i\omega - \lambda|$ for both the base and mean flows.

Only resonant mechanisms are active in cylinder flow whereas in wall-bounded

turbulence, both resonant and pseudo-resonant mechanisms are relevant. Three wavenumber triplets, representative of the near-wall cycle, VLSM's, and stationary disturbances, highlight the competing influences of viscous dissipation, mean flow advection, and mean shear on not only the most amplified modes but also the spectrum and pseudospectrum. The importance of each term depends significantly on the wall-normal height where the perturbations are localized. In the inner region where there is very high mean shear and viscosity is most important, amplification is primarily due to pseudo-resonant mechanisms. Forcing in the spanwise/wall-normal plane leads to a large response in the streamwise direction as seen for the near-wall cycle mode. The eigenvalues are highly damped yet the resolvent norm is on the order of unity due to the high sensitivity of the spectrum to perturbations. In the log region, mean shear and consequently lift-up are weaker yet the declining importance of viscosity results in eigenvalues which are closer to the imaginary axis. The most amplified disturbance for wavenumbers corresponding to the VLSM, nevertheless, is primarily due to pseudo-resonance. Consequently, there is poor agreement between non-normality quantified by $|\hat{\phi}_1^* \hat{\psi}_1|^{-1}$ and $\sigma_1 |i\omega - \lambda|$ as it is clear that amplification can no longer be attributed to a single eigenvalue. Mean flow advection results in the Orr mechanism (see McKeon, 2017) and hence less overlap between the forcing and response modes.

Stationary disturbances are the globally most amplified disturbances by effectively leveraging mean shear. The perturbation energy is almost exclusively concentrated in the v - and w -components of the forcing mode and in the u -component of the response mode. Assuming streamwise constant disturbances eliminates the mean flow advection term from the resolvent operator and hence suppresses the Orr mechanism. All of the non-normality, consequently, can be classified as a component-type non-normality, in contrast to the cylinder flow, and the eigenvalues of the LNS operator are real. Non-normality quantified by $|\hat{\phi}_1^* \hat{\psi}_1|^{-1}$ agrees well with $\sigma_1 |i\omega - \lambda|$ since amplification can be attributed to the eigenvalue closest to the imaginary axis. For a generic wavenumber triplet, the applicability of the rank-1 approximation can be approximated by the ratio of the largest term in the resolvent operator to the other entries. Finally, the distribution of energy among various velocity components may be useful when considering how the nonlinear term, which can be computed from resolvent response modes (McKeon et al., 2013), projects onto the optimal resolvent forcing modes.

3.8 Summary and Contribution

Stability and resolvent analysis have been juxtaposed to highlight the types of amplification mechanisms they are likely to identify. The two can be formally related through a dyad expansion of the resolvent operator. When the resolvent identifies eigenmodes as the most amplified disturbance, the forward eigenmodes are proportional to the resolvent response modes and the adjoint eigenmodes are proportional to the resolvent forcing modes which, consequently, contain sensitivity information. This formulation also elucidates how to interpret the real part of eigenvalues belonging to the mean LNS operator. It plays a role in the degree to which a disturbance is amplified by the resolvent and it separates eigenvalues in the spectrum leading to low-rank behavior. This is particularly evident in cylinder flow where the eigenvalue corresponding to the shedding mode is isolated from the other eigenvalues and is so close to the imaginary axis that it dominates the resolvent norm for a broad range of frequencies.

The contributions to the resolvent norm have been split into a resonance part and a non-normal part. The distance between the eigenvalue and a particular point on the imaginary axis quantifies the resonant contribution and its product with the resolvent norm is one scalar measure of non-normality. Non-normal amplification can also be computed by the inverse of the inner product between the most amplified resolvent forcing and response modes, i.e. $|\hat{\phi}_1^* \hat{\psi}_1|^{-1}$. If the two scalar measures of non-normality agree, it implies that amplification is due to a discrete eigenvalue and mean stability analysis is valid; otherwise, the amplification is due to pseudo-resonance and $|\hat{\phi}_1^* \hat{\psi}_1|^{-1}$ is the best measure of non-normality. This also suggests that to maximize the resolvent norm, it is desirable to minimize the overlap of the forcing and response modes, which biases the resolvent in favor of selecting amplification mechanisms which are as non-normal as possible.

Contours of the pseudo-spectrum have been overlaid with the spectrum for the mean LNS operator associated with cylinder and turbulent channel flow. It can be seen that the spectrum is very sensitive to perturbations in the latter flow due to high shear which suggests that the least stable eigenvalues of the mean operator viewed in isolation are less informative than they are for the two-dimensional case. Moreover, the mode shapes and characteristics of the pseudo-spectrum can be predicted from the mean profile or wavenumber triad selected as they have implications on which terms in the LNS operator are most important.

Chapter 4

SCALING CHARACTERISTICS AND NONLINEAR FORCING OF CYLINDER FLOW

In this chapter, the resolvent analysis of the mean cylinder wake is examined in greater detail. It was observed in the previous chapter that the resolvent operator is very low-rank at temporal frequencies near the shedding frequency. It is revealed here that the resolvent always one dominant structure even when the flow is three-dimensional. As such, resolvent analysis can be considered as a useful tool in the context of data-assimilation since it is able to identify the energetic structures that interact nonlinearly to produce the necessary Reynolds stresses that sustain the mean profile. Additionally it contains information about the scaling of the mean profile and how the dominant amplifications are sustained by nonlinear interactions. In the context of reduced-order modeling, resolvent analysis assists with the placement of flow sensors which are most helpful when placed in regions of the flow where the fluctuations are energetic.

The implications of this case study go beyond its utility in flow reconstruction and estimation. The optimal resolvent response modes are shown to agree closely with DMD modes which are computed from snapshots of the flow. For frequencies where the resolvent operator is not low-rank, the most amplified modes predicted from resolvent analysis do not match the DMD modes. These observations are similar to those made by Towne et al. (2018) for SPOD modes. Consequently, it is not efficient to consider triadic interactions of resolvent modes outside the range where the rank-1 approximation is valid. Instead, it is possible to predict higher frequency harmonics by computing directly the nonlinear forcing $\mathbf{u} \cdot \nabla \mathbf{u}$ and using it as an input to the resolvent operator. This can be achieved by computing the triadic interactions of resolvent modes where the rank-1 approximation is valid. To demonstrate these ideas, the chapter focuses primarily on the $Re = 100$ case although other Reynolds numbers are considered at the beginning when analyzing the scaling behavior of the resolvent modes and mean profile.

4.1 Scaling of the Mean Profile and Resolvent Norm

The peak of the resolvent norm for the mean wake around a circular cylinder occurs at the shedding frequency ω_s , as seen in Figure 4.1, where the leading singular

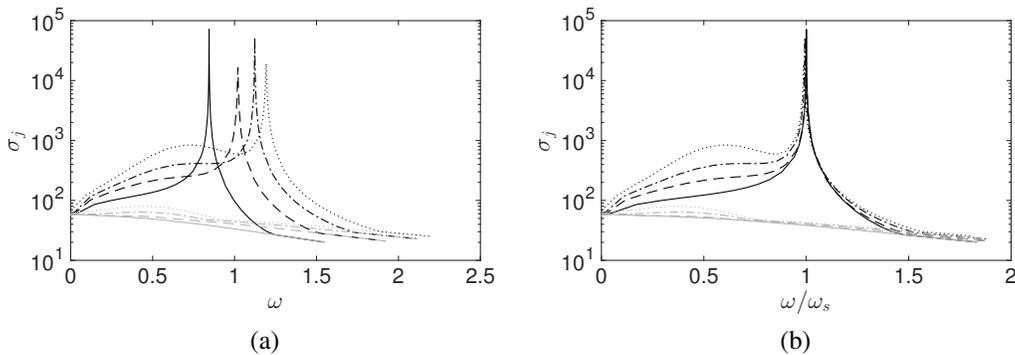


Figure 4.1: The first singular value σ_1 (black) and second singular value σ_2 (grey) for $Re = 60$ (—), $Re = 100$ (---), $Re = 140$ (- · -), and $Re = 180$ (···). Singular values are plotted versus ω in (a) and ω normalized by ω_s in (b).

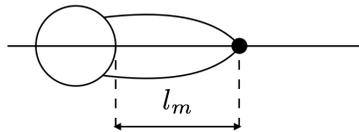


Figure 4.2: Schematic of the cylinder where the recirculation length l_m is defined as the distance between the edge of the cylinder and the end of the recirculation bubble.

values for several Reynolds numbers are plotted. The location of the peak in the resolvent norm collapses in Figure 4.1 when the frequency axis is scaled by ω_s . A scaling for the peak resolvent norm as a function of the Reynolds number cannot be determined since the real part of the eigenvalue is nearly zero; consequently, the resolvent norm is very sensitive to the spatial resolution and temporal convergence of the mean flow.

The mean profile $\bar{u}(x, y)$ scales with $St = \omega_s D / U_\infty$ at the recirculation point x_r , or the streamwise location where the mean flow along the centerline switches direction from negative to positive, implying $\bar{u}(x = x_r, y = 0) = 0$. l_m is the recirculation length and is defined as the distance from the edge of the cylinder to the recirculation point (see schematic in Figure 4.2). The mean profiles for the Reynolds numbers considered earlier are plotted in Figure 4.3(a-d) with and without the y -coordinate scaled by St . The profiles for \bar{u} in panels (a-b) collapse almost perfectly as does the shape of \bar{v} in panels (c-d) even though the peak velocity for \bar{v} depends on Reynolds number.

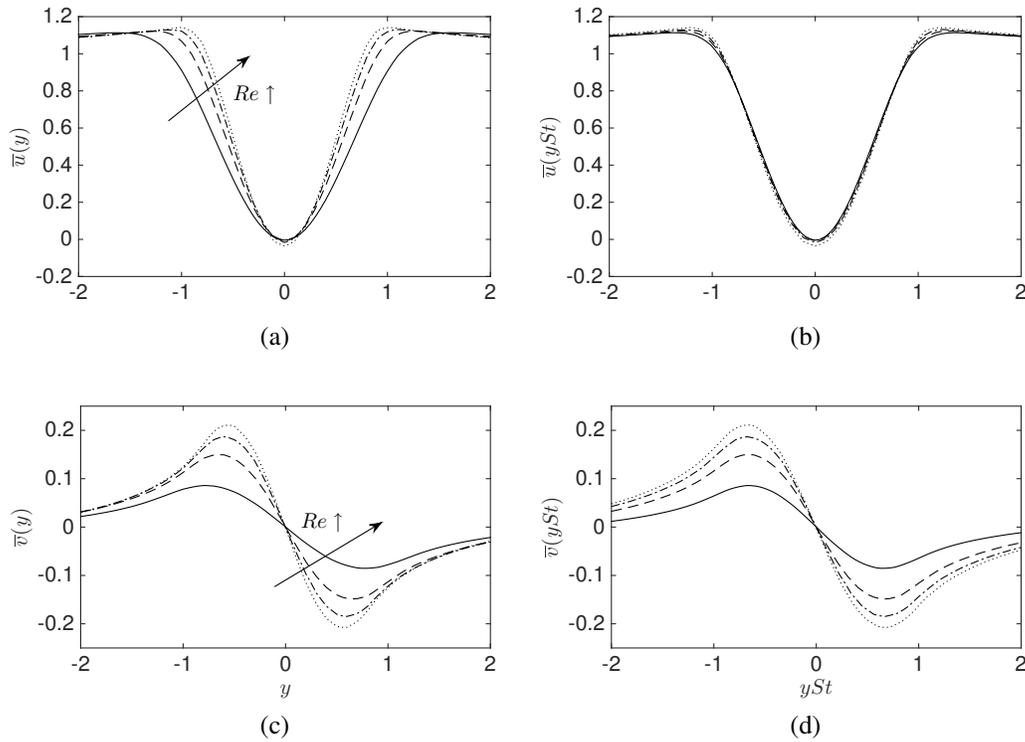


Figure 4.3: Mean profiles \bar{u} (a) and \bar{v} (b) at the recirculation point for $Re = 60$ (—), $Re = 100$ (---), $Re = 140$ (- · -), and $Re = 180$ (···). The y -axis is rescaled by St in (c) and (d) for \bar{u} and \bar{v} , respectively.

The mean profiles for a square cylinder and the three-dimensional cylinder wake are also computed to see whether or not the same scaling applied. The results for \bar{u} only are plotted in Figure 4.4(a-b) and follow the same trends as the two-dimensional cylinder wake. Similar scaling of the location of the peak resolvent norm is observed for these flows as for the two-dimensional wake in Figure 4.1, underscoring the importance of the similarity characteristics of the mean flow for the response modes. These have not been plotted in the interest of brevity.

It is interesting to consider how the profile for the base flow \mathbf{U}_0 for the circular cylinder scales at the recirculation point. It is well known that the base flow cannot predict the frequency of the vortex shedding (Barkley, 2006). The profiles for $U_0(x = x_r, y)$ scaled by St , where the frequency is equal to that of the least stable eigenvalue, are plotted in Figure 4.5. Their shapes are not quite identical since the growth rate of the least stable eigenvalue increases with Reynolds number although the width of the profile remains roughly constant. Barkley (2006) showed that the frequency predicted by the base flow is almost constant above the critical Reynolds

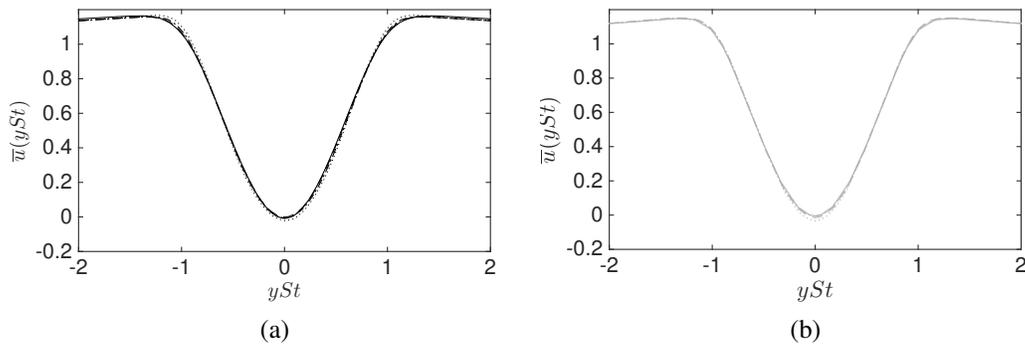


Figure 4.4: Rescaled mean profiles \bar{u} at the recirculation point for the square cylinder (a) and three-dimensional cylinder wake (b). The Reynolds numbers are $Re = 60$ (—), $Re = 90$ (---), $Re = 120$ (-·-), and $Re = 150$ (···) (black) for the square cylinder and $Re = 200$ (—), $Re = 230$ (---), $Re = 260$ (-·-), and $Re = 290$ (···) (grey) for the three-dimensional cylinder wake.

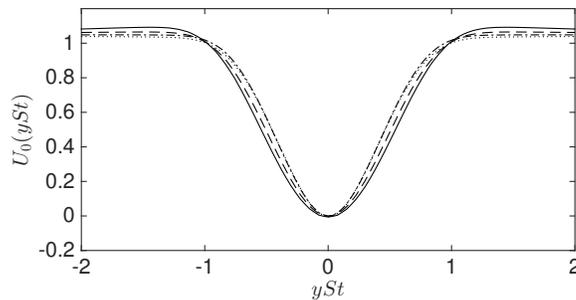


Figure 4.5: Base flow profiles U_0 at recirculation point for circular cylinder $Re = 60$ (—), $Re = 100$ (---), $Re = 140$ (-·-), and $Re = 180$ (···).

number, so the profiles in Figure 4.5 have been rescaled by nearly the same St . It can be concluded that the frequency predicted from stability analysis of the unstable base flow is determined by the width of the profile at the recirculation point, which remains roughly constant for $Re > Re_c$. The main change that occurs when increasing the Reynolds number is an elongation of the mean recirculation bubble which scales linearly with Reynolds number (Zielinska et al., 1997).

4.2 Scaling and Convection Velocity of Resolvent Modes

The signature of ω_s can also be observed in the optimal resolvent forcing and response modes. The scaling of the former is simpler and is addressed first. The amplitude of the forcing modes along the centerline of the domain for $Re = 60$ and $Re = 180$ is plotted in Figure 4.6(a). The spatial support is limited to upstream

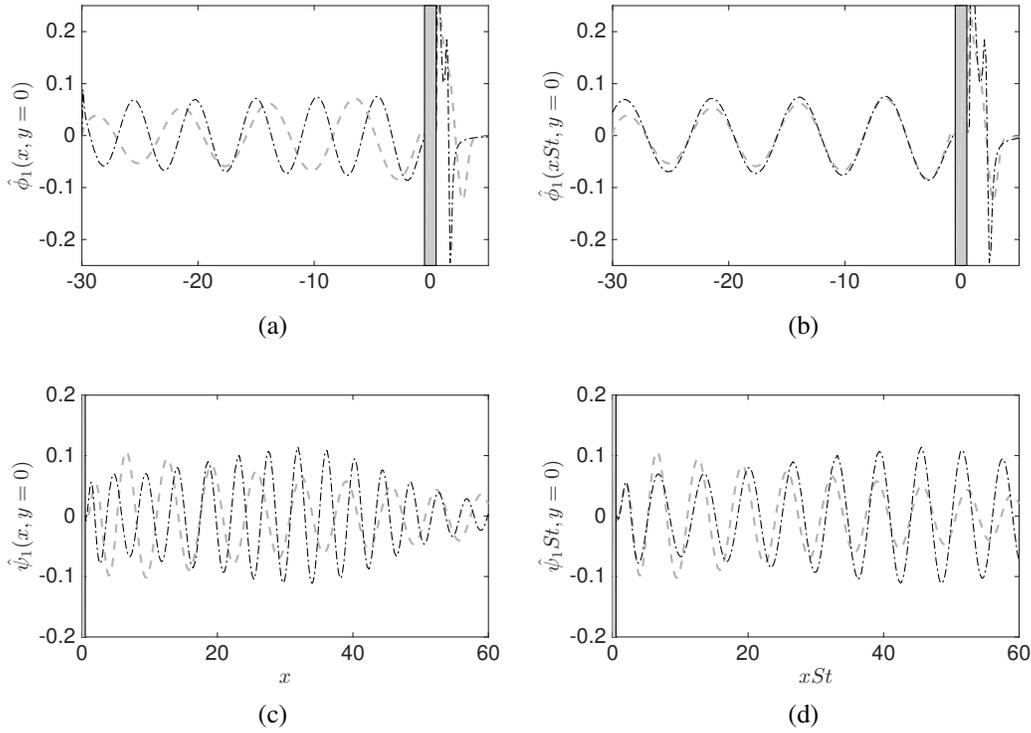


Figure 4.6: Profiles of the optimal response mode $\hat{\psi}_1$ and forcing mode $\hat{\phi}_1$ in (a) and (c), respectively, along the centerline $y = 0$. The modes are rescaled by St in (b) and (d). The grey, dashed line is $Re = 60$ while the black, dashed-dotted line is $Re = 180$. The cylinder location is delineated by the shaded rectangle.

of the cylinder as well as the region immediately behind it. A convection velocity, defined as

$$U_c = \frac{\omega_s}{k(x)} = \frac{q(x)}{2\pi} \omega_s, \quad (4.1)$$

where $q(x)$ is the spatial wavelength, is computed for the forcing mode in Figure 4.6(a) sufficiently upstream of the cylinder where it is constant, i.e. $q \neq q(x)$. For $Re = 60$, $q = 7.44$ and $\omega_s = 0.844$ yielding a convection velocity of $U_c = 1$. In fact, the convection velocity for all Reynolds numbers is equal to unity for the forcing mode sufficiently upstream of the cylinder and this is compatible with the fact that $\bar{u}(x, y = 0) = 1$ in this region. Consequently, an adjoint perturbation is transported upstream at a convection velocity equal to the local mean velocity. The forcing modes along the centerline for the circular cylinder, where the mean velocity is unity upstream of the cylinder, collapse when scaled by St in the streamwise direction as seen in Figure 4.6(b).

The scaling of the resolvent response modes is less straightforward due to the com-

plexity of the mean profile behind the cylinder. The analogous response modes for $Re = 60$ and $Re = 180$ are plotted in Figure 4.6(c). Even though they scale with St in Figure 4.6(d) for a couple of periods, the profiles diverge further downstream. This can be attributed to the fact that U_c is not constant behind the cylinder and this is apparent in Figure 4.7(a) where U_c has been plotted for several Reynolds numbers. As observed by Thompson et al. (2014), the vortex street transitions from spaced-out vortices to a bunched-up street with more diffused vortices at a stream-wise location of approximately $x = 10$. This is less pronounced for the lower Reynolds number cases as U_c does not decrease once it reaches a maximum while it does for the higher Reynolds numbers. Nevertheless, the maximum value of U_c is constant for all Reynolds numbers.

U_c is also calculated by tracking the local minima and maxima of the fluctuating v field in the DNS snapshots. The minima correspond to the edges of the vortices and these results match the values obtained using the mode shape. It can be concluded that the most amplified global resolvent mode contains almost all of the dynamics of the flow. It can be expected that the decrease in U_c somehow manifests itself in the mean profile. Extracting the mean velocity along the centerline as seen in Figure 4.7(b) illustrates that even though the maximum velocity is approximately constant with Reynolds number, there is a noticeable decline for higher Reynolds numbers downstream of the cylinder. Similar to U_c , the mean velocity along the centerline remains constant for the $Re = 60$ and $Re = 100$ cases but declines for higher Reynolds numbers, particularly for $Re = 180$. It should be noted that the mean centerline velocity is not equivalent to the convection velocity of the vortices. Lin and Hsieh (2003) determined that the path of the vortex cores more closely follows $y = b/2$ where b is the half-velocity-defect. The centerline mean velocity still represents the trends in U_c and is consistent with the results from Cimbala et al. (1988), Williamson and Prasad (1993) and Thompson et al. (2014).

The low-rank nature of the cylinder flow results in a close relationship between the mean profile and the most amplified resolvent mode. Furthermore, it has been shown that ω_s plays a significant role in their scaling. Less attention has been given to the scaling of the resolvent modes in the region immediately behind the cylinder as well as the role of the recirculation length l_m . These topics are discussed in the next section.

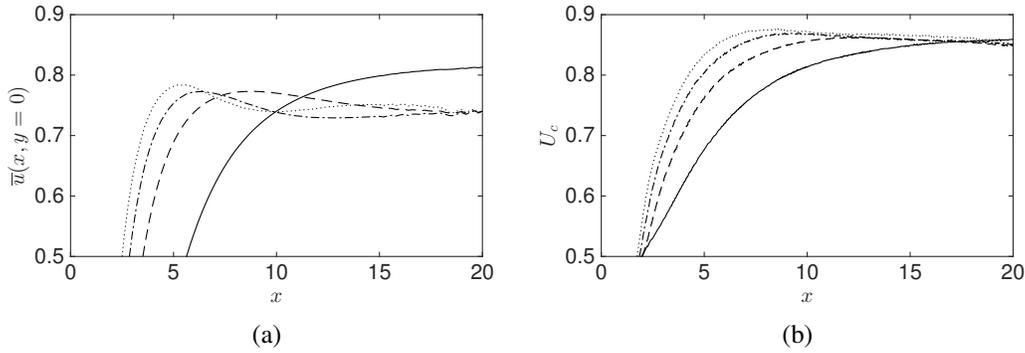


Figure 4.7: Mean profile along centerline (a) and convection velocities of the vortices (b) for $Re = 60$ (—), $Re = 100$ (---), $Re = 140$ (- · -), and $Re = 180$ (···).

4.3 The Wavemaker and Recirculation Length

Figure 4.7(a) shows that U_c increases behind the cylinder before it either reaches a constant value in the low Reynolds number cases or attains a maximum before decreasing in the high Reynolds number cases. This implies that the vortices are accelerating after forming behind the cylinder. It also means that the wavelength of the resolvent response mode is constantly increasing as a function of streamwise distance. The initial wavelength is proportional to the recirculation length as seen in Figure 11 of Balachandar et al. (1997), or the size of the vortex immediately before it is shed from the cylinder. Since there are two vortices of opposite sign per wavelength of the response mode, the initial wavelength is thus $2l_m$. The convection velocity at the recirculation point, consequently, is Reynolds number dependent and is proportional to $l_m\omega_s$.

To achieve the convection velocity at x_r , the vortex needs to accelerate over a distance l_m from an initial speed of zero. Assuming the acceleration a is constant, which is a reasonable assumption since the slope of U_c in Figure 4.7(a) is also approximately constant at the beginning of the vortices' trajectory, signifies that simple kinematics can be used to determine acceleration:

$$U_c^2 = U_i^2 + 2al_m, \quad (4.2)$$

where U_i is the initial velocity of the vortex and is set to zero. The final value of U_c is proportional to $l_m\omega_s$ so acceleration scales in the following manner:

$$a \sim \omega^2 l_m. \quad (4.3)$$

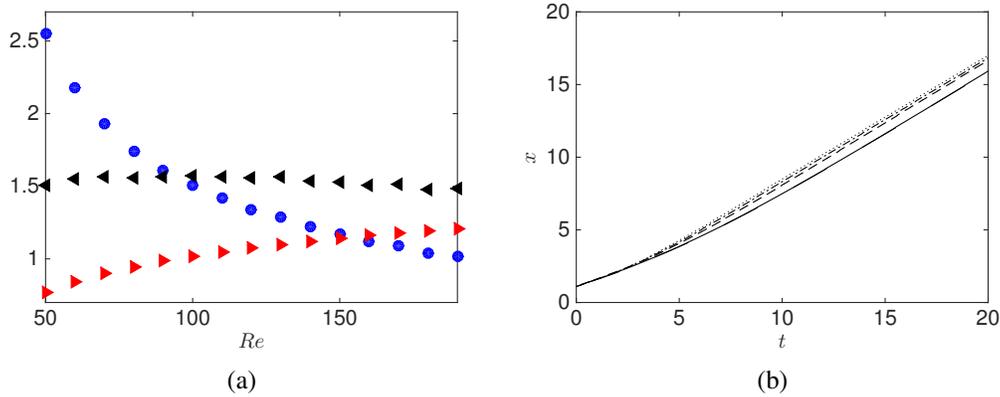


Figure 4.8: Comparison of the recirculation bubble length l_m (blue circles), shedding frequency ω_s (red right-facing triangles) and $\omega_s^2 l_m$ (black left-facing triangles) in (a). Position of the vortex cores as a function of time in (b) for $Re = 60$ (—), $Re = 100$ (---), $Re = 140$ (- ·), and $Re = 180$ (· · ·).

The value of the acceleration in (4.3) is the acceleration rescaled such that it is independent of Reynolds number. For larger Reynolds numbers where the vortices are smaller, they need to accelerate more quickly since the recirculation length is much shorter.

The time scale associated with the vortices is ω_s while the length scale is q . Ignoring the fact that smaller vortices eventually slow down due to the decay of the von Kármán vortex street, the convection velocity is constant across Reynolds numbers sufficiently downstream of the cylinder. The time scale associated with the acceleration, on the other hand, is l_m , or q immediately behind the cylinder while the time scale is still ω_s as there are no other time scales in the flow. The maximum convection velocity given by $q\omega_s$ is constant across all Reynolds numbers while the acceleration $l_m\omega_s^2$ is also constant. Consequently, there is a relationship between a mean quantity (l_m) and a frequency of fluctuation (ω_s). The respective values for l_m , ω_s , and $l_m\omega_s^2$ are plotted in Figure 4.8(a) to show $l_m\omega_s^2$ is constant across Reynolds number as well as the trajectories of the vortices in Figure 4.8(b) to show that the initial acceleration and final convection velocity are equal across Reynolds numbers. The trajectories are not identical since the larger vortices at lower Reynolds numbers take longer to reach their maximum speed.

The fact that l_m is directly related to the frequency of the unsteady fluctuations supports the notion that the wavemaker \mathcal{W} , which can be approximated by

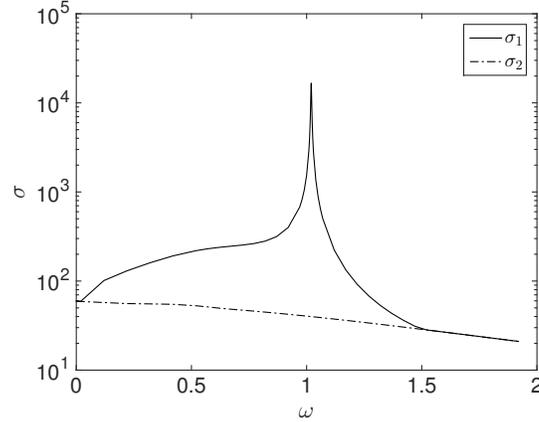


Figure 4.9: The first two singular values plotted as a function of ω for the mean cylinder flow at a Reynolds number of $Re = 100$.

$$\mathcal{W}(\mathbf{x}_0) \approx \|\hat{\psi}_1(\mathbf{x}_0, \omega_s)\| \|\hat{\phi}_1(\mathbf{x}_0, \omega_s)\|, \quad (4.4)$$

where $\|(\cdot)\|$ denotes the magnitude (Symon et al., 2018), represents the region of instability in the flow, e.g. Giannetti and Luchini (2007). A physical analogy can be noted between the bluff body flows and the simple pendulum. When visualizing contours of the streamwise velocity, there is a region of fluid, the length of which is approximately l_m , with low streamwise velocity directly behind the cylinder swinging up and down at a frequency of ω_s . This is similar to a simple pendulum whose radial frequency Ω depends solely on its length L (i.e. $\Omega = \sqrt{g/L}$). Since g is a constant representing acceleration due to gravity, it can be shown that $\Omega^2 L$ is constant which is analogous to $\omega_s^2 l_m$ for the cylinder flow.

4.4 Singular Values and Mode Branches

For the remainder of the chapter, the discussion will focus on the $Re = 100$ case. Figure 3.7(b) is reproduced below in Figure 4.9 with the axes switched such that the trends for the first two singular values are easier to identify. As seen in Chapter 3, the behavior of σ_1 near the shedding frequency $\omega_s = 1.02$ is primarily associated with the marginally stable eigenvalue. The behavior of σ_2 , on the other hand, has little to no variation in ω , suggesting it corresponds to a different amplification mechanism which is not captured by the eigenvalues. For frequencies where $\omega > 1.49\omega_s$, this physical mechanism becomes the most amplified and there is no longer any separation between the singular values.

The mode shapes $\hat{\psi}_1$ and $\hat{\psi}_2$ are plotted in Figure 4.10 for $\omega \in [0.5\omega_s, 2.0\omega_s]$. Two

classes of modes can be identified: a wake mode which resembles the vortex shedding and a critical-layer mode whose spatial support is in the free-stream region above and below the wake. These are referred to as mode W and mode C, respectively, to avoid confusion with the three-dimensional modes mode A and mode B. In Figure 4.9, it can be noted that mode W is dominant when $\sigma_1 \gg \sigma_2$, or when the resolvent is low-rank and there is a separation of singular values. At a frequency of $\omega = 1.49\omega_s$, Mode C becomes dominant and the resolvent operator is no longer low-rank. One can interpret σ_2 for $0.55\omega_s < \omega < 1.49\omega_s$ and σ_1 for $\omega > 1.49\omega_s$ as a mode ‘branch’ which is continuous. The most amplified mode shape switches at $\omega = 1.49\omega_s$. Mode branch W is demarcated by σ_1 until $\omega = 1.49\omega_s$, at which point it no longer appears in Figure 4.9. Branch W still exists, but it falls off so dramatically that for $\omega = 1.5\omega_s$, it is the fifth most amplified structure and hence would appear as σ_5 had more singular values been plotted. Dergham et al. (2013) observed similar behavior for the flow over a two-dimensional backward facing step at $Re = 600$. They found several different branches and categorized them as shear layer dynamics, which dominate at lower frequencies, and free-stream dynamics, which dominate at higher frequencies.

Figure 4.9 suggests that the resolvent will struggle to identify coherent structures at harmonics of the shedding frequency since there is virtually no separation of singular values. The nonlinear forcing to the resolvent is likely to be biased in a manner such that it excites one particular structure at $2\omega_s$ which is not one of the most amplified resolvent modes. It also appears that the physical signature of the shedding mode is significant in the mean profile. To explore this further, it can be seen from Figure 4.11 that $\hat{\psi}_1$ for suboptimal frequencies is a stretched or shrunk version of the shedding mode when $\omega < 1.49\omega_s$. Thus, one can surmise that the convection velocity of the vortices $U_c = \omega/q(x)$, where $q(x)$ is the spatial wavelength of the mode as a function of x , is constant for each ω . In Figure 4.11, the profile of the modes along the centerline is plotted as a function of x . When x is rescaled by the factor ω/ω_s , the mode shape collapse, implying that the convection velocity is the same as a function of x .

The convection velocity can also be measured for the C branch modes. Unlike the W branch modes, they have a roughly constant convection velocity equal to unity or the free-stream velocity. Based on this finding, it is possible to predict the structure of the mode at $3\omega_s$; it would resemble the $2\omega_s$ mode but the spatial wavelength would be a factor of 50% lower, resulting in a denser wavepacket in the

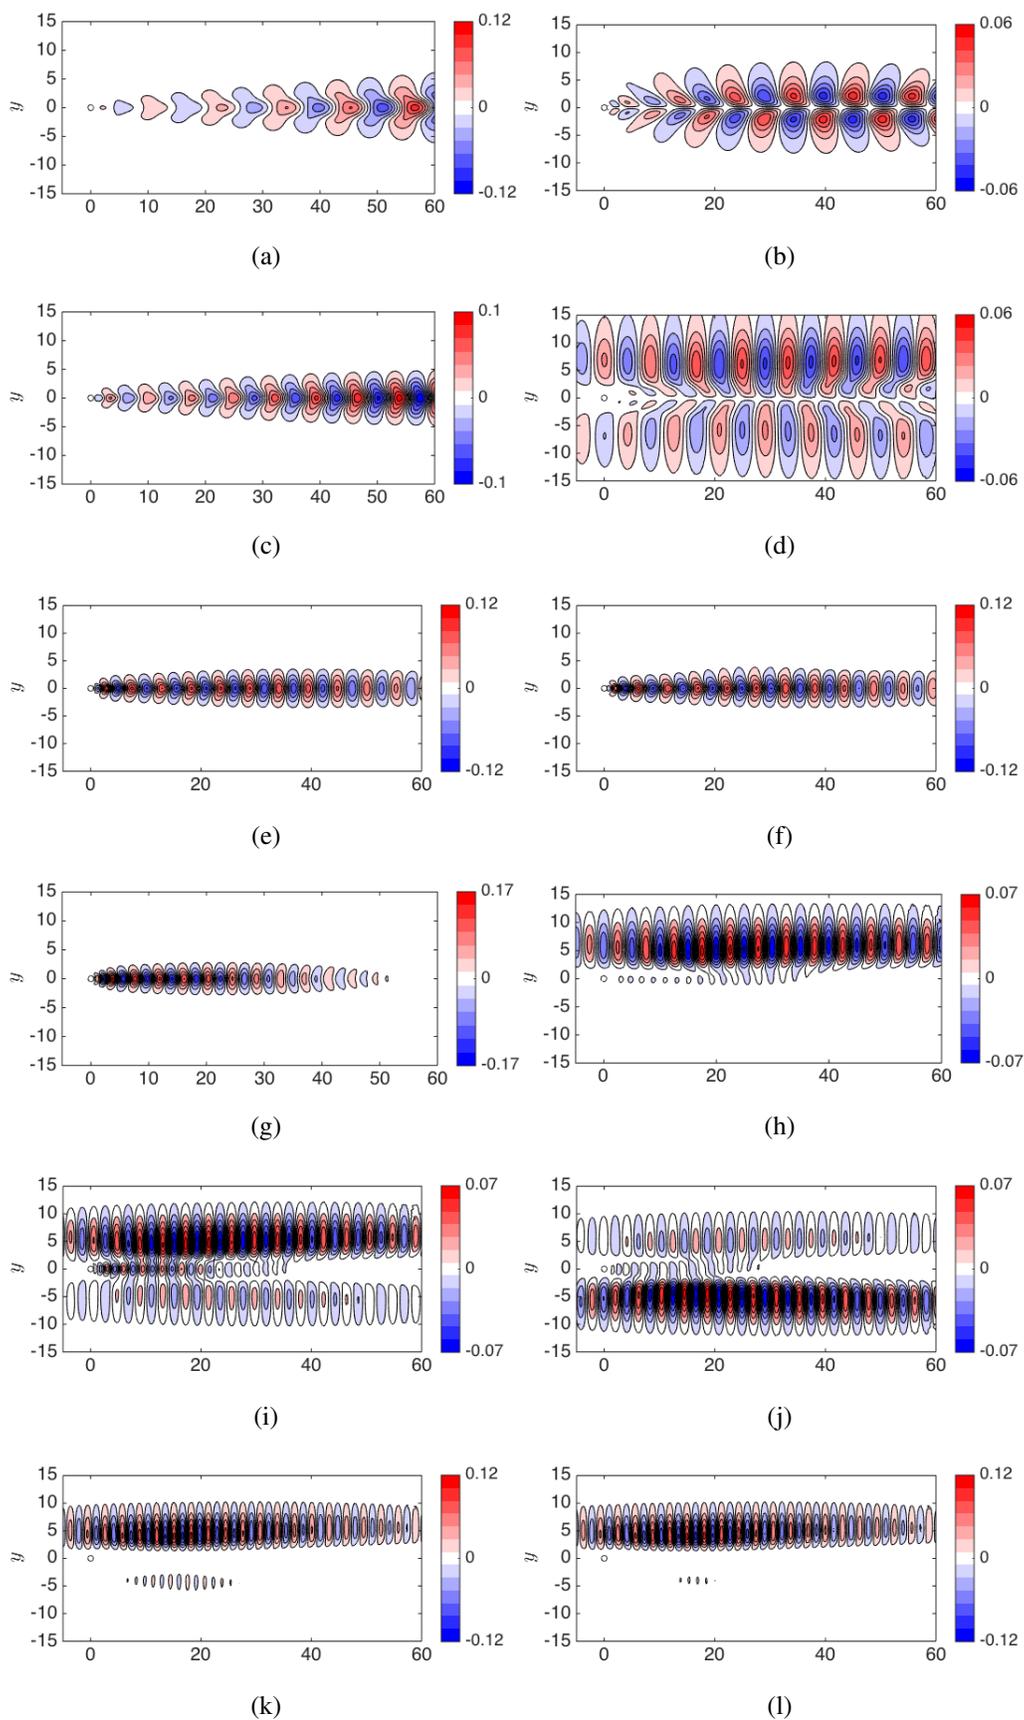


Figure 4.10: First two resolvent response modes $\hat{\psi}_1$ (left) and $\hat{\psi}_2$ (right) for $\omega = 0.5\omega_s$ (a,b), $\omega = 0.75\omega_s$ (c,d), $\omega = 1.0\omega_s$ (e,f), $\omega = 1.25\omega_s$ (g,h), $\omega = 1.5\omega_s$ (i,j), and $\omega = 2\omega_s$ (k,l).

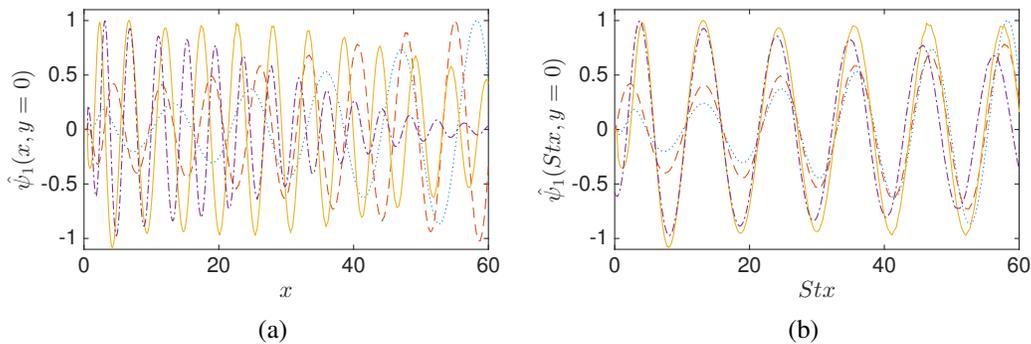


Figure 4.11: Unscaled (a) and scaled (b) profiles along the centerline of the resolvent response modes for the $Re = 100$ mean cylinder wake. The frequencies plotted are $\omega = 0.5\omega_s$ (\cdots), $\omega = 0.75\omega_s$ ($---$), $\omega = 1.00\omega_s$ ($-$), and $\omega = 1.25\omega_s$ ($- \cdot -$).

free-stream.

4.5 Comparing Resolvent and DMD Modes

Figure 4.12 compares the resolvent response modes to the DMD modes at ω_s , the globally most amplified frequency, and its first harmonic $2\omega_s$, where the resolvent operator is not low-rank. There is very good agreement for ω_s although they are not exactly the same as seen in Figure 4.13, where the difference between the two is plotted. The amplitude of the DMD mode decays in the streamwise direction more quickly than it does for the resolvent mode. Another difference is that the spatial support of the u -component is slightly further from the centerline in the DMD modes. The v -component, on the other hand, does not suffer from this problem since it is symmetric across the centerline.

There is no agreement between the two sets of modes when $\omega = 2\omega_s$ as seen in Figures 4.12 and 4.13. The resolvent mode predicts a structure in the free-stream with a convection velocity which is significantly higher than that of the DMD mode. It is incorrect to say that resolvent analysis does not work for this frequency. Instead, it suggests that there is no linear mechanism leading to significant amplification. Assuming a white-in-time stochastic forcing will not reproduce the modes observed in the flow, which is hardly surprising since the operator is anything but low-rank. Similar observations have been made by Zare et al. (2017) in the context of turbulent, wall-bounded shear flows and Towne et al. (2018) and Schmidt et al. (2017a) in the context of turbulent jets.

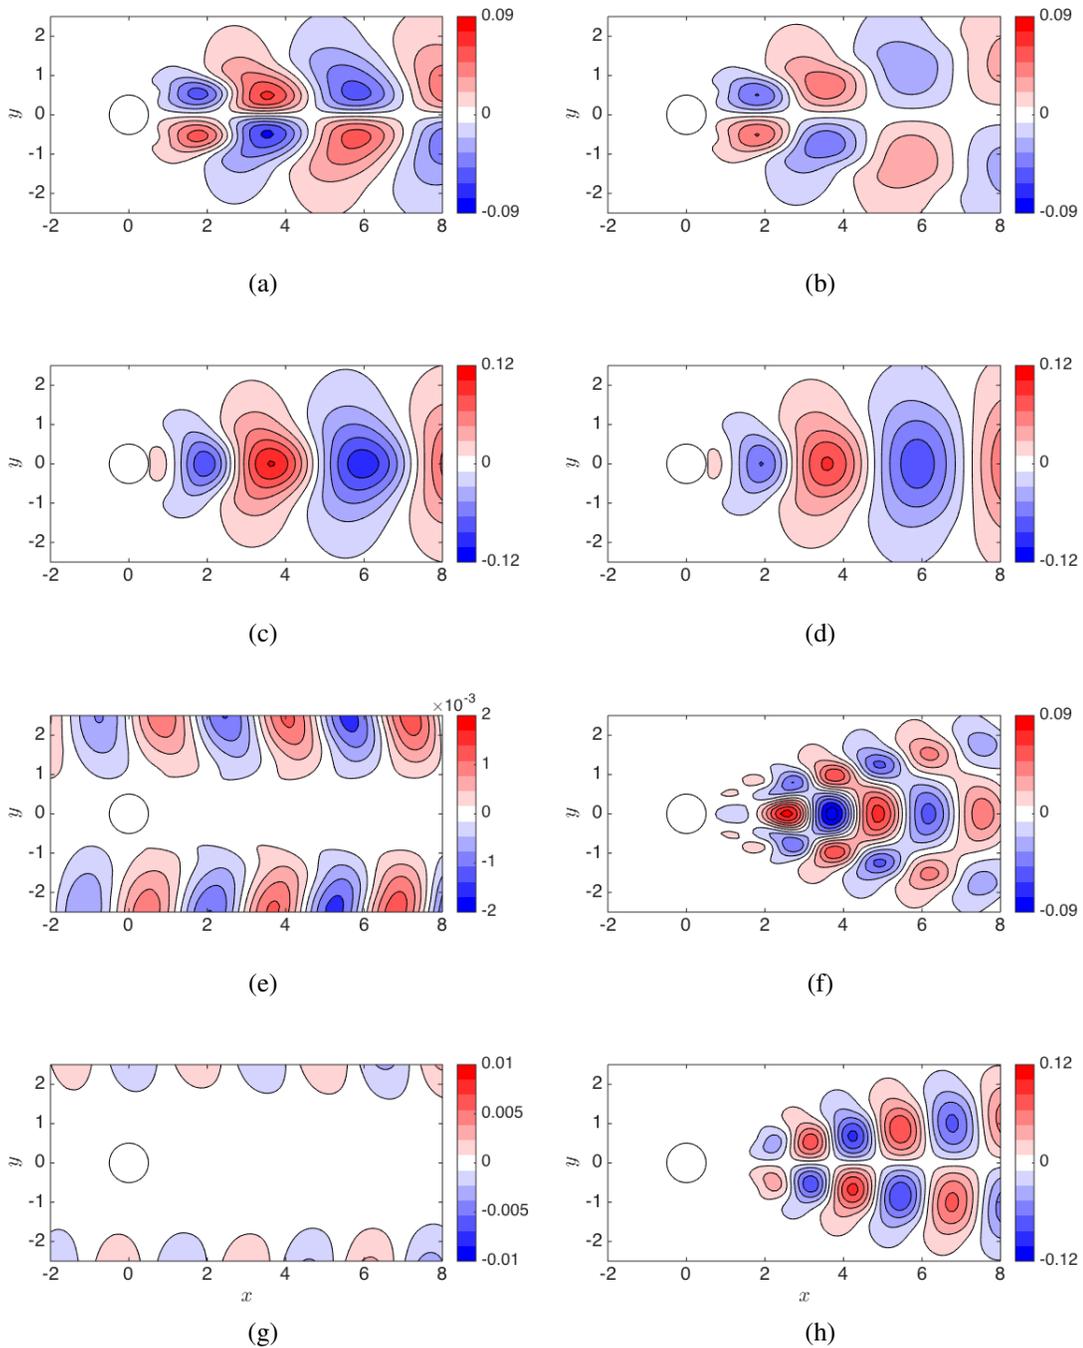


Figure 4.12: Comparison of resolvent response (left) and DMD (right) modes for $\omega = \omega_s$ in (a-d) and $\omega = 2\omega_s$ in (e-h). The streamwise velocity component is plotted in (a-b,e-f) while the vertical velocity component is in (c-d,g-h).

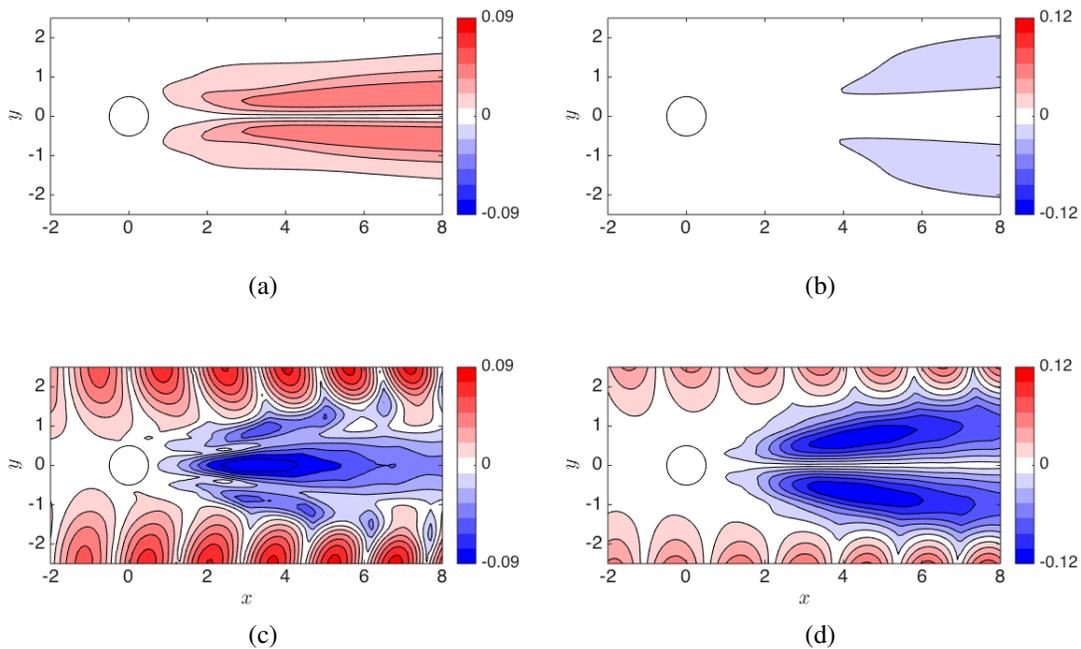


Figure 4.13: Discrepancy between the modulus of the resolvent and DMD modes for ω_s (top) and $2\omega_s$ (bottom). The u -component is on the left while the v -component is on the right.

4.6 Parasitic Modes and Triadic Interactions

It is proposed that the DMD mode for $\omega = 2\omega_s$ can be recovered by attempting to compute directly the nonlinear forcing at this temporal frequency. The quadratic nonlinearity of the NSE leads to the well-known triadic compatibility constraint for scale interactions (McKeon, 2017):

$$\mathbf{f}_k = \begin{pmatrix} f_{ku} \\ f_{kv} \\ f_{kw} \end{pmatrix} = \int \int \int_{k=k'+k'', k' \neq -k''} -(\mathbf{u}_{k'} \cdot \nabla) \mathbf{u}_{k''} dk, \quad (4.5)$$

where \mathbf{k} is the wavenumber vector (k_x, k_z, ω) . For global modes, only the temporal frequencies need to be considered and they must satisfy $\omega_1 + \omega_2 = \omega_3$. One possibility, therefore, is simply to have

$$\hat{\mathbf{f}}(2\omega_s) \approx \hat{\mathbf{u}}(\omega_s) \cdot \nabla \hat{\mathbf{u}}(\omega_s) \approx \hat{\boldsymbol{\psi}}_1(\omega_s) \cdot \nabla \hat{\boldsymbol{\psi}}_1(\omega_s), \quad (4.6)$$

since $\omega_s + \omega_s = 2\omega_s$. Equivalently, the nonlinear forcing for $\omega = -2\omega_s$ is the self-interaction of $\hat{\boldsymbol{\psi}}_1(-\omega_s)$, or $-\omega_s + -\omega_s = -2\omega_s$. A similar idea was explored by Rosenberg, 2018 for exact coherent solutions in Couette flow.

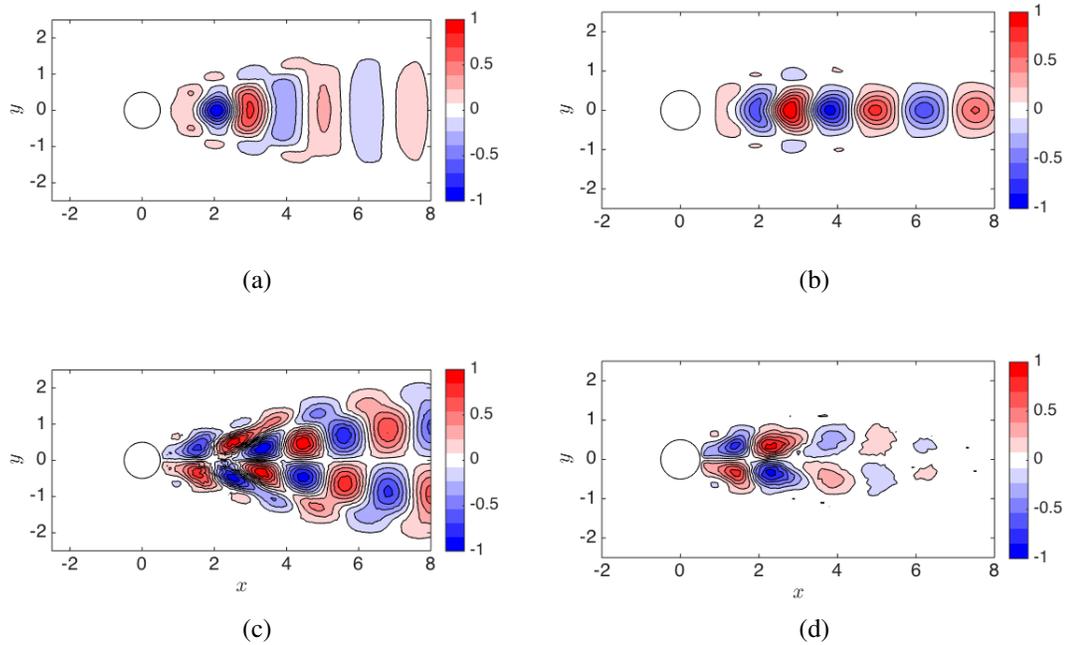


Figure 4.14: Nonlinear forcing $\hat{\mathbf{f}}(2\omega_s)$ computed directly from the DNS (left) and approximated by the self-interaction of $\hat{\boldsymbol{\psi}}_1(\omega_s)$ (right). The x -component is plotted in (a,b) while the y -component is in (c,d).

To test this idea, the nonlinear forcing $\mathbf{u} \cdot \nabla \mathbf{u}$ is computed from the DNS snapshots and is Fourier-transformed in time. The nonlinear forcing at $2\omega_s$ is compared to the nonlinear forcing generated by the self-interaction of $\hat{\boldsymbol{\psi}}_1(\omega_s)$ and very good agreement can be observed in Figure 4.14 in the region behind the cylinder or $0.5 < x < 3$. The agreement is not as good further downstream. Multiplying the resolvent operator for $\omega = 2\omega_s$ by this nonlinear forcing, or

$$\hat{\mathbf{u}}(2\omega_s) = \mathcal{H}(2\omega_s)(\hat{\boldsymbol{\psi}}_1(\omega_s) \cdot \nabla \hat{\boldsymbol{\psi}}_1(\omega_s)), \quad (4.7)$$

results in a velocity response mode which resembles the structure of the DMD mode from Figure 4.12 quite closely as seen in Figure 4.15. The discrepancy between the forced resolvent modes and the DMD modes is plotted in Figure 4.16. The u -component is underestimated around $x = 2$ and overestimated when $x > 5$. The spreading of the mode in the y -direction, furthermore, is not perfectly captured as the spatial support of the resolvent mode is concentrated closer to the centerline than the DMD mode. The v -component suffers from the same problems although the amplitudes of the discrepancy are slightly lower than those of the u -component. In spite of these differences, the forced resolvent mode is a significantly better match

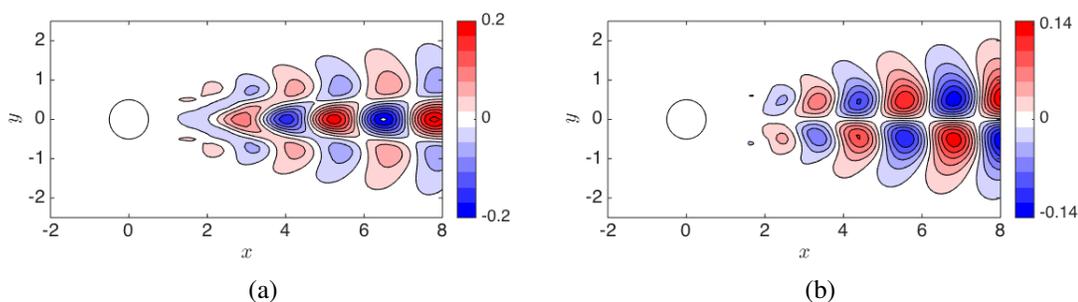


Figure 4.15: The velocity response $\hat{u}(2\omega_s)$ in (a) and $\hat{v}(2\omega_s)$ in (b) computed by forcing the $\mathcal{H}(2\omega_s)$ resolvent operator with $\hat{\psi}_1(\omega_s) \cdot \nabla \hat{\psi}_1(\omega_s)$ for comparison with Figure 4.12(f) and (h), respectively.

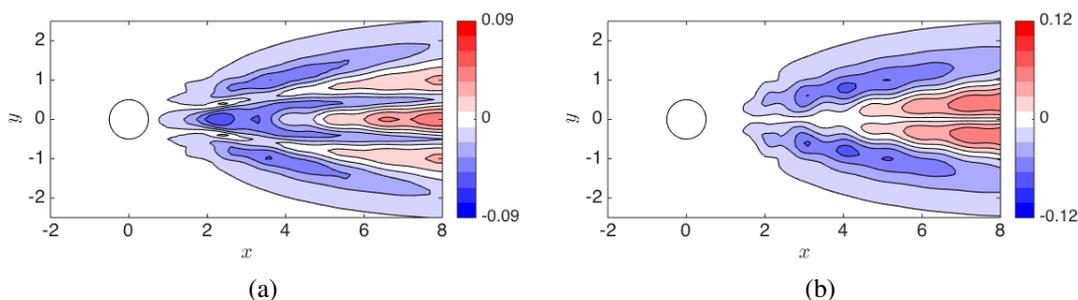


Figure 4.16: Discrepancy between the modulus of the forced resolvent and DMD modes for $2\omega_s$. The u -component is on the left while the v -component is on the right.

than the optimal response mode $\hat{\psi}_1(2\omega_s)$.

Even though Equation 4.6 is an approximation of the nonlinear forcing at $2\omega_s$, it is able to produce a good approximation of the structure associated with $\omega = 2\omega_s$. The rest of the nonlinear forcing $\hat{f}(2\omega_s)$ comes from two other sources including additional triadic interactions of velocity modes which sum to $2\omega_s$ such as $3\omega_s - \omega_s$ and the interactions of suboptimal resolvent modes with nontrivial amplitudes. This method of approximating the nonlinear forcing to obtain a mode shape which closely resembles the data is much more efficient than attempting to consider suboptimal modes of the resolvent operator at $\mathcal{H}(2\omega_s)$. The computation of the first ten singular modes requires over 100 cpu hours and none of the modes resembles the DMD mode. The above method, on the other hand, takes 0.1 cpu hours.

For this two-dimensional case, the resolvent identifies an instability mechanism as opposed to non-normal growth which is seen more commonly in wall-bounded

shear flows. The results of Schmidt et al. (2017a) reinforce this idea as the leading resolvent mode is representative of a Kelvin-Helmholtz instability in turbulent jets. Temporal frequencies which do not correspond to a linear amplification mechanism are parasitic in the sense that they are driven by the nonlinear interactions of highly amplified resolvent modes at resonant frequencies. The resolvent operator is still capable of predicting the velocity response but the forcing cannot be assumed to be arbitrary. In fact, knowledge of the nonlinear forcing, which can nevertheless be predicted from nonlinear interactions of highly amplified modes, is necessary at parasitic frequencies.

The approximation $\hat{f}(2\omega_s) \approx \hat{\psi}_1(\omega_s) \cdot \nabla \hat{\psi}_1(\omega_s)$ is not without precedent. Sipp and Lebedev (2007), for example, performed a weakly nonlinear expansion of the cylinder wake near the bifurcation threshold and obtained a linear equation, forced by the self-interaction of the ω_s mode, for the $2\omega_s$ mode. This analysis utilized the base flow at Re_c and approximated changes to the base flow and Reynolds stresses as the Reynolds number increased from terms in the expansion. As such, its range of validity does not extend far beyond Re_c unlike the current method, which needs the mean profile as an input yet is valid for any Reynolds number.

Turton et al. (2015) also alluded to the forcing terms for higher-order harmonics although their emphasis was on the global frequency only. Following their methodology, one can rewrite the linearized NSE as

$$in\omega u_n = \mathbf{L}u_n + \mathcal{N}_n, \quad (4.8)$$

where $\mathcal{N} = () \cdot \nabla()$, $\mathcal{N}_n = \sum_{m \neq 0, n} \mathcal{N}(u_m, u_{n-m})$, and the subscripts denote the order of the harmonic. Assuming the form $\|u_n\| \sim \epsilon^{|n|}$, they were able to show that the equation for $n = 1$ reduces to

$$i\omega_s u_1 = \mathbf{L}u_1, \quad (4.9)$$

signifying that u_1 is an eigenvector of the neutrally stable mean operator \mathbf{L} . This is known as RZIF, or an operator containing a real part zero, imaginary part frequency eigenvalue.

Here, their analysis can be extended to now consider $n = 2$, which, when assuming $\|u_n\| \sim \epsilon^{|n|}$, results in

$$\underbrace{2i\omega_s u_2}_{\epsilon^2} = \underbrace{\mathbf{L}u_2}_{\epsilon^2} + \underbrace{\mathcal{N}(u_1, u_1)}_{\epsilon^2} + \underbrace{\mathcal{N}(u_3, u_{-1})}_{\epsilon^4} + \underbrace{\mathcal{N}(u_{-1}, u_3)}_{\epsilon^4} + \dots \quad (4.10)$$

Retaining terms on the order of ϵ^2 results in

$$u_2 \approx (2i\omega_s \mathbf{I} - \mathbf{L})^{-1} \mathcal{N}(u_1, u_1), \quad (4.11)$$

which can be simplified to

$$\hat{\mathbf{u}}(2\omega_s) \approx \mathcal{H}(2\omega_s)(\hat{\boldsymbol{\psi}}_1(\omega_s) \cdot \nabla \hat{\boldsymbol{\psi}}_1(\omega_s)). \quad (4.12)$$

A significant number of the terms on the right-hand side of the equals sign in Equation 4.8 can be ignored to simplify the problem.

Higher-order harmonics can also be computed by considering only the largest terms on the right-hand side. To obtain the mode shape at $3\omega_s$, the nonlinear forcing must arise from both highly amplified frequencies as well as parasitic frequencies. Since ω_s and $2\omega_s$ are the most energetic frequencies in the wake, the nonlinear forcing at $3\omega_s$ is approximated as

$$\hat{\mathbf{f}}(3\omega_s) \approx \hat{\mathbf{u}}(2\omega_s) \cdot \nabla \hat{\boldsymbol{\psi}}_1(\omega_s) + \hat{\boldsymbol{\psi}}_1(\omega_s) \cdot \nabla \hat{\mathbf{u}}(2\omega_s). \quad (4.13)$$

The velocity response $\hat{\mathbf{u}}(3\omega_s)$ is computed by running $\hat{\mathbf{f}}(3\omega_s)$ through the resolvent operator $\mathcal{H}(3\omega_s)$ and the results are plotted in Figure 4.17. Despite truncating many other triads which could contribute to $\hat{\mathbf{f}}(3\omega_s)$, there is modest agreement between the u -velocity component of the resolvent and DMD modes. The agreement for the v -component, which has a less complicated structure, is slightly better. The locations where the perturbation energy is most concentrated (i.e. the spatial locations where the modes attains their highest amplitudes) is not accurately captured by the ‘forced’ resolvent mode which predicts much higher activity further downstream in the wake.

4.7 Nonlinear Forcing

The nonlinear forcing at the shedding frequency, i.e. $\hat{\mathbf{f}}(\omega_s)$, can also be computed from the nonlinear interactions of resolvent modes and even DMD modes, e.g.,

$$\hat{\mathbf{f}}(\omega_s) = \hat{\boldsymbol{\psi}}_1(-\omega_s) \cdot \nabla \hat{\mathbf{u}}(2\omega_s) + \hat{\mathbf{u}}(2\omega_s) \cdot \nabla \hat{\boldsymbol{\psi}}_1(-\omega_s), \quad (4.14)$$

since $2\omega_s - \omega_s = \omega_s$ is triadically consistent. Equation 4.14 distinguishes between $\hat{\boldsymbol{\psi}}_1$ and $\hat{\mathbf{u}}$ since the former is obtained from the singular value decomposition of $\mathcal{H}(\omega_s)$ while the latter arises from $\mathcal{H}(2\omega_s)$ forced by a specific, structured forcing which was computed earlier in Equation 4.7. In Figure 4.18, the true nonlinear forcing at ω_s , which is computed from the DNS snapshots, is plotted in panels (a-b) and compared to the predictions from resolvent analysis in (c-d) and DMD in (e-f). The resolvent mode prediction agrees fairly well for the y -component of the forcing although the amplitudes along the centerline are noticeably higher. The spatial

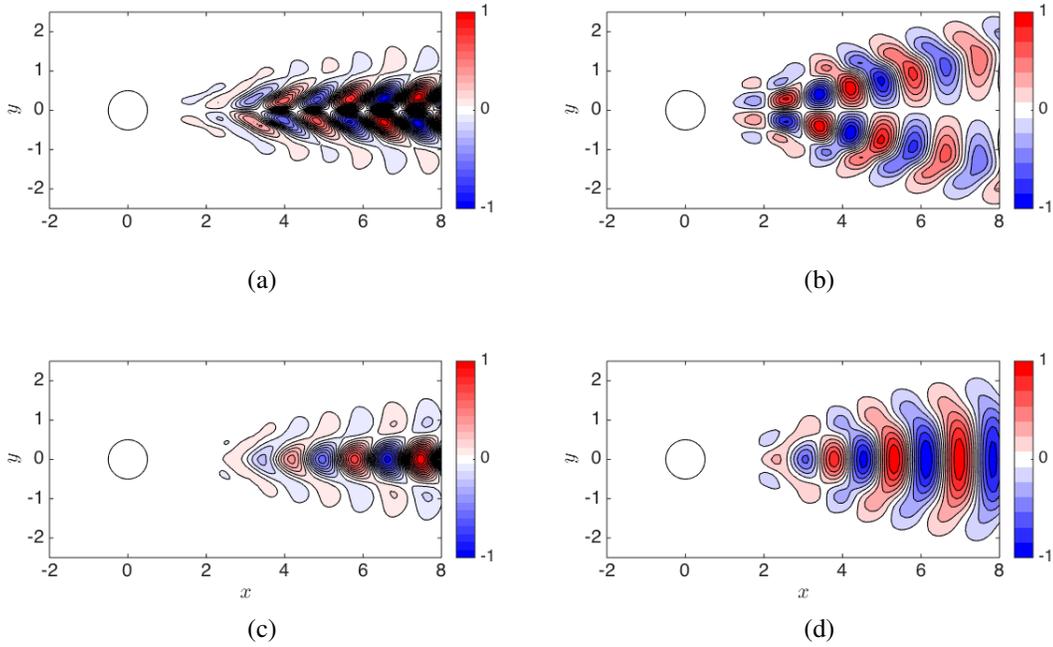


Figure 4.17: The velocity response $\hat{u}(3\omega_s)$ in (a) and $\hat{v}(3\omega_s)$ in (c) computed by forcing the $\mathcal{H}(3\omega_s)$ resolvent operator with an approximated $\hat{f}(3\omega_s)$. The modes are normalized so that the maximum contours are unity and are compared to the DMD modes in (b) and (d).

support of the resolvent prediction, furthermore, is much closer to the centerline, an observation which was already made for the resolvent modes when compared to the DMD modes in Figure 4.12. The x -component of the resolvent prediction, which does not agree with the DNS as well, deteriorates as one moves further downstream from the cylinder. The prediction immediately behind the cylinder, however, is acceptable. As will be seen in the next development, this region turns out to be the most important. $\hat{f}(\omega_s)$ predicted from nonlinear interactions of the DMD modes is in excellent agreement with the true nonlinear forcing.

In order for the vortex shedding to be sustained, the nonlinear forcing at the shedding frequency must have a nontrivial projection onto the optimal resolvent forcing mode:

$$\hat{\mathbf{u}}(\omega_s) = \sum_j \hat{\psi}_j(\omega_s) \sigma_j(\omega_s) \langle \hat{\mathbf{f}}(\omega_s), \hat{\phi}_j(\omega_s) \rangle = \sum_j \hat{\psi}_j(\omega_s) \sigma_j(\omega_s) \bar{\chi}_j(\omega_s), \quad (4.15)$$

where $\bar{\chi}_j$ is the weight associated with the j th response mode (McKeon et al., 2013). Since the separation of singular values for this flow is significant ($\sigma_1 \gg \sigma_2$),

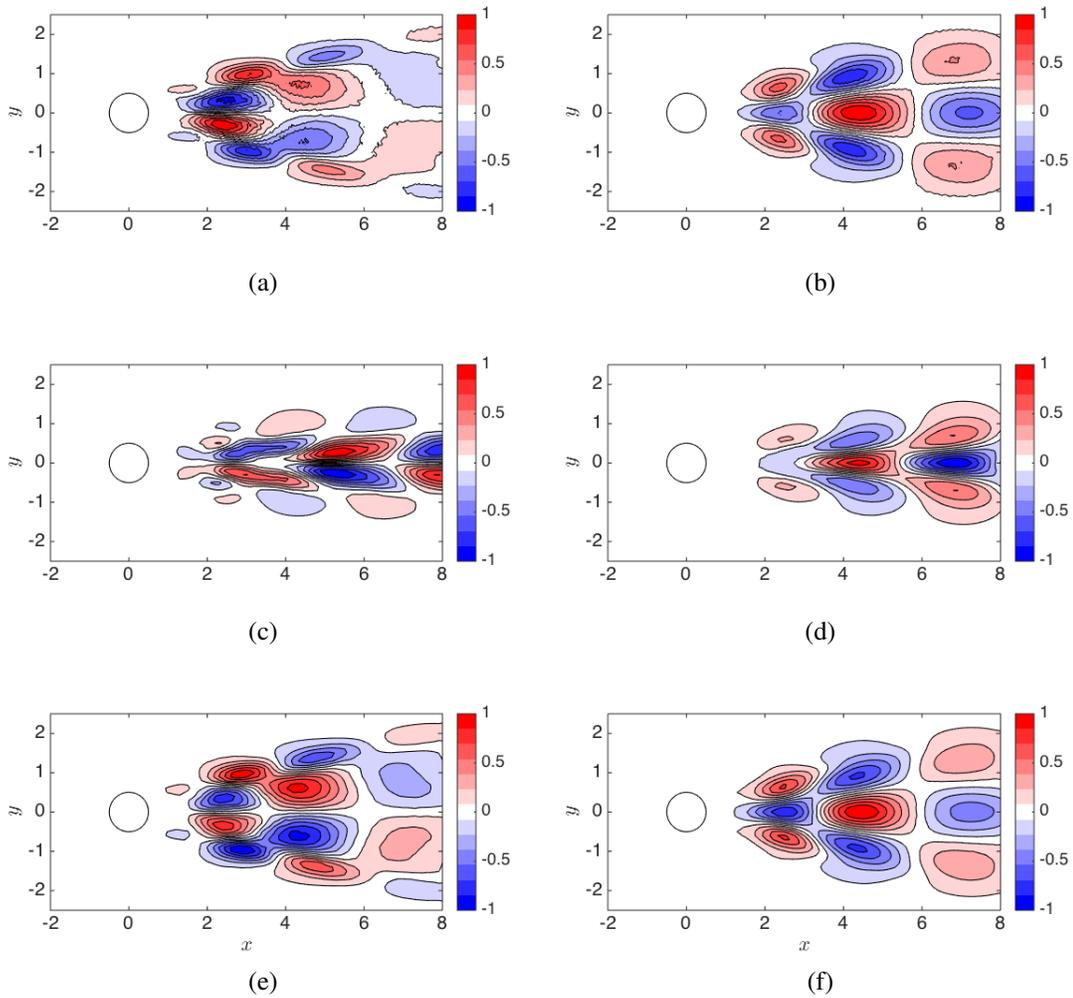


Figure 4.18: Nonlinear forcing at the shedding frequency $\hat{\mathbf{f}}(\omega_s)$ computed from the DNS in (a,b), triadically consistent resolvent response modes in (c,d) and triadically consistent DMD modes in (e,f). The x -component is plotted in panels (a,c,e) and the y -component in panels (b,d,f).

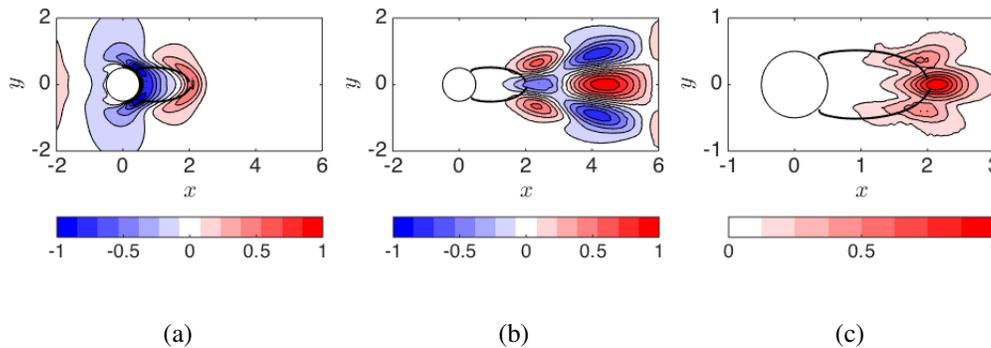


Figure 4.19: The optimal resolvent forcing mode (a), actual nonlinear forcing (b) and the magnitude of their overlap (c) for $Re = 100$. Only the v -component has been plotted in (a) and (b) while (c) is the projection computed using both velocity components. Thick black lines indicate the boundary of the recirculation bubble and the contours have been normalized such that the maximum amplitude is unity.

the projection of the nonlinear forcing onto the first forcing mode only is considered and, due to the simplicity of this flow, it is possible to visualize this projection. The resolvent forcing mode $\hat{\phi}_1(\omega_s)$, nonlinear forcing $\hat{f}(\omega_s)$, and the local magnitude of their projection $|\hat{\phi}_1(\omega_s) \cdot \hat{f}(\omega_s)|$ are plotted in Figure 4.19. The projection is zero everywhere except at the edge of the recirculation bubble.

This overlap location is consistent across all Reynolds numbers and reinforces the notion that the edge of the recirculation bubble coincides with the streamwise station of marginal absolute instability (Pier, 2002). The wavemaker portion ahead of the overlap is the heart of the instability and the rest of the flow convectively amplifies the perturbation (see Chomaz et al., 1988; Monkewitz, 1988). The schematic in Figure 4.20 delineates the region of absolute instability in the flow where the resolvent forcing and response modes overlap. When the mean flow advection along the centerline is no longer positive, the flow transitions from absolute to convective instability and there is no feedback between the forcing and response modes. Nonlinearity is responsible for saturating the growth of the instability at the downstream edge of the wavemaker.

4.8 Summary and a Closed Loop

The resolvent analysis for the low Reynolds number bluff body flows considered here identifies one dominant structure which is highly amplified at the shedding frequency even when the unsteady flow is three-dimensional. A Strouhal number based on the shedding frequency can be used to collapse the profile of σ_1 as a func-

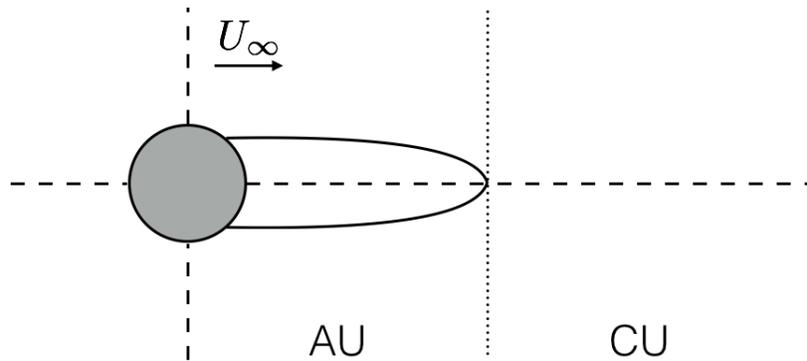


Figure 4.20: Instability regimes in cylinder flow where the flow is absolutely unstable (AU) immediately behind the cylinder in a region that coincides with the mean recirculation region and convectively unstable (CU) further downstream.

tion of ω and the mean profiles at the recirculation point can be scaled by ω_s . The shedding frequency also manifests itself in the resolvent modes and the convection velocities of the shed vortices. Sufficiently upstream of the cylinder, the convection velocity of adjoint perturbations is equal to the local mean velocity which is unity. The convection velocity of direct perturbations, or the vortices themselves, is also equal to the local mean velocity. Immediately behind the cylinder, the wavelength of the resolvent response modes is proportional to the recirculation length.

For frequencies where $\omega < 1.49\omega_s$, there is a separation of singular values and the resolvent operator is approximately low-rank. The most amplified mode belongs to the wake or W branch and the phase speed of the leading response mode matches the convection velocity of the vortices. The response modes at suboptimal frequencies, therefore, can be rescaled such that their profiles collapse along the centerline. At higher frequencies, the resolvent identifies a free-stream mode which exhibits critical-layer-like behavior. The phase speed matches the mean velocity and the modes belong to the C branch. For the lower frequency cases, the second resolvent mode $\hat{\psi}_2$ belong to this branch so they switch in order of precedence at $\omega = 1.49\omega_s$. The velocity response at $\omega = 2\omega_s$ cannot be predicted by the most amplified resolvent mode from a singular value decomposition of the operator.

Nonlinear interactions of resolvent response modes need to be considered to close the system and these are summarized in Figure 4.21. Forcing $\mathcal{H}(2\omega_s)$ by the self-interaction of $\hat{\psi}_1(\omega_s)$ yields the correct velocity response $\hat{u}(2\omega_s)$. This mode interacts with the $\hat{\psi}_1(-\omega_s)$ to produce the nonlinear forcing, which overlaps with the optimal forcing at the recirculation point, that is necessary to sustain the vortex

shedding. As the vortices form behind the cylinder, they accelerate before attaining their maximum convection velocity further downstream. The convection velocities scale with ω_s and q while the acceleration scales with ω_s^2 and l_m establishing a link between ω_s , the timescale of the fluctuations, and l_m , a statistical length. The wavemaker scales with l_m and coincides with the maximum of the projection of $\hat{f}(\omega_s)$ onto $\hat{\phi}_1(\omega_s)$.

The edge of the wavemaker, consequently, is an important region of the flow. The following chapter takes this into account in order to determine the minimum domain of measurements which is necessary to successfully data-assimilate the mean flow. The results of this chapter also suggest when it is or is not appropriate to use the most amplified resolvent modes to reconstruct the flow. Linear instability mechanisms, such as the vortex shedding for the cylinder, tend to dominate the fluctuations and leave a significant footprint on the mean profile. The NSE linearized around the mean, therefore, contain marginally stable eigenvalues which correspond to these phenomena.

The structures at temporal frequencies beyond the bandwidth of instability mechanisms cannot be easily identified from a singular value decomposition of the resolvent operator. Instead, they can be determined by forcing the resolvent operator by nonlinear interactions of highly amplified resolvent response modes. In the case of cylinder where there is a single instability mechanism at a discrete temporal frequency, the higher-order harmonics arise from the self-interaction of the globally most amplified mode. In the context of flow reconstruction, which is discussed in Chapter 7, it is important to respect when the rank-1 approximation is valid and obtain the structure of parasitic modes by approximating the structured forcing to the resolvent operator at those frequencies.

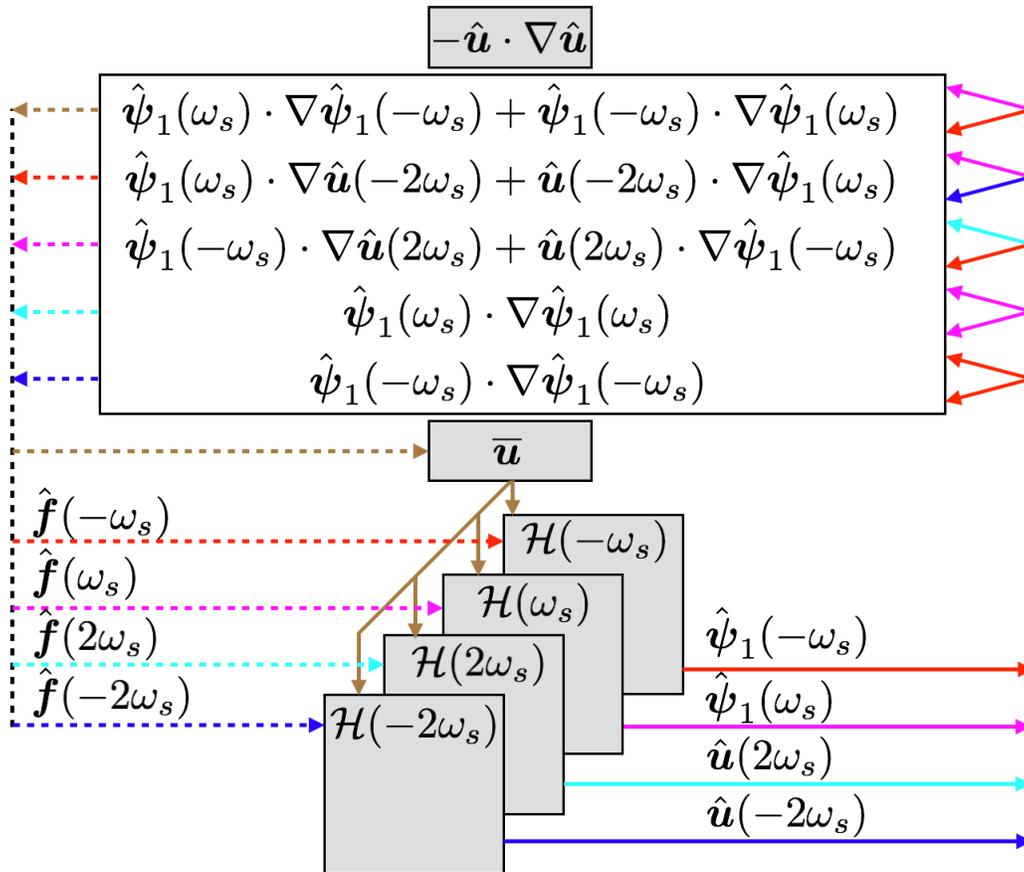


Figure 4.21: Block diagram inspired by Moarref et al., 2014 which portrays the triadic interactions needed to sustain cylinder flow at low Reynolds numbers, e.g. $Re = 100$. The mean profile is needed as an input to form the resolvent operators for each temporal frequency. A singular value decomposition of $\mathcal{H}(\omega_s)$ and $\mathcal{H}(-\omega_s)$ reveals the globally most amplified velocity response. Their nonlinear interaction forms the Reynolds stresses needed to support the mean profile while their self-interactions provide the nonlinear forcing for the first harmonics $2\omega_s$ and $-2\omega_s$. The resulting velocity response from this forced system interacts nonlinearly with the most amplified resolvent modes to form the nonlinear forcing which sustains the vortex shedding instability.

DATA-ASSIMILATION OF NUMERICAL DATA

In this chapter, the mean profiles for flow around a circular cylinder at $Re = 100$ and $Re = 300$ are obtained from data-assimilation of (subsamped) time-averaged DNS data. The flow physics and scaling identified in the previous chapter guide where measurements of the mean velocity are most effective for the algorithm to be successful. The motivation behind choosing a Reynolds number of $Re = 100$ stems from the fact that Mantić-Lugo et al. (2014) were able to obtain a self-consistent model that predicts the amplitude of the most unstable eigenmode in the flow. Connections are made between this approach and data-assimilation. Finally, a Reynolds number of $Re = 300$ is chosen to demonstrate that it is possible to data-assimilate the mean of a two-dimensional slice of the flow as long as the 2D continuity equation is satisfied. Since the flow is not 2D but periodic in the spanwise direction (Williamson, 1996), the span-averaged wake fulfills this requirement.

Table 5.1 summarizes the various cases that are considered in this chapter. The first two cases are full knowledge of the velocity field (velocity is known everywhere in the domain) and full-state knowledge (velocity and pressure are known everywhere in the domain). These are denoted as Cases 1A and 1B (see Table 5.1), respectively. Cost function A corresponds to Case 1A while cost function B corresponds to Case 1B. The final values of the cost function, denoted by the superscript t , are normalized by their initial values, denoted by the superscript 0

$$\mathcal{E}_A(\bar{\mathbf{u}}) = \mathcal{E}_A^t / \mathcal{E}_A^0, \quad (5.1a)$$

$$\mathcal{E}_B(\bar{\mathbf{u}}, \bar{p}) = \mathcal{E}_B^t / \mathcal{E}_B^0. \quad (5.1b)$$

Even though the domains are truncated for some cases, \mathcal{E}_A and \mathcal{E}_B will be used to make fair comparisons among all the cases. In other words, when measurements are available in a subset of the domain, the algorithm should still be able to reconstruct the flow in the entire domain.

5.1 Full Domain

Excellent agreement is obtained between the data-assimilated mean velocity profile for Case 1A and the DNS. The cost function declines to 2.14e-05% of its initial

Case	Domain	N	Weighting	States	\mathcal{E}_A	\mathcal{E}_B	λ
1A	Full	313,048	none	\bar{u}	2.14e-5		1.81e-3 + 1.02i
1B	Full	365,358	none	\bar{u}, \bar{p}		3.77e-5	1.76e-3 + 1.02i
2	Surface Pressure	313,080	none	\bar{u}, \bar{p}_{surf}	1.12e-4		1.81e-3 + 1.02i
3A	$x \in [-2, 6] \cup y \in [-2, 2]$	16,780	none	\bar{u}	3.29		1.83e-3 + 1.02i
3B	$x \in [-2, 6] \cup y \in [-2, 2]$	19,591	none	\bar{u}, \bar{p}		4.54	1.28e-3 + 1.02i
4A	$x \in [-1, 3] \cup y \in [-2, 2]$	9,502	none	\bar{u}	5.05		1.83e-3 + 1.02i
4B	$x \in [-1, 3] \cup y \in [-2, 2]$	11,102	none	\bar{u}, \bar{p}		6.56	1.63e-3 + 1.02i
5A	$x \in [-1, 2] \cup y \in [-2, 2]$	5,176	none	\bar{u}	9.73		1.31e-2 + 1.01i
5B	$x \in [-1, 2] \cup y \in [-2, 2]$	6,053	none	\bar{u}, \bar{p}		10.6	1.08e-2 + 1.01i
6A	$x \in [1, 3] \cup y \in [-2, 2]$	4,976	none	\bar{u}	5.31		3.52e-3 + 1.02i
6B	$x \in [1, 3] \cup y \in [-2, 2]$	5,815	none	\bar{u}, \bar{p}		6.9	1.71e-3 + 1.02i
7A	$x \in [2, 4] \cup y \in [-2, 2]$	4,902	none	\bar{u}	5.31		1.84e-3 + 1.02i
7B	$x \in [2, 4] \cup y \in [-2, 2]$	5,722	none	\bar{u}, \bar{p}		6.4	1.29e-3 + 1.02i
8	Full	313,048	wavemaker	\bar{u}	1.79e-5		1.76e-3 + 1.02i
9	Full	313,048	res. mode	\bar{u}	1.83e-5		1.77e-3 + 1.02i
10	Full	313,048	stability	\bar{u}	2.11e-5		1.75e-3 + 1.02i

Table 5.1: The domain, number and weighting of points, assimilated states, final normalized cost function, and least stable eigenvalue of the various cases studied.

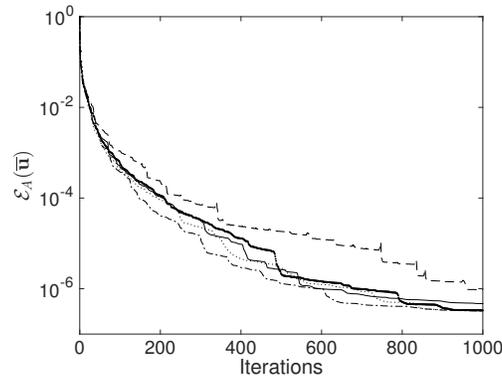


Figure 5.1: The cost functional $\mathcal{E}(\bar{\mathbf{u}})$ plotted for Case 1A (thin solid line), Case 1B (dashed line), Case 8 (dotted line), Case 9 (dashed-dotted line), and Case 10 (thick solid line). Only the first 1000 iterations have been shown.

value which is similar to the optimization performed by Foures et al. (2014). Case 1B is able to reconstruct the mean pressure and its cost function, which includes the mean pressure discrepancies, declines to $3.77e-05\%$ of its original value. The values of \mathcal{E}_A and \mathcal{E}_B are plotted in Figure 5.1 for the first 1000 iterations of the procedure. Figure 5.2(a,d,g) compares the mean pressure for Case 1A, Case 1B, and the DNS. As expected, Case 1B is almost an identical match with the DNS since it is given the mean pressure information in the whole domain. Case 1A, on the other hand, does well in most regions of the flow with the exception of the wake behind the recirculation zone from $2 \leq x \leq 4$. It is encouraging that it does this well without any knowledge of the mean pressure field.

Figure 5.2 also presents the x - and y - components of the mean forcing \mathbf{f} in panels (b,e,h) and (c,f,i), respectively. Case 1B exactly reconstructs both components of the mean forcing as it should since the adjoint velocity field is not constrained to be divergence-free. Case 1A, on the other hand, is limited to capturing only the solenoidal component of the mean forcing. The agreement between f_x for Case 1A in panel (b) and Case 1B in panel (e), nevertheless, is remarkable although there are small differences between them such as the location along the centerline of the peak mean forcing and the small lobe of negative mean forcing immediately behind the cylinder which does not appear in Case 1A. Notably f_y is virtually zero for Case 1A as seen in panel (c) when compared to Case 1B in panel (f).

In Case 2, pressure measurements on the cylinder surface are known in addition to full field velocity. As seen in Table 5.1, there is little impact on the quality of the velocity field reconstruction when measured by \mathcal{E}_A . Figure 5.3, which is a more

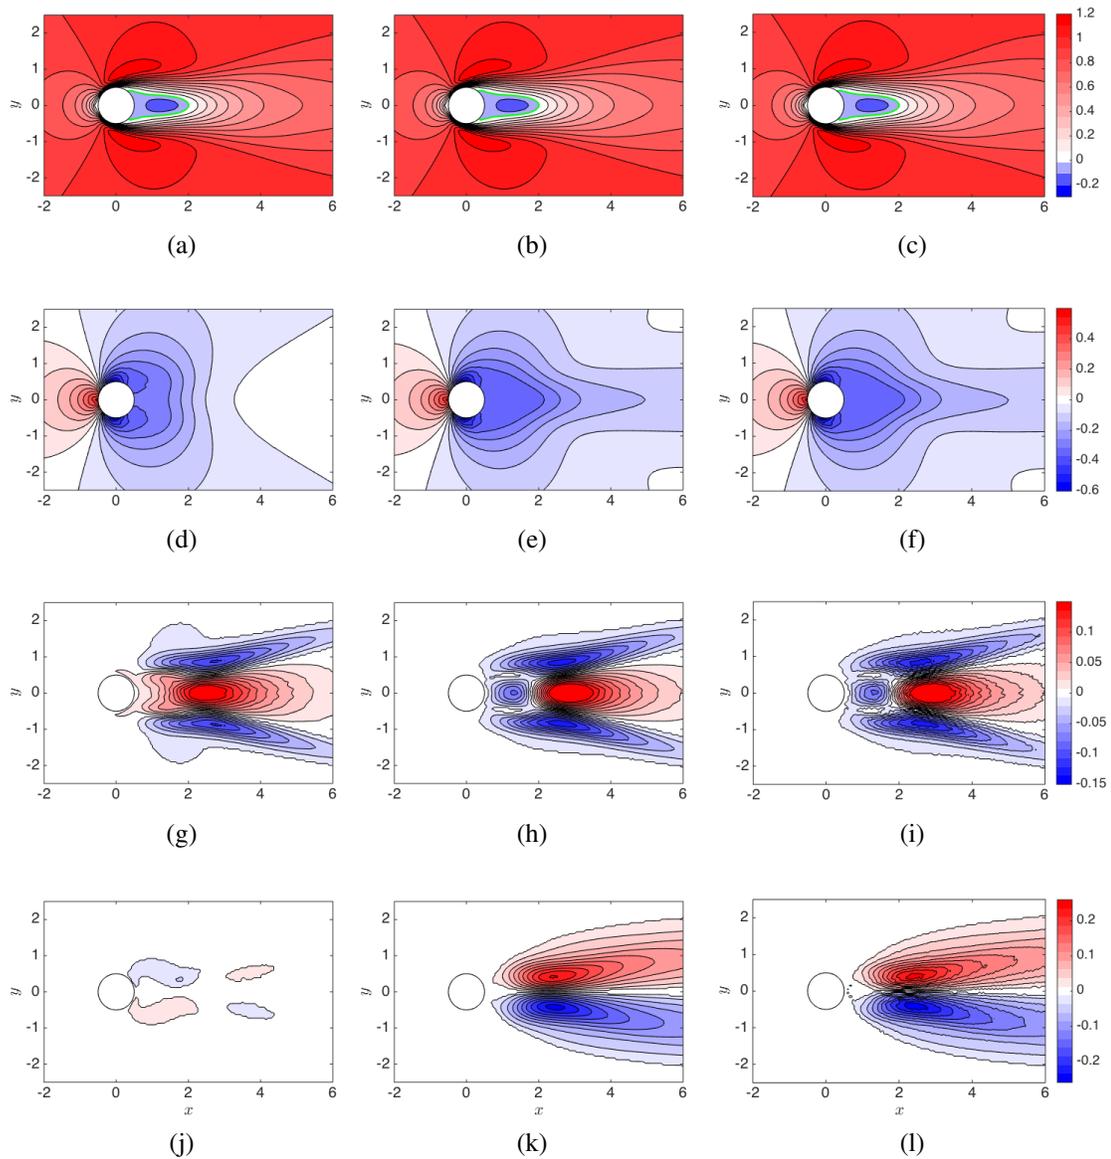


Figure 5.2: Contours of \bar{p} (left column), f_x (center column), and f_y (right column) for Cases 1A (top row), and 1B (middle row). These results are compared to DNS (bottom row).

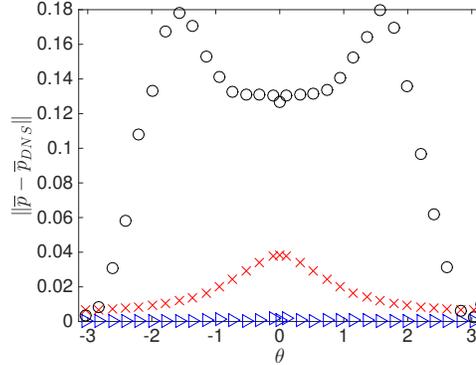


Figure 5.3: Absolute pressure discrepancy for the initial guess (open black circles), Case 1A (red crosses) and Case 2 (blue triangles). The angle θ represents where on the surface the measurement is taken. $\theta = \pm\pi$ corresponds to the leading edge stagnation point and $\theta = 0$ corresponds to the trailing edge stagnation point.

detailed version of Figure 5.2, confirms that the pressure along the cylinder surface is forced to agree with the measurements. In Case 2, data-assimilation correctly assimilates the mean pressure on the cylinder surface and the residual discrepancy between \bar{p} and \bar{p}_{DNS} is essentially zero. The residual discrepancy is also plotted for Case 1A in Figure 5.3 and, surprisingly, it agrees quite closely with the DNS. This is not true for the base flow, or initial guess, where there are significant deviations from the true mean pressure. It can be concluded that even without any pressure measurements, the data-assimilation algorithm provides a reasonable reconstruction of the mean pressure near the cylinder surface and, consequently, the mean forces.

5.1.1 Implications for experimental data

The full domain results suggest that the irrotational component of the forcing for this 2D flow is primarily in the y -direction while the solenoidal component is primarily in x . Such an observation can be explained by considering a simple case: parallel viscous flow in a channel where $\bar{\mathbf{u}} = \bar{u}(y)$. Due to the symmetries in the flow, the divergence of the Reynolds stress tensor is zero in the z -direction so $\mathbf{f} = [f_x \ f_y \ 0]^T$. Decomposing $\mathbf{f} = \nabla\xi + \mathbf{f}_s$ and assuming $\xi = \hat{\xi}(y)e^{i(k_x x + k_z z - \omega t)}$, one obtains

$$\mathbf{f} = \begin{pmatrix} ik_x \hat{\xi} \\ \partial \hat{\xi} / \partial y \\ ik_z \hat{\xi} \end{pmatrix} + \mathbf{f}_s. \quad (5.2)$$

Computing the divergence of (5.2) eliminates \mathbf{f}_s and results in

$$\nabla \cdot \hat{\mathbf{f}} = \begin{pmatrix} -ik_x^2 \hat{\xi} \\ \partial^2 \hat{\xi} / \partial y^2 \\ -ik_z^2 \hat{\xi} \end{pmatrix}. \quad (5.3)$$

The mean forcing corresponds to $k_x = k_z = 0$ and so the only non-zero component left in (5.3) is in y . For a parallel viscous flow, one can conclude that $\nabla \xi = f_y$ and $\mathbf{f}_s = f_x$.

The mean flow around the cylinder is not parallel although the transverse velocity component is outweighed by the streamwise component in most regions of the flow. The only portions of the flow where this is not true are the front half of the cylinder where the fluid is diverted around the body and the wake immediately behind the cylinder where there is a mismatch in f_x between Cases 1A and 1B. In general, as long as the flow is weakly non-parallel, f_x is accurately captured by the data-assimilation algorithm using velocity-only measurements. In experimental settings where the mean pressure is not measured, one could compare f_x computed from the experiment with its data-assimilation counterpart and expect a reasonable match assuming the flow is weakly non-parallel (see next chapter). It would not be necessary to compare the curl of the forcing (i.e. $\nabla \times \mathbf{f}$), as done in Symon et al. (2017), which requires computing two gradients of the Reynolds stress fields in order to assess the degree of success of the assimilation.

5.1.2 Stabilizing feedback

Data-assimilation can be interpreted as a controller which determines the feedback, in this case \mathbf{f} , which stabilizes the mean profile. A linear stability analysis of the base flow reveals a pair of unstable eigenvalues where $\lambda = 0.118 \pm 0.749i$ for $Re = 100$. A stability analysis of the mean profile is performed over various iterations of the optimization procedure for Case 1A. The real and imaginary components of the least stable eigenvalue are plotted in Figure 5.4. As the number of iterations increases, the growth rate declines until it reaches zero while the imaginary part increases until it reaches the frequency of the saturated vortex shedding. These results are reminiscent of Mantič-Lugo et al. (2014) who constructed a self-consistent model of the mean cylinder wake by adjusting the amplitude of the vortex shedding mode until the mean profile was marginally stable. They modeled \mathbf{f} using $2\text{Real}(\tilde{\mathbf{u}}_1 \cdot \nabla \tilde{\mathbf{u}}_1^*)$, where $\tilde{\mathbf{u}}_1$ is the vortex shedding mode. In this approach, it is not assumed that \mathbf{f} arises due to a single frequency but the validity of approximating it from a single mode can be useful, as will be discussed in the next section.

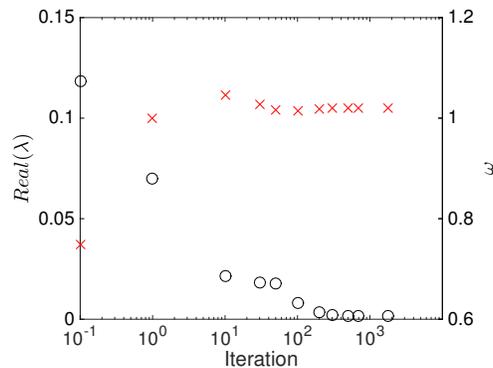


Figure 5.4: Real (open black circles) and imaginary (red crosses) components of the assimilated flow’s least stable eigenvalue λ at various iterations, the zeroth of which appears as 10^{-1} .

As an additional measure of whether the reconstructed mean velocity field is accurate, the least stable eigenvalue λ is reported in Table 5.1. A poorly data-assimilated field does not correctly identify λ . For the cases discussed in this section, λ is very accurately predicted which is to be expected since the algorithm has knowledge of the velocity everywhere in the domain.

5.2 Mean Pressure Correction

Since the global resolvent mode corresponding to ω_s dominates the flow, the right-hand side of (2.63) is treated as the correctly weighted triadic interaction of $\hat{\psi}_1(\omega_s)$ with its conjugate counterpart $\hat{\psi}_1(-\omega_s)$:

$$\nabla^2 \xi \approx \nabla \cdot \left[2\chi_1^2 \text{Real}(\hat{\psi}_1(\omega_s) \cdot \nabla \hat{\psi}_1(-\omega_s)) \right]. \quad (5.4)$$

The amplitude χ_1 of the resolvent response mode is calibrated using knowledge of the fluctuating velocity field at a single point in the flow, e.g., Gómez et al., 2016a. Solving for ξ results in a pressure field which ‘corrects’ the mean pressure computed from velocity-only data-assimilation. Reconstructing the unsteady pressure field can be accomplished from resolvent analysis by retaining the pressure term in the operator (e.g. Luhar et al., 2014; Gómez et al., 2016b).

The approximated pressure discrepancy computed from solving (5.4) is compared with the discrepancy between Cases 1A and 1B in Figure 5.5. The agreement is surprisingly good although the resolvent mode prediction is slightly narrower in the y -direction, similar to the resolvent modes when compared to the DMD modes in Figure 4.12. The other difference is that the resolvent prediction does not correct

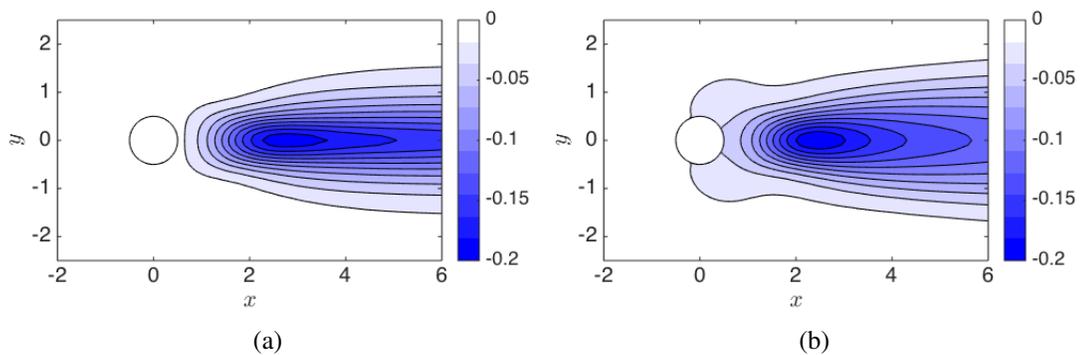


Figure 5.5: (a) The approximated mean pressure field discrepancy compared with (b) the difference between Case 1A and the DNS.

the mean pressure near the surface of the cylinder so it would not improve the mean pressure estimate of Case 1A in Figure 5.3. Despite the success of this approach, it is unlikely to work as well for higher Reynolds numbers since Mantič-Lugo et al. (2014) could not stabilize the cylinder wake using a single mode for $Re > 110$. Given the success of rank-1 models at predicting the spectra at various locations in the flow (e.g. Beneddine et al., 2016), it is possible to obtain an estimate of the Reynolds stresses which can be compared with the rotational forcing from velocity-only measurements. Regions of the flow where there is a difference between these two indicates where the mean pressure from the data-assimilation deviates from the true mean pressure.

5.3 Domain Truncation and Impact on Reconstructed Forcing

The physics and scaling from the previous section can now be capitalized on to identify where to measure the flow. The dynamically most significant region is arguably the downstream edge of the wavemaker, where the mean profile scales with the shedding frequency and the resolvent forcing mode overlaps with the nonlinear forcing. Some of the cases intentionally include or omit this region. Cases 3A and 3B are the minimum domain sizes ($x \in [-2, 6] \cup y \in [-2, 2]$) where the artificial truncation of the experimental domain has a minimal impact on the reconstructed forcing. The cost function for Case 3A declines to 3.29% of the original value of cost function A (recall this is the same cost function used for Case 1A) while Case 3B declines to 4.54% of the original value of cost function B. Figure 5.6(a-c) contains f_x for Case 3A, f_x for Case 3B, and f_y for Case 3B, respectively. It does not include f_y for Case 3A since it cannot be recovered from velocity-only measurements. All of the forcings match well with their full domain counterparts in

Figure 5.2. As such, they can be considered near the ‘truth’ against which the other cases can be compared.

5.3.1 Capturing the Wavemaker

Cases 4A and 4B include the wavemaker region ($x \in [-1, 3] \cup y \in [-2, 2]$) but omit a portion of the flow where the pressure discrepancy between the base and mean flows is greatest. The data-assimilated velocity is close to the true mean for both Cases 4A and 4B as cost function A reaches 5.05% of its initial value. The pressure reconstruction is somewhat impacted as it declines to 6.56%. The truncation of the domain has a clear impact on the reconstructed forcings which are plotted in Figure 5.6(d-f). For Case 4A, f_x agrees closely with the truth despite truncating the domain behind $x = 3$. For Case 4B, the effect of measuring only up until $x = 3$ has a greater effect since the maximum f_x is concentrated at $x > 3$ instead of $x = 2.75$. One can identify where measurements are available simply by looking at the spatial support of f_y although the overall shape is in slightly better agreement with its DNS counterpart than f_x .

Cases 5A and 5B, where the domain is limited to $x \in [-1, 2] \cup y \in [-2, 2]$, are considerably worse as they do not capture the downstream edge of the wavemaker. Cost function A reaches 9.73% and B reaches 10.6%, values which are approximately a factor of two higher than the other cases. Moreover, the least stable eigenvalue predicted from a stability analysis is further from the truth than all the other cases (the real part is an order of magnitude higher) which do consider the entire wavemaker region. The problems associated with truncating the domain manifest themselves even more dramatically when plotting the forcing for these two cases as seen in Figure 5.6(g-i). The spatial support of f_x and f_y is squeezed into a small region immediately behind the cylinder. To compensate, the algorithm overestimates the forcing amplitudes near the downstream edge of the measurement zone.

5.3.2 Nonlinear Forcing and Maximum Pressure Discrepancy

Limiting the domain to the region where the nonlinear forcing and optimal forcing mode overlap in Cases 6A and 6B ($x \in [1, 3] \cup y \in [-2, 2]$) is shown to be sufficient for reconstruction of the mean flow. Assimilation of the mean velocity is quite successful (the cost function declines to 5.31%) but the mean pressure is less so (6.9%). Indeed, Figure 5.6(j-l) highlights how f_x is better reconstructed than f_y for this domain. Cases 7A and 7B limit the domain to where the initial pressure discrepancy $p_0 - \bar{p}$ is greatest ($x \in [2, 4] \cup y \in [-2, 2]$) and this yields improved

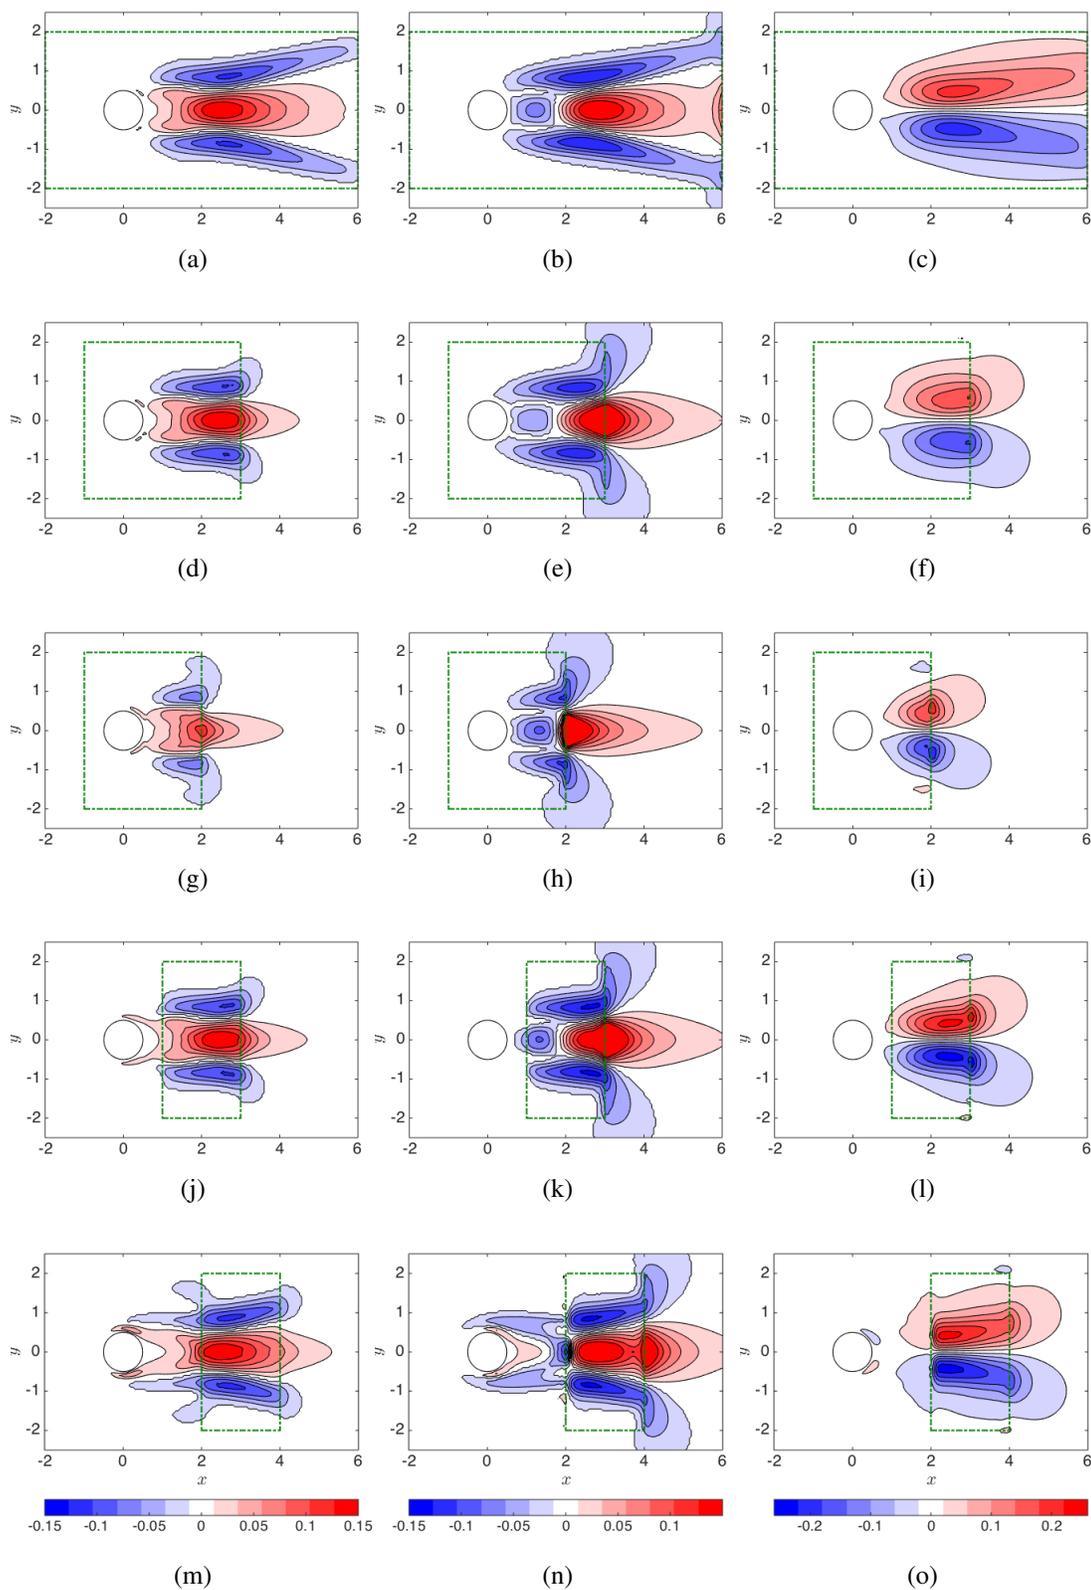


Figure 5.6: The reconstructed f_x (left column) using velocity-only measurements, the reconstructed f_x (centre column) using velocity and pressure measurements, and the reconstructed f_y (right column) using velocity and pressure measurements for various measurement domains. The results for each case have been plotted in order such that the top row corresponds to Cases 3A and 3B while the bottom row corresponds to Cases 7A and 7B.

pressure reconstruction (6.4%). The quality of the forcing, which is plotted in Figure 5.6(m-o), is better for f_y than it is for f_x but there is a qualitative improvement for both when compared to Cases 6A and 6B. Measurement domains confined to the region upstream of the cylinder or in the far wake, which perform considerably worse than the cases considered here (see Foures et al., 2014), are not discussed here.

It can be concluded that mean velocity measurements are most important at the edge of the wavemaker where the nonlinear forcing and resolvent forcing mode overlap. Additionally, as demonstrated by Leontini et al., 2010, a local stability analysis of the mean wake at the edge of the recirculation bubble predicts the frequency of the vortex shedding, signifying that it is essential for the data-assimilation algorithm to reproduce the profile at this particular streamwise location. The mean pressure, on the other hand, is most successfully recovered by obtaining pressure measurements where the discrepancy between the initial guess and the true mean velocity is greatest ($x \in [2,4]$). There is no pressure equivalent to PIV which makes it difficult to obtain pressure measurements in the wake. As demonstrated earlier, it is possible to utilize the correctly weighted resolvent response modes to ‘correct’ the mean pressure.

5.4 Weighting Measurement Points

Measurements may also be weighted less if they are more susceptible to noise contamination, for example. The weight for any point is given by

$$w = 1 + 2 \frac{\zeta}{\max(\zeta)}, \quad (5.5)$$

where ζ is the absolute value of some quantity of interest such as the wavemaker. Points far away from these regions, therefore, are still assigned a weight of at least unity. The multiplicative factor of two in Equation 5.5 is an arbitrary choice yet it does not play a significant role on the final outcome.

Up until this point, the algorithm has weighted all the points in the domain equally. It is now investigated as to whether weighting the points based on the physical and modeling insight of the previous sections can improve the efficiency of the procedure. The effect of weighting the measurements in three ways is considered. In Case 8 the points are weighted by the value of the wavemaker computed from the DNS mean and in Case 9 they are weighted by the modulus of the most amplified resolvent response mode computed from the DNS mean. Since the domain on

which experimental measurements is known may not be large enough to compute global resolvent modes, Case 10 considers the wavemaker computed from the assimilated flow at each iteration of the optimization. The wavemaker for the zeroth iteration, for example, is computed from the leading resolvent modes of the base flow. As the recirculation region shrinks during the data-assimilation procedure, the wavemaker also shrinks.

The functional $\mathcal{E}(\bar{\mathbf{u}})$ is plotted as a function of iteration for Cases 8-10 in Figure 5.1 alongside Cases 1A and 1B. While weighting the points is slightly beneficial in terms of the speed at which the assimilated field converges to the mean, there is no benefit in terms of the overall quality of the assimilation. Case 9 is arguably the best in terms of speed which is reasonable since the Reynolds stresses can be approximated from the resolvent response modes as discussed in the previous subsection. It can be concluded that the domain where experimental measurements are obtained is far more important than how points are weighted, which is encouraging since it is not necessary to redo the procedure once the weights are determined from resolvent analysis of the data-assimilated mean.

5.5 Extension to 3D Wake

The final objective of this chapter is to data-assimilate the mean cylinder wake for $Re = 300$, a Reynolds number at which Mode B is active and the wake is three-dimensional (Williamson, 1996; Barkley and Henderson, 1996). In addition to being time-averaged, the data are also averaged in the spanwise direction. The \bar{w} velocity component, therefore, is zero since the flow is periodic in the spanwise direction and the resulting mean profile $\bar{\mathbf{u}}$ satisfies $\partial_x \bar{u} + \partial_y \bar{v} = 0$. It is possible to compute a two-dimensional base flow solution at this Reynolds number but this is not feasible at higher Reynolds numbers (Sipp et al., 2010), particularly for the experimental data in the next chapter. As such, the algorithm begins by computing a base flow at $Re = 200$. Once the data-assimilated flow has converged to the mean at $Re = 200$, the Reynolds number is increased to $Re = 300$ using the data-assimilated at $Re = 200$ as an initial guess. The procedure continues at $Re = 300$ until convergence is achieved.

The final data-assimilated flow is compared to its DNS counterpart in Figure 5.7. The profiles are virtually indistinguishable which reinforces how converging at a lower Reynolds number before data-assimilating at a higher Reynolds number is a valid approach. Since pressure measurements are not included in the assimilation,

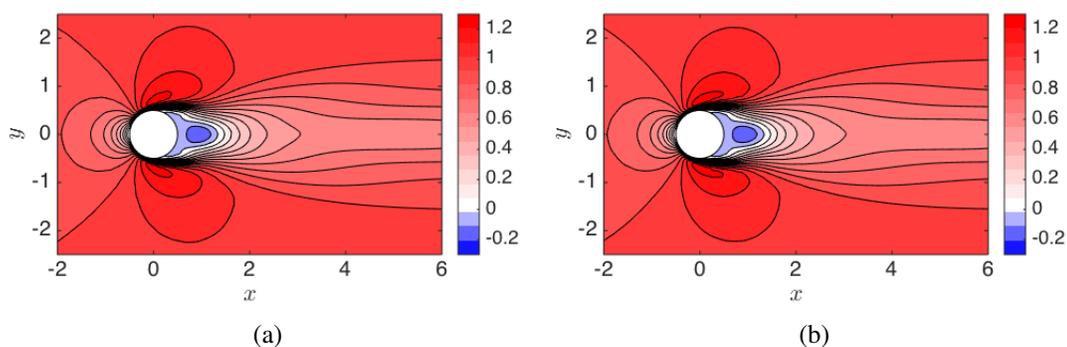


Figure 5.7: Contours of \bar{u} from the data-assimilation compared to (b) \bar{u} from the DNS.

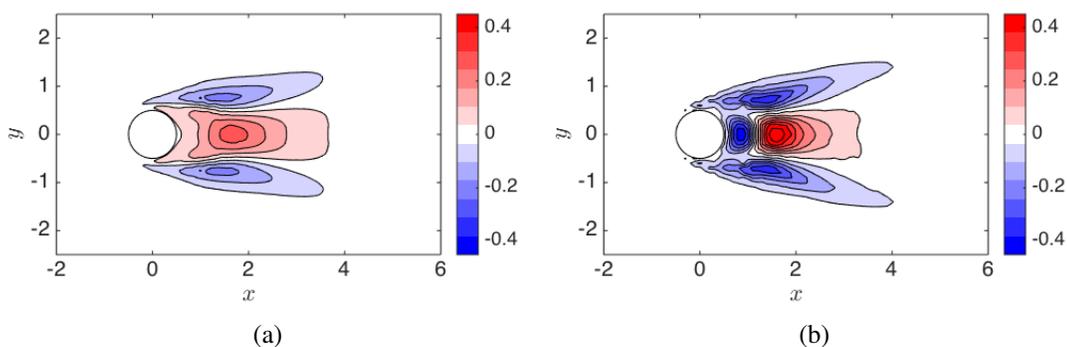


Figure 5.8: Contours of (a) f_x from the data-assimilation compared to (b) f_x from the DNS.

the x -component of the forcing f_x is compared between the data-assimilation and DNS in Figure 5.8. There is less agreement between the two when compared to the $Re = 100$ case in Figure 5.2. The magnitude of f_x is noticeably smaller than that of the DNS. It is likely that the effects of non-parallelism are stronger for the $Re = 300$ case since the recirculation bubble is shorter.

Instead of comparing f_x , it is possible to compare $\nabla \times \mathbf{f}$ as seen in Figure 5.9 since this removes the irrotational component of \mathbf{f} . Very good agreement can be observed between the two and the degree to which the contour levels match is significantly better than Figure 5.8.

5.6 Summary

Data-assimilation has been performed for the mean flows around a circular cylinder at $Re = 100$ and $Re = 300$. The measurements, which are an input to the algo-

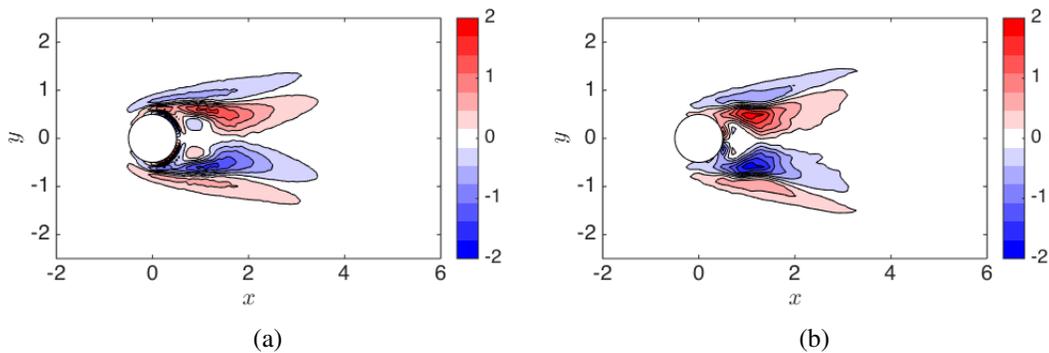


Figure 5.9: Contours of (a) $\nabla \times \mathbf{f}$ from the data-assimilation compared to (b) $\nabla \times \mathbf{f}$ from the DNS.

rithm, were varied by including or omitting pressure, truncating the domain where they were known, and weighting regions to emphasize important flow physics. It can be concluded that the minimum measurements needed for successful reconstruction of an unsteady flow are mean velocity measurements in the domain where the Reynolds stress gradients are concentrated as well as knowledge of the fluctuations at a single point to calibrate resolvent modes (more knowledge might be known had the flow been Fourier-transformed in space, e.g. Gómez et al., 2016a). Pressure measurements for incompressible flows are not needed as the pressure fluctuations can be accurately captured by the resolvent response modes. The mean pressure obtained from data-assimilation of velocity-only measurements can be corrected by solving a Poisson equation which is forced by the divergence of correctly weighted resolvent response modes interacting nonlinearly with their complex conjugate counterparts. The advantage of this method is that it avoids having to compute gradients of Reynolds stresses which amplify underlying noise in the experiment. Finally, a two-dimensional slice of a three-dimensional flow can be successfully assimilated as long as it satisfies the 2D continuity equation.

Data-assimilation would work for a flow field where there is a mean velocity component in the third dimension, but the recovered Reynolds stresses would not be correct in order to compensate for three-dimensionality. It should be verified that the mean flow approximately satisfies the two-dimensional continuity equation before applying this framework. One limitation of the method is the resolution of the input measurements which, up until this point, have been available at every mesh point. The next chapter investigates the minimum resolution required for successful flow reconstruction and makes use of the smoothing procedure since the

measurements are significantly more sparse than the number of mesh points. In terms of solving for the mean pressure, the methodology assumes that the first resolvent mode can sufficiently represent the fluctuating kinetic energy at a given ω . The flows considered in this thesis generally satisfy this requirement, but there are instances where multiple modes need to be considered to faithfully represent the dynamics, e.g., Schmidt et al. (2017a).

Chapter 6

DATA-ASSIMILATION OF EXPERIMENTAL DATA

Symon, S., Dovetta, N., McKeon, B. J., Sipp, D., and Schmid, P. J. (2017). “Data assimilation of mean velocity from 2D PIV measurements of flow over an idealized airfoil”. In: *Exp. in Fluids* 58.5, p. 61.

In this chapter, the data-assimilation algorithm is adapted for experimental data where the mean velocity measurements come from time-averaged, two-dimensional, two-component PIV. The beginning of the chapter focuses on the minimum quality of PIV data needed to successfully data-assimilate the mean for the idealized airfoil. In particular, the resolution and field of view are artificially altered from the original dataset to assess their impact on the reconstruction. The second half of the chapter considers the NACA 0018 airfoil at zero and nonzero angles of attack. The experimental parameters such as field of view and vector resolution are guided by the idealized airfoil results. These data-assimilated profiles are obtained to build reduced-order models in Chapter 7.

6.1 Base Flow and PIV Results

Before introducing the results of the data-assimilation algorithm, it is useful to compare the time-averaged PIV flow with the base flow solution at $Re = 2500$, which is used as an initial guess to the optimization. Figure 6.1 includes contour plots of the streamwise velocity for both the mean flow and the base flow as well as contour plots of vorticity overlaid with streamlines of the flow. Note that experimental data below the airfoil are not available since the laser sheet does not illuminate both sides of the airfoil. The most significant difference between the two fields is the streamwise length of the recirculation bubble, which for the base flow extends multiple chord lengths downstream of the airfoil. This is consistent with the observation that recirculation bubble lengths for base flows around closed bodies tend to scale linearly with the Reynolds number even when the flows become unstable as reported by Zielinska et al., 1997 for the cylinder wake.

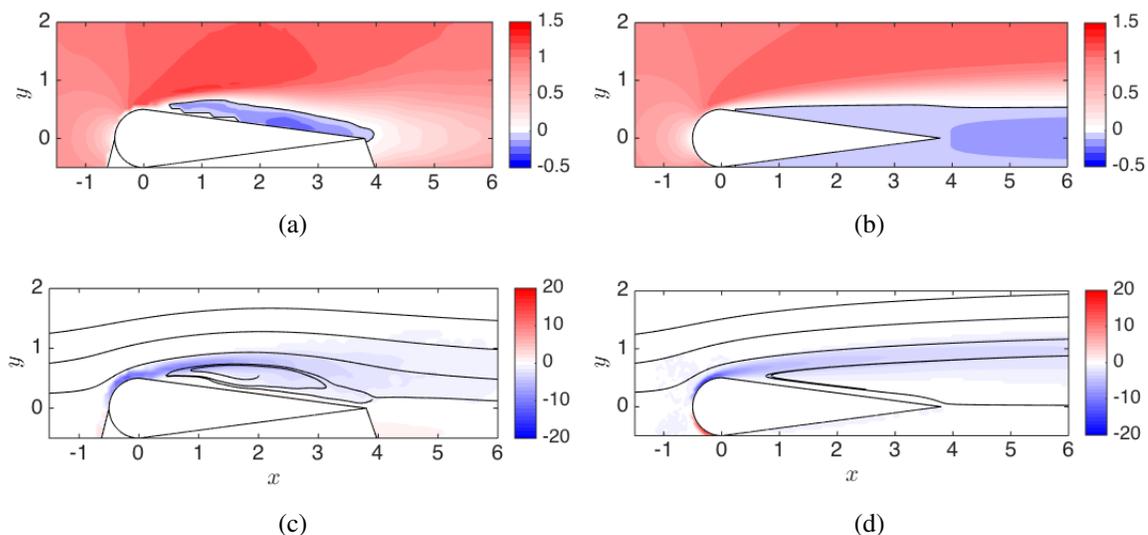


Figure 6.1: Base flow: A comparison between the mean flow (left) obtained from PIV and the base flow for $Re = 2500$ (right) around the idealized airfoil. The flows are visualized by contours of the streamwise velocity (top) and vorticity (bottom). The black contour in (a) and (b) corresponds to $\bar{u} = 0$ and streamlines are included in (c) and (d). Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

Case	Description	Number of meas. points	Q_{end}	Q_r	\mathcal{E}_r
11	Full-field	3669	6.76%	25.9%	40.1%
12	Half resolution	934	8.17%	31.3%	36.1%
13	Third resolution	416	11.5%	44.1%	47.2%
14	Small-field	1765	7.73%	26.1%	48.6%
15	Half-small	465	8.79%	33.7%	39.9%

Table 6.1: Summary of the data-assimilation results for the idealized airfoil.

6.2 Full-Field Information

The data-assimilation procedure is first conducted using full-field knowledge (3669 PIV vectors) or setting \mathcal{S} equal to the identity matrix in Equation 2.76. This is referred to as Case 11. Figure 6.2 compares the mean flow with the data-assimilated flow in a fashion similar to the comparison between the mean and base flows. Overall, very good agreement is observed between both mean velocity fields, particularly when comparing the contours of the streamwise velocity. The size, streamwise position, and shape of the recirculation bubble are also captured quite well although there is a slight discrepancy with the center of the bubble. It is difficult to determine this location from the PIV data due to the lack of spatial resolution in the

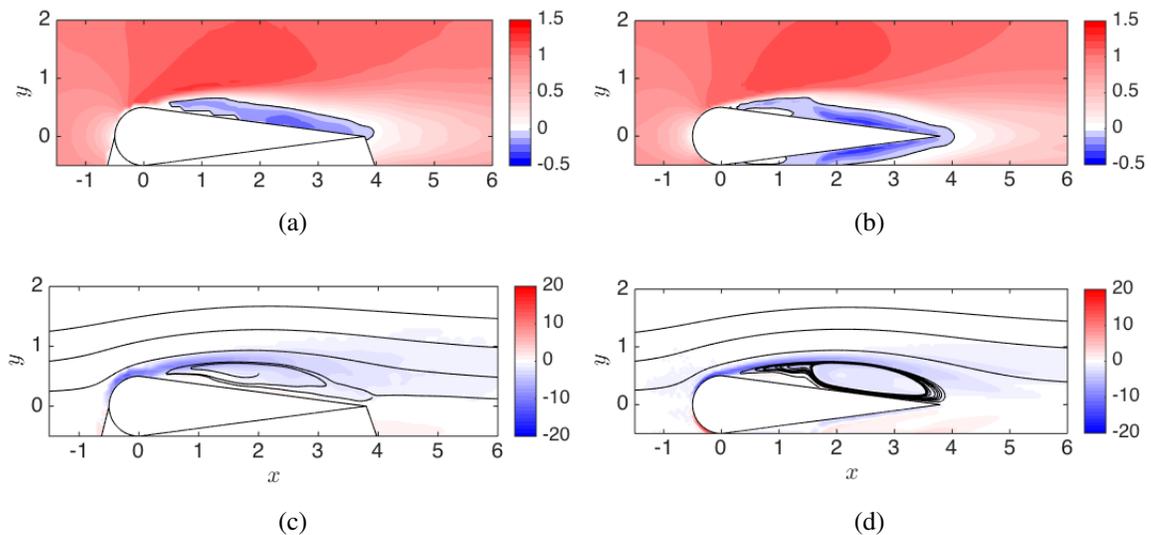


Figure 6.2: Case 11 results: A comparison between the mean flow (left) obtained from PIV and the data-assimilated flow (right) around the idealized airfoil. The flows are visualized by contours of the streamwise velocity (top) and vorticity (bottom). The black contour in (a) and (b) corresponds to $\bar{u} = 0$ and streamlines are included in (c) and (d). Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

transverse direction. The $\bar{u} = 0$ contour does not, however, match near the front of the bubble and this discrepancy is discussed in Section 6.4.

For the full-field case (Case 11), the cost functional decreases by $\mathcal{E}_r = 40.1\%$ and the experimental mismatch is $Q_{exp} = 6.76\%$, down from an initial value of 26.1% when the experimental mismatch is calculated with respect to the base flow ($Q_r = 25.9\%$). These results are summarized in Table 6.1. It is worth noting that the initial mismatch is misleadingly low since the PIV data do not extend far downstream ($x \leq 6$) and the base flow is at a Reynolds number lower than that of the experiment. The real base flow, which is very difficult to compute, would be even more unphysical than the current one. The region of the flow behind the airfoil is where the largest discrepancies are concentrated since the recirculation region is overestimated by the base flow.

The last result for the full-field information case is to compare the forcing of the assimilated flow with that of the experiment. In Figure 6.3, both components of \mathbf{f} are displayed. Similar to the cylinder case, f_y is not captured by the data-assimilation algorithm at all. The agreement for f_x , furthermore, is underwhelming since the contour levels for the experiment are higher than those for the assimilation although

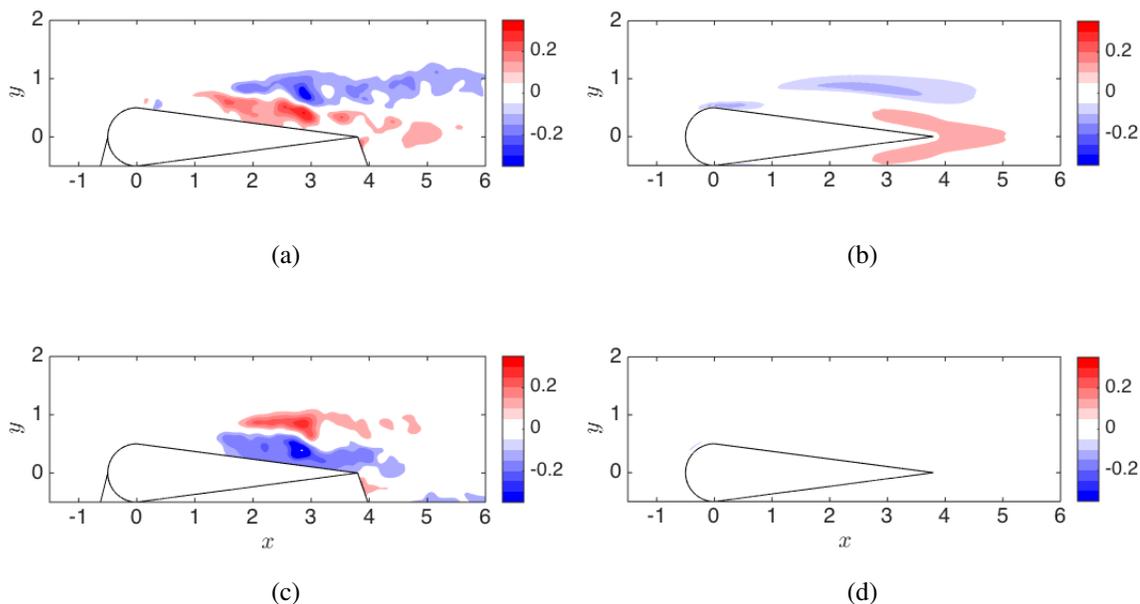


Figure 6.3: Case 11 results: A comparison between f_x (top) and f_y (bottom) for the experimental velocity field (left) and assimilated velocity field (right).

the overall structure is similar.

Since the algorithm can only capture the solenoidal part of the forcing, the curl of the forcing field $\nabla \times \mathbf{f}$ is presented in Figure 6.4 instead of f_x and f_y individually. To compute this quantity from the experiment, however, requires two spatial derivatives of the Reynolds stresses. A circular averaging or ‘disk’ filter in Matlab is used to smooth out the Reynolds stresses before computing gradients. No smoothing is necessary for the assimilated forcing. There is good agreement between the structure of $\nabla \times \mathbf{f}$ and the magnitudes of the peaks in the shear layer. It should be noted that complete agreement is not expected since the PIV data are not ideal and there is underlying noise in the measurements. The assimilation is also able to reconstruct the forcing at the leading edge of the airfoil, information which is not available from the experimental data since there are too few PIV measurements to resolve the high Reynolds stress gradients in this region. It is striking to observe how clean the assimilated forcing looks in comparison to the experimental one. The results reinforce how one of the primary motivations of data-assimilation is to remove noise and produce more highly resolved flow quantities. Assimilated fields are particularly useful at reproducing fields which need to be differentiated since experimental derivatives tend to amplify the underlying noise in a measurement.

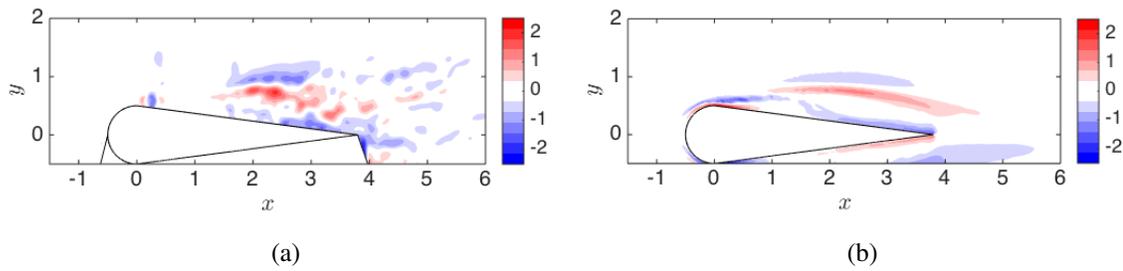


Figure 6.4: Case 11 results: A comparison between curl of the forcing $\nabla \times \mathbf{f}$ for the experimental velocity field (left) and assimilated velocity field (right). Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

6.3 Minimum Resolution and Domain

Using the projection operator \mathcal{S} , it is possible to investigate the effect of decreasing the full field resolution and its impact on the assimilated flow fields. Three different resolutions are tested by removing points in both the streamwise and transverse directions in a fixed pattern. To obtain an experimental field with half the resolution of the original one, for example, every other point is removed in both the streamwise and transverse directions. This process is repeated to obtain data sets with one third of the original resolution as well as one fourth of the original resolution.

Results for the assimilated flows at the lower resolutions are presented in Figure 6.5 for the one-half and one-third case (Case 12 and Case 13, respectively). The one-fourth resolution case failed to converge at higher Reynolds numbers, so the results have been omitted. The experimental mismatch calculations are summarized in Table 6.1. It is evident that there is a noticeable decrease in the assimilated flow field quality for the one-third case in comparison to the one-half case. The recirculation bubble is not smooth and the velocity contours have significantly deviated from the mean. This is largely due to the fact that the cell size over which the smoothing procedure is applied is nine times larger than it was in the full-field case. As a result, the range of velocities in a given cell can be quite large, particularly when the cell encompasses measurements above the shear layer where the velocity is close to the free stream and below the shear layer where it is close to zero. Nevertheless, the one-third case correctly reproduces the main features of the mean field, particularly the length and height of the recirculation bubble. It is also manages to recover an acceptable approximation of the forcing field as shown in Figure 6.6.

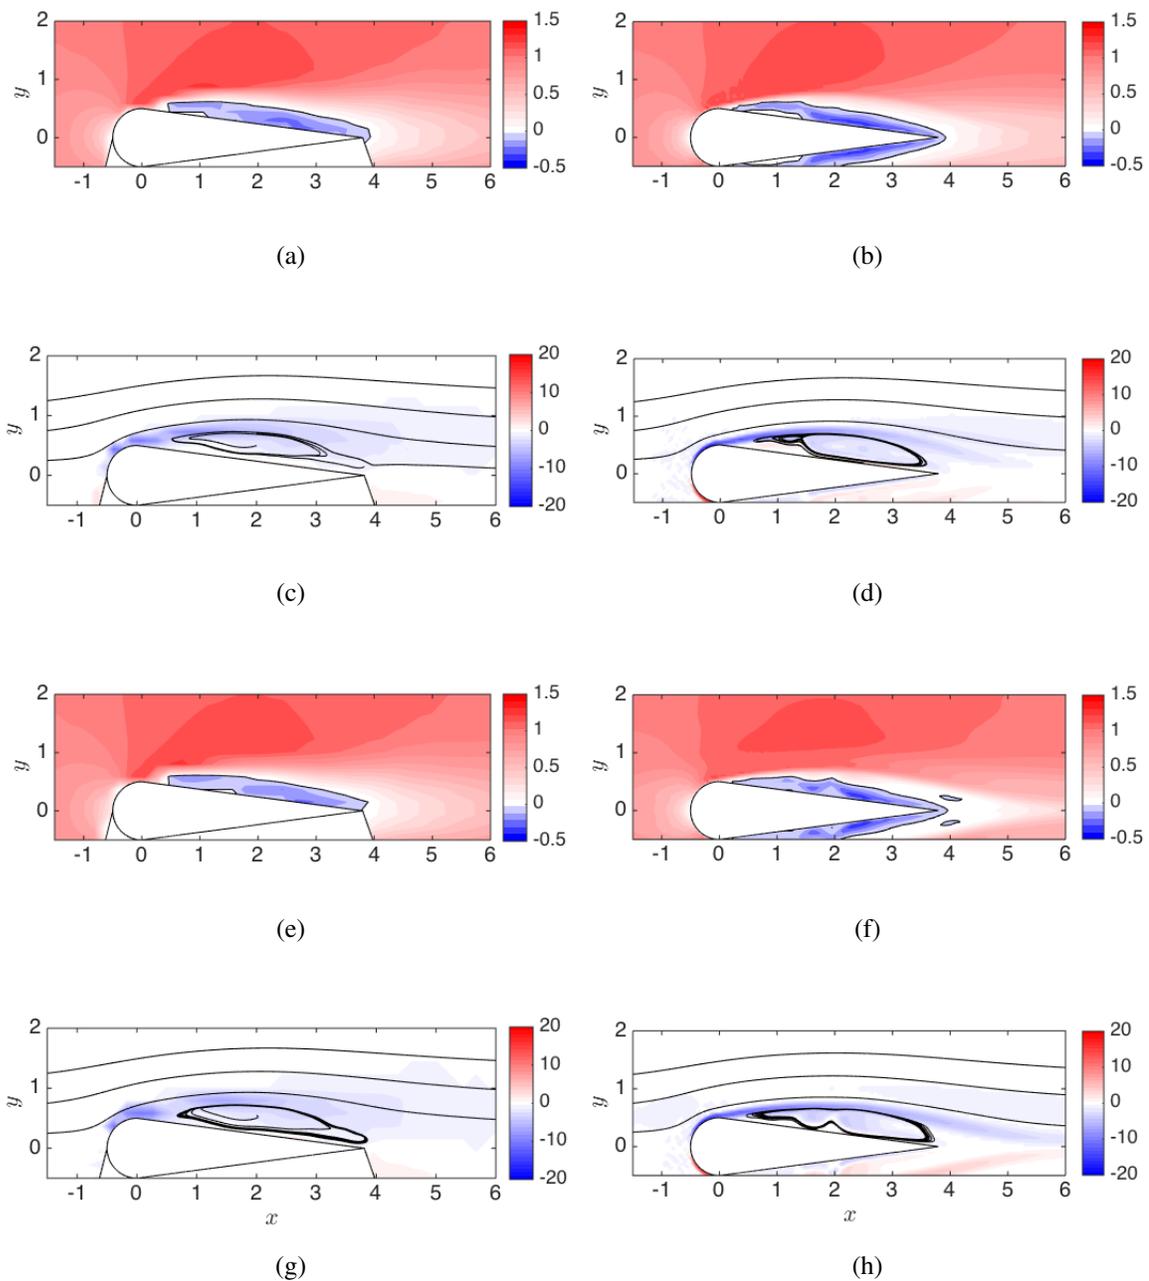


Figure 6.5: Case 12 and 13 results: The projected (experimental) mean flows are shown by (a) and (c) for the one-half resolution case (Case 12) and (e) and (g) for the one-third resolution case (Case 13). The corresponding assimilated flows are (b) and (d) for the one-half resolution case and (f) and (h) for the one-third resolution case. The flows are visualized by contours of the streamwise velocity and vorticity. Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

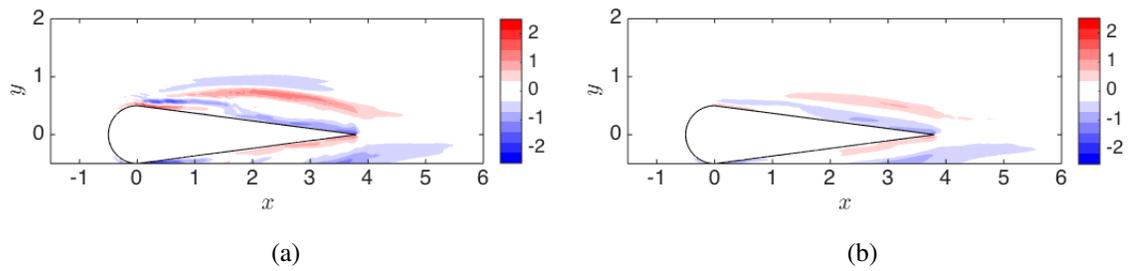


Figure 6.6: Case 12 and 13 results: Curl of the forcing $\nabla \times \mathbf{f}$ for the one-half resolution case (a) and the one-third resolution case (b). Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

It is important to note that the discrepancy field calculated by Equation 2.76 is not smooth since all the mesh points in a given cell have the same value. This field becomes decreasingly smooth as the resolution of the PIV data decreases since the cell size over which the averaging operator applies increases. Despite how discontinuous this smoothing may be, the solution to the adjoint equations, and consequently $\nabla \times \mathbf{f}$, is quite smooth as seen in Figure 6.6. It is worth noting, however, that the one-third case does not accurately locate the maximum contours of $\nabla \times \mathbf{f}$. The maximum positive contour in the forcing field, for example, is located further downstream when compared to the full-field and one-half cases.

Another way to reduce the number of experimental measurements is to truncate the PIV field of view. In the interest of brevity, only the smallest field of view for which there is no major sacrifice in data-assimilation quality is presented. The domain is limited such that PIV points outside the range $-1.5 < x < 6.0$, $0 < y < 1.5$ are excluded. The results are quantified in Table 6.1 as Case 14. As observed in Chapter 5, the measurement domain for the algorithm needs to include the spatial support of the Reynolds stresses. The domain for Case 14 is sufficiently large to capture this region as seen in Figure 6.4.

The trends observed in Table 6.1 are consistent with intuition - the more experimental measurements available, the closer the assimilated flow field is to the experiment. A promising result, however, is that decreasing the resolution or field of view of the reference data set does not significantly affect the assimilated flow field quality up to a point. As mentioned earlier, there is a significant difference between the one-half and one-third resolution cases. It is difficult to know *a priori*

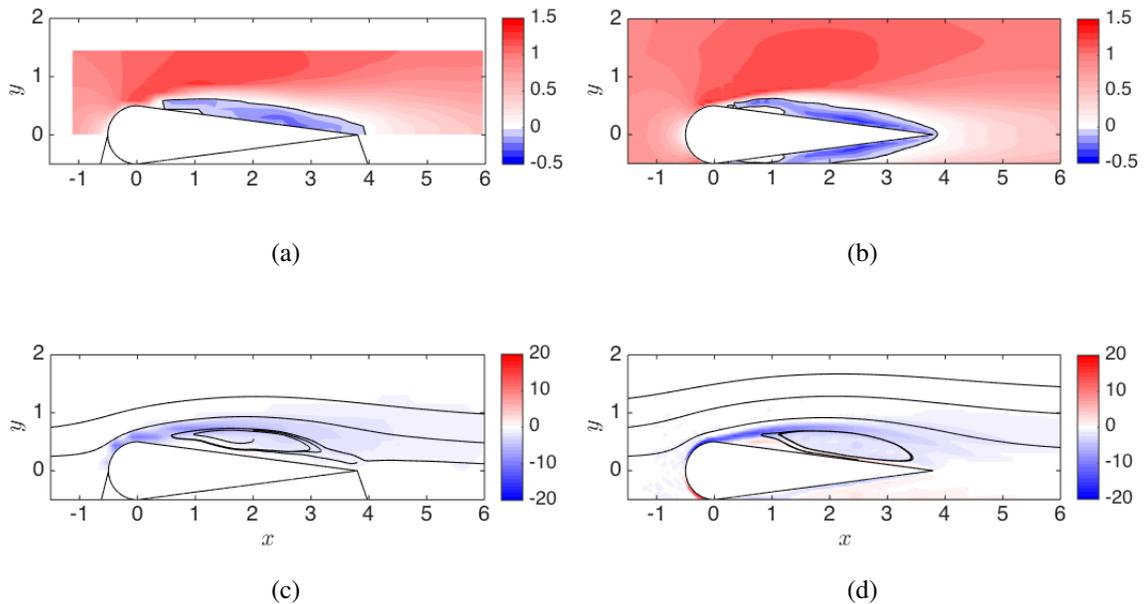


Figure 6.7: Case 15 results: A comparison between the projected (experimental) mean flow (left) obtained from PIV and the data-assimilated flow (right) around the idealized airfoil. The flows are visualized by contours of the streamwise velocity (top) and vorticity (bottom). The black contour in (a) and (b) corresponds to $\bar{u} = 0$ and streamlines are included in (c) and (d). Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

how much resolution is needed as this depends largely on the flow structures in the mean flow. The one-third resolution case struggles to resolve the shear layer while the one-fourth resolution case (not shown due to convergence problems) used grid cells which were too large to resolve the mean recirculation bubble.

To reduce the number of measurements further, it is possible to combine the two approaches mentioned so far by truncating the field of view and reducing the resolution of the reference data. Combining the small domain with the half resolution case decreases the number of points by approximately a factor of ten and is referred to as Case 15. A comparison of the assimilated flow using the approach with the mean flow is presented in Figure 6.7. The forcing fields are also displayed in Figure 6.8 to demonstrate that reduction in points by a factor of approximately ten has only a minimal impact on the results.

Data-assimilation works well within the current framework and reproduces smooth velocity and forcing fields. There are, however, limitations to this analysis that are worth mentioning. To begin with, the algorithm needs several hundred spatial

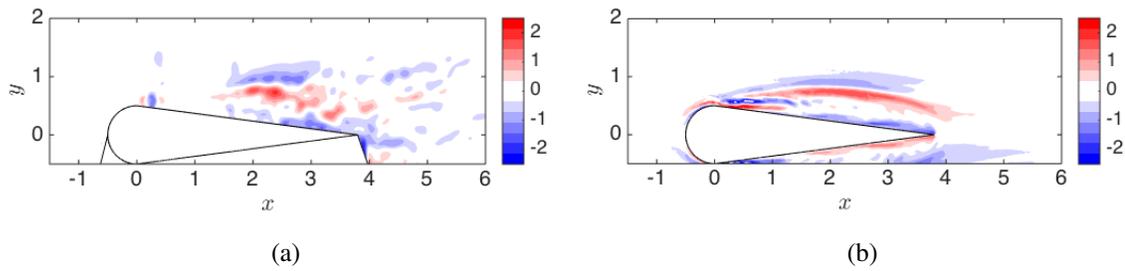


Figure 6.8: Case 15 results: A comparison between the curl of the forcing $\nabla \times \mathbf{f}$ for the experimental velocity field (left) and the assimilated velocity field (right). Note that the projection operator has not been applied to the experimental forcing since this would exacerbate the level of noise. Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

measurements for successful mean flow reconstruction. It is certainly possible to devise projection operators which divide the experimental data into variable cell sizes so that regions of the flow with high velocity gradients are more highly resolved than other regions. Experimental data, especially from PIV, are typically not collected on non-uniform grids so this study limits the approach to projection operators which can be reproduced by simply altering the experimental parameters. There are also temporal averaging requirements since a sufficient number of measurements is needed to obtain converged statistics. For bluff bodies, this might require some *a priori* knowledge of the shedding frequency, while for turbulent flows it is necessary to determine u_τ and the eddy turnover time.

6.4 Residual Discrepancy

An assessment of the remaining discrepancy between the experimental and assimilated mean velocity fields is discussed in this section. The influence of three-dimensionality of the flow is considered as well as elements of the data-assimilation framework.

6.4.1 Three-dimensionality

To quantify three-dimensional effects, the divergence of the experimental velocity is calculated. If the 2D continuity equation is not satisfied then this constraint may not be appropriate. The same calculation is performed for the data-assimilated flow field from Case 11, which is constrained to be divergence-free in 2D. This field is interpolated onto the PIV grid to determine what the permissible range of values

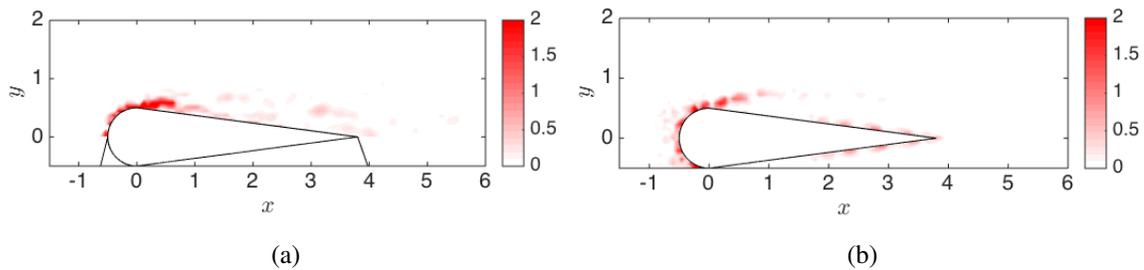


Figure 6.9: Contours of $\nabla \cdot \bar{\mathbf{u}}$ computed for the experiment (left) and the data-assimilated field (right) which is constrained to satisfy the 2D continuity equation. Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

would be for a flow field obeying the 2D continuity equation. Figure 6.9 compares $\nabla \cdot \bar{\mathbf{u}}$ for the experiment and the assimilation.

The plots indicate that $\nabla \cdot \bar{\mathbf{u}}$ is nearly zero everywhere except near the leading edge. Since $\nabla \cdot \bar{\mathbf{u}}$ is also high for the data-assimilated field at the leading edge it is reasonable to suggest that the grid resolution is not sufficient to capture the large velocity gradients in this region of the flow.

As mentioned earlier, 3D effects would be compensated for by the unknown momentum forcing term \mathbf{f}_s . It is clear from the results above that \mathbf{f}_s does not deviate significantly from the experiment. This means it is not being corrupted by spanwise velocity gradients which have been assumed to be zero. From these two observations, the role of 3D effects is negligible and does not play a significant role in the residual discrepancy.

6.4.2 Model simplifications

There are two modeling simplifications in the data-assimilation framework: boundary conditions and the coordinate system mapping between the experiment and the simulation. The boundary conditions for the simulation are very general and do not take into account the blockage ratio, which for this experimental configuration is approximately 7%. The PIV domain, furthermore, does not extend sufficiently far in the transverse direction to observe effects from the wall. Since there is excellent agreement between the assimilation and experiment outside the recirculation bubble of the airfoil, the effects of the boundary conditions and blockage ratio are negligible.

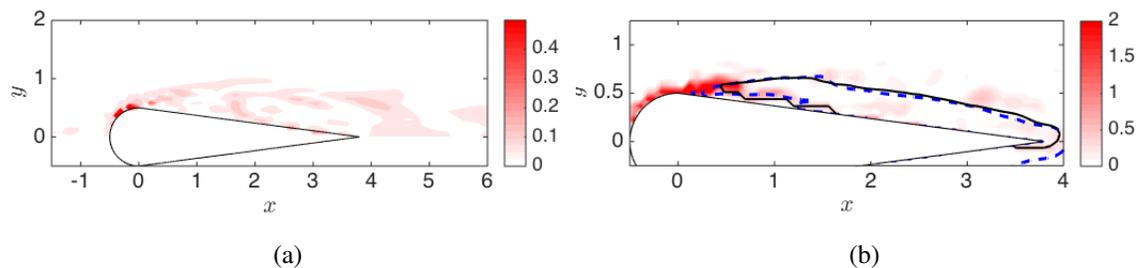


Figure 6.10: Residual discrepancy between the experimental and assimilated mean velocity fields for Case 11. The left-hand side quantifies the magnitude of the velocity discrepancy while the right-hand side shows the $\bar{u} = 0$ contour for the experiment (black, solid line) and assimilation (blue, dotted line) overlaid with contours of $\nabla \cdot \bar{\mathbf{u}}$. Reprinted by permission from Springer Nature: Springer, *Experiments in Fluids*, Symon et al. (2017), Copyright 2017.

The most significant source of the residual discrepancy, as seen in Figure 6.10, is the mapping between the experiment and the simulation. The largest contribution comes from several points along the leading edge of the airfoil while there are smaller contributions in parts of the recirculation bubble as well as the airfoil boundary near the trailing edge. There are three reasons to account for these results. First, it is difficult to precisely determine the airfoil location from the PIV data. Second, there are imperfections in the airfoil shape which is modeled as a cylinder followed by a wedge of half angle 7.5° . Finally, the flow separates in this region and PIV has difficulty pinpointing the separation point when the flow is laminar. This is discussed in greater detail at the end of the chapter

6.4.3 Choice of cost function

This study considers simple cost functions where the objective is to minimize the discrepancy at all PIV points subject to the incompressible RANS equations. No regularization parameters are introduced and all the PIV vectors are weighted equally. While effective, the major drawback to such an approach is that low velocity regions of the flow are treated as less important. The discrepancy between the velocity in the recirculation bubble, for example, is low even for the base flow case. Consequently, the procedure is biased towards high speed regions of the flow and it tries to correct these regions first.

The right-hand plot of Figure 6.10 compares the $\bar{u} = 0$ contour of Case 11 with its counterpart from the PIV data. The plot also includes contours of $\nabla \cdot \bar{\mathbf{u}}$ to emphasize

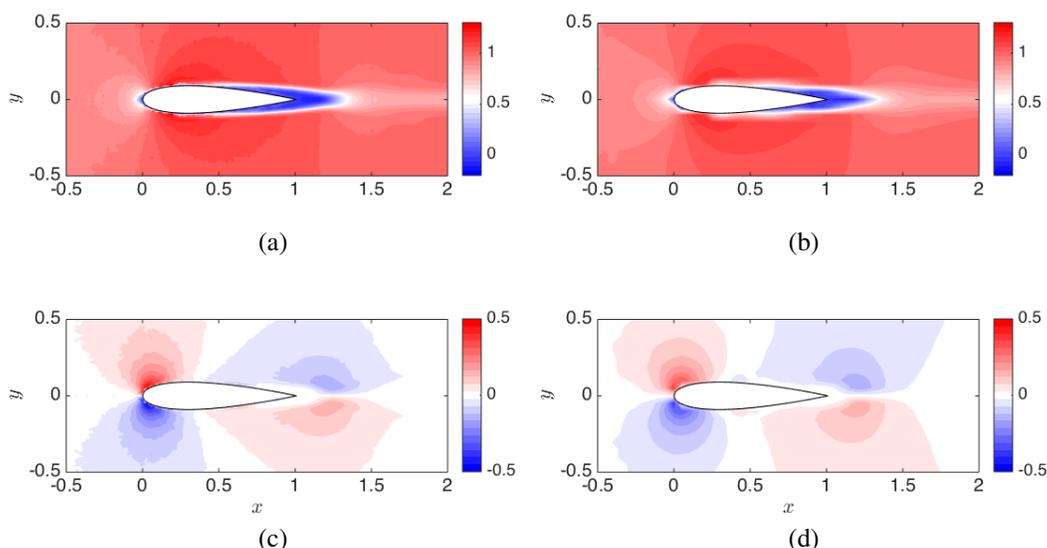


Figure 6.11: A comparison between the mean flow (left) obtained from PIV and the data-assimilated flow (right) around a NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$. The flows are visualized by contours of the streamwise velocity (top) and transverse velocity (bottom).

that three-dimensionality is not the leading cause of the discrepancy. There are small deviations along the entire contour which could be addressed by modifying the cost function to penalize these discrepancies to a higher degree.

6.5 Zero Angle of Attack Case

The PIV resolution and field of view for the NACA 0018 airfoil are selected based on the results from the idealized airfoil. Instead of investigating the effect of S on the quality of the reconstruction, S (Equation 2.76) is set to the identity matrix for the remainder of the thesis. The objective is simply to data-assimilate the mean profiles from the experiment so that they can be used as an input to resolvent analysis in Chapter 7. Unlike the idealized airfoil case, neither \mathbf{f} nor $\nabla \times \mathbf{f}$ from the PIV data are spatially filtered for any of the cases below.

6.5.1 $Re = 10250$

Both components of the data-assimilated mean velocity for the $\alpha = 0^\circ$, $Re = 10250$ case are plotted in Figure 6.11. Since the flow does not separate at this angle of attack, neither the streamlines nor the mean vorticity have been plotted. There is very good agreement between the PIV and data-assimilated fields for both velocity components. The quality of the PIV data is quite good although the contours

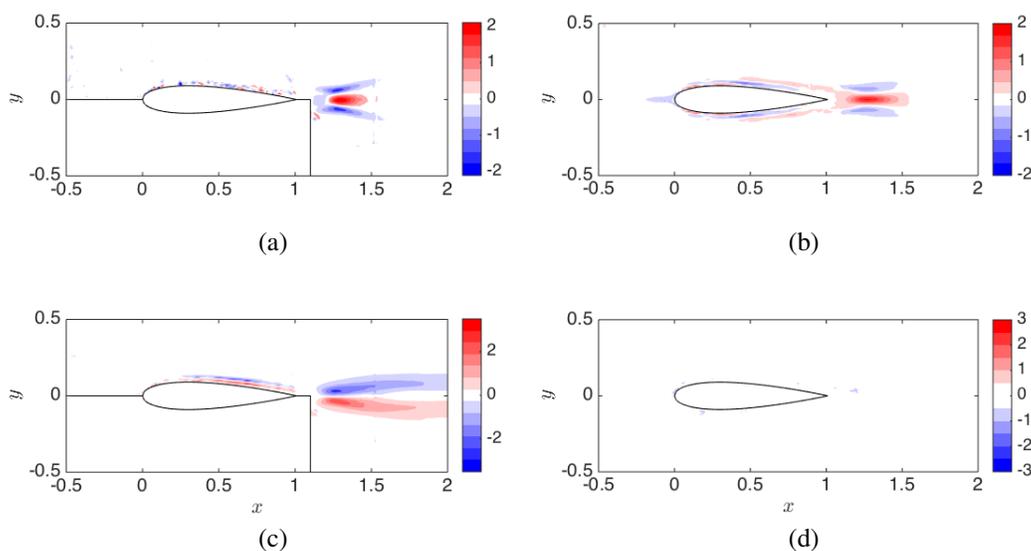


Figure 6.12: A comparison between f_x (top) and f_y (bottom) for the experimental velocity field (left) and assimilated velocity field (right) of flow around a NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$. The shadow regions below the airfoil have been masked out in panels (a) and (c).

of the data-assimilated means are slightly smoother than their experimental counterparts. The most significant difference between the two is the resolution of the mean velocity field near the airfoil surface. It is impossible to resolve the boundary layer yet have a field of view large enough to capture the Reynolds stress gradients. The benefit of data-assimilating the mean velocity field onto a higher resolution mesh where the no-slip boundary condition is satisfied will be observed in Chapter 7 when computing resolvent forcing modes.

Both components of the mean forcing \mathbf{f} are plotted in Figure 6.12. These results mirror those in Figure 5.2 and 6.3 in that f_x is reasonably captured although the contour levels are slightly lower than those from PIV. f_y , on the other hand, is not captured at all. One last point that is worth remarking is that the shapes of f_x and f_y observed from the PIV data are very similar to those for the circular cylinder at $Re = 100$. f_x is symmetric and has three lobes. The central lobe is positive while the upper and lower lobes are negative. f_y is anti-symmetric and has just two lobes like the cylinder. The similarity suggests that the dynamics of the airfoil at $\alpha = 0^\circ$ are very similar to low Reynolds number cylinder flow. It is likely that most of the fluctuating kinetic energy is contained in a vortex shedding mode, which forms the largest contribution to the Reynolds stresses needed to sustain the mean profile.

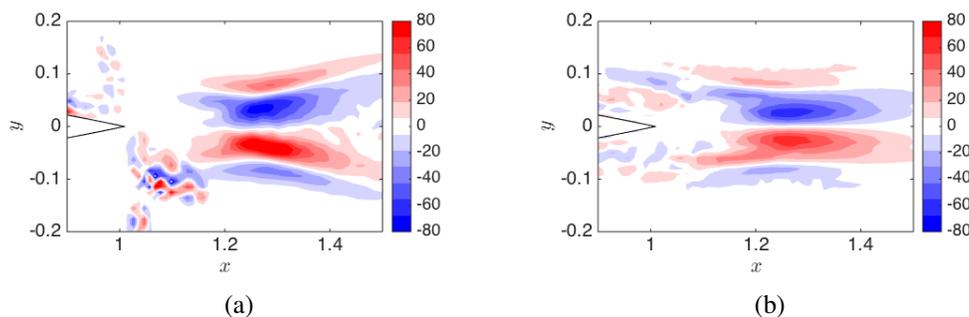


Figure 6.13: A comparison between the curl of the forcing $\nabla \times \mathbf{f}$ of the experimental velocity field (left) and assimilated velocity field (right) for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$.

Chapter 7 will exploit the similarity between the two flows.

The curl of the forcing $\nabla \times \mathbf{f}$ is presented in Figure 6.13. There is good agreement between the two and again it can be remarked that the shape is quite similar to the cylinder case as seen in Figure 5.9. In comparison to the idealized airfoil case, there is less noise in Figure 6.13(a) in the wake ($x \in [1, 2]$) where comparisons between the experiment and data-assimilation can be made. Over the airfoil surface, which has not been plotted, the experimental data contain a lot of noise since it cannot resolve the sharp gradients of the fluctuating velocity fields.

6.5.2 $Re = 20700$

The mean velocity field at $Re = 20700$ is also assimilated for the flow around a NACA 0018 airfoil. The data-assimilated streamwise and transverse mean velocities are compared to the PIV results in Figure 6.14. While the overall agreement is good, the algorithm struggles to reconstruct the flow near the airfoil surface. This is particularly evident when viewing the blue contours in Figure 6.14(b) for $0.5 < x < 1.0$. The resolution of the PIV is not sufficient to resolve low-speed regions. There are discrepancies when viewing the transverse velocity component as well which predicts stronger flow towards the body than the true mean. The impact of this discrepancy on the resolvent modes is analyzed in Chapter 7.

Here, the impact on the curl of the reconstructed forcing $\nabla \times \mathbf{f}$ is considered in Figure 6.15. f_x and f_y have not been plotted in the interest of brevity. The magnitudes and overall shapes agree between the PIV and the experiment, but the match is not as good as it was for $Re = 10250$ in Figure 6.13. The data-assimilated result is less smooth than its experimental counterpart with sharper contour lines and multiple lo-

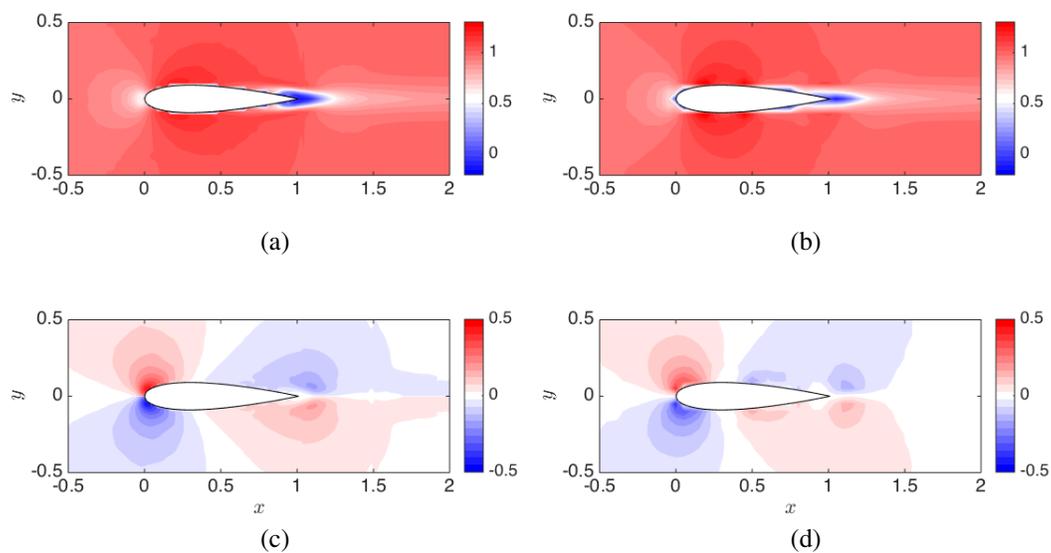


Figure 6.14: A comparison between the mean flow (left) obtained from PIV and the data-assimilated flow (right) around a NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 20700$. The flows are visualized by contours of the streamwise velocity (top) and transverse velocity (bottom).

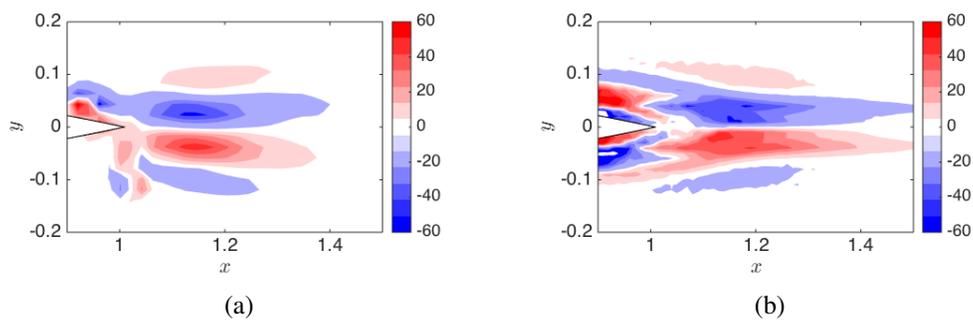


Figure 6.15: A comparison between the curl of the forcing $\nabla \times \mathbf{f}$ of the experimental velocity field (left) and assimilated velocity field (right) for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 20700$.

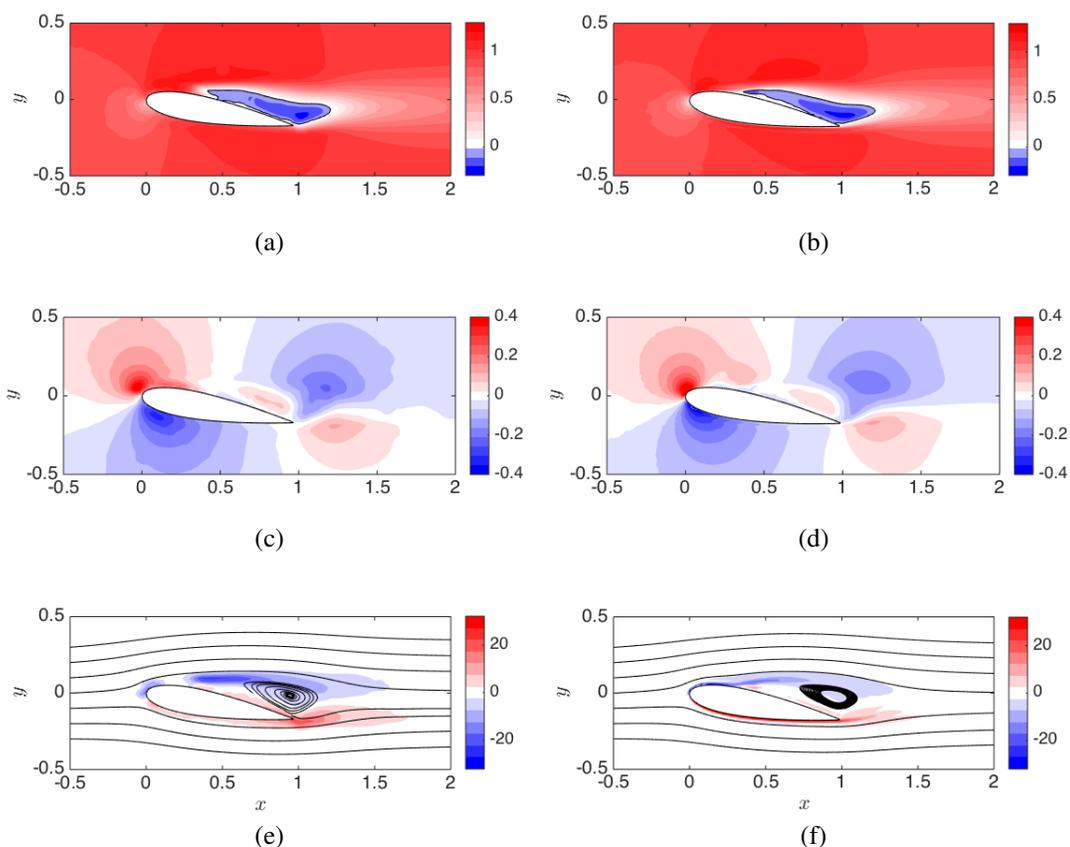


Figure 6.16: A comparison between the experimental (left) and data-assimilated (right) means for the NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$. There are contours of the streamwise velocity in (a,b), transverse velocity in (c,d), and vorticity in (e,f). The $\bar{u} = 0$ contours is in black for panels (a,b) while streamlines are demarcated by black lines in panels (e,f).

cal minima and maxima rather than a global minimum and maximum at $x = 1.15$, $y = \pm 0.03$. The results could be improved by collecting experimental data at a higher resolution on a smaller domain behind the airfoil since this coincides with the spatial support of the Reynolds stress gradients.

6.6 $\alpha = 10^\circ$ Case

The data-assimilated mean flow for the $\alpha = 10^\circ$ and $Re = 10250$ case is presented in Figure 6.16. Similar to the idealized airfoil, Figure 6.16 includes contours of the vorticity in panels (e,f) and streamlines to point out the size and location of the recirculation bubble. The contours of the streamwise and transverse velocities match very well particularly with respect to the reverse flow region as seen in Figure 6.16(b) and the structure of the transverse velocity contours in Figure 6.16(d).

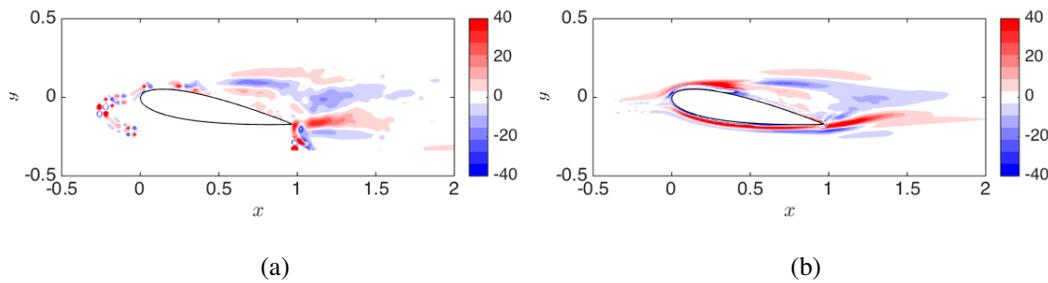


Figure 6.17: A comparison between the curl of the forcing $\nabla \times \mathbf{f}$ of the experimental velocity field (left) and assimilated velocity field (right) for the NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$.

Even the low-speed transverse velocity contours in the separated flow region immediately above the airfoil are recovered by the algorithm.

Similar to the idealized airfoil case, the individual components of \mathbf{f} are difficult to compare since the flow is highly separated. Therefore, only the curl of the reconstructed forcing $\nabla \times \mathbf{f}$ is plotted in Figure 6.17. The experimental result in panel (a) consists of the PIV data for $\alpha = 10^\circ$ only so the pressure side data are not included. Since the PIV data are so poorly resolved near the airfoil surface on the pressure side, they add no meaningful information to Figure 6.17 and have not been plotted. There are three alternating ‘bands’ of positive-negative-positive contours which coincide with the shear layer. This structure is similar to that of the idealized airfoil, and the agreement between the PIV and assimilation is good in this region. Immediately behind the airfoil, there is strong positive forcing in both the PIV and data-assimilated fields and patches of weaker, negative forcing above and below.

Finally, it is worth remarking that the data-assimilated field is more smooth than its data-assimilation counterpart. The same could be said for the idealized airfoil case earlier. The PIV results for the $\alpha = 0^\circ$ cases, however, are slightly more smooth than their data-assimilation counterparts even though it is difficult to access Reynolds stress information in regions near the airfoil surface. It can be concluded that the data-assimilation algorithm works slightly better for separated flows where the gradients of the mean velocity field are less severe. When the flow is attached, as it is for the $\alpha = 0^\circ$ cases, the algorithm has a harder time data-assimilating the flow at the airfoil surface and so the quality of the reconstructed Reynolds stress gradients suffers as seen for the $\alpha = 0^\circ$, $Re = 20700$ case.

As noted by Raffel et al. (2018), there are difficulties associated with PIV for lam-

inar flow around airfoils. The effects of strong centrifugal forces around the airfoil leading edge and strong shear result in the outward movement of tracer particles in a direction perpendicular to the curved streamlines. This loss of seeding compromises measurements in the region close to the wall. In Figure 6.16(a), for example, the separation point from the PIV is predicted to occur at $x \approx 0.4$ which is too far downstream. The loss of seeding also noticeably impacts Figure 6.16(e) where the contours of negative vorticity are separated in the region $0.1 < x < 0.25$. The data-assimilation results yield a significantly better prediction of the separation point and can fill in data near the leading edge where the PIV data cannot be trusted.

6.7 Summary

The data-assimilation framework based on Foures et al. (2014) is adapted to recover the mean flow and unknown momentum forcing around an idealized airfoil at a Reynolds number of $Re = 13500$. It is also applied to flows around a NACA 0018 airfoil at two angles of attack and two Reynolds numbers. The experimental data sets originate from time-resolved PIV data on a uniform Cartesian grid with a spatial resolution far lower than that of the numerical simulation. It is possible, nevertheless, to compute an assimilated flow field which is in very good agreement with the experiment. The recovered forcing fields are also in good agreement with their experimental counterparts and generally contain less noise since the experimental forcing is computed by differentiating Reynolds stress fields.

Three modifications to the original framework are necessary for practical implementation of the algorithm. First, it is necessary to use a smoothing procedure to ensure that the adjoint equations are being forced at all mesh points instead of just those which lie closest to the experimental grid points. Second, discrepancy velocity measurements are computed above the centerline only, when the flow is symmetric, and are reflected to account for the shadow cast by the airfoil. For nonzero angles of attack, the means for $\pm\alpha$ are obtained and stitched together. Finally, it is necessary to begin with a base flow computed at a lower Reynolds number than that of the experiment due to the difficulty of computing a base flow at Reynolds numbers much higher than the critical one. Once the data-assimilation process is complete for a lower Reynolds number, it can be incrementally increased until it matches the experimental conditions.

As expected, the full field case where all PIV vectors are used to guide the mean flow reconstruction yields the best reconstructed fields. It is encouraging, never-

theless, that truncating the PIV field of view or decreasing the resolution does not significantly impact the quality of the results. Improved smoothing procedures or weighting measurement points where important mean flow features such as the recirculation bubble are present in the flow could help reduce the number of necessary points further. The quality of the results could also be improved by implementing bounds on how close the numerical values match the experiment via the fitting criterion to account for measurement uncertainty.

Chapter 7

FLOW RECONSTRUCTION

All of the tools from the previous chapters are combined here to reconstruct unsteady flows realized in the laboratory using the smallest number of measurements possible. The schematic in Figure 7.1, which is an expanded version of the general approach shown in Figure 1.5, outlines the procedure which is followed in this chapter. The original contributions to this framework, which was first introduced by Gómez et al. (2016a) and Bennedine et al. (2016), include treatment of the nonlinear term, data-assimilation, and obtaining mean pressure. To begin with, the procedure introduced in Chapter 1 is described in greater detail, highlighting the contributions from the previous chapters. Next, the flow is reconstructed for a NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$. As noted in Chapter 6, this flow is quite similar to the low Reynolds number cylinder flow studied in Chapter 4 in as much as the airfoil experiences vortex shedding resulting in very similar Reynolds stresses. Preliminary results for the resolvent analysis of the data-assimilated $\alpha = 10^\circ$ case appear in Appendix B as the flow reconstruction procedure has not yet been applied to this flow. The benefits of data-assimilating the mean profile, however, are much more significant for this angle of attack than the one presented in this chapter.

7.1 Reconstruction Procedure

The procedure, as illustrated in Figure 7.1, begins by collecting the experimental data (top row in the figure). In the general case, a ‘crude’ mean profile is obtained using non-time-resolved PIV, for example, as well as a single sensor which contains time-resolved information such as a hot-wire or a pitot tube. The specific reconstruction performed here uses the mean profile obtained from PIV (see Chapter 2) and since the data are time-resolved, the fluctuations are known at each PIV point. The location of the probe point, furthermore, is guided by Chapter 3 where the spatial support of the most energetic resolvent modes can be predicted by features of the mean profile. The resolution and field of view with respect to the PIV have been guided from the results in Chapters 5 and 6. The data-assimilated velocity field is the input for resolvent analysis and a discrete Fourier-transform of the probe velocity $\check{v}(\mathbf{x}_0)$ is used to compute the local frequency spectrum (second row). It should be noted that the v -component of velocity is being measured for this flow

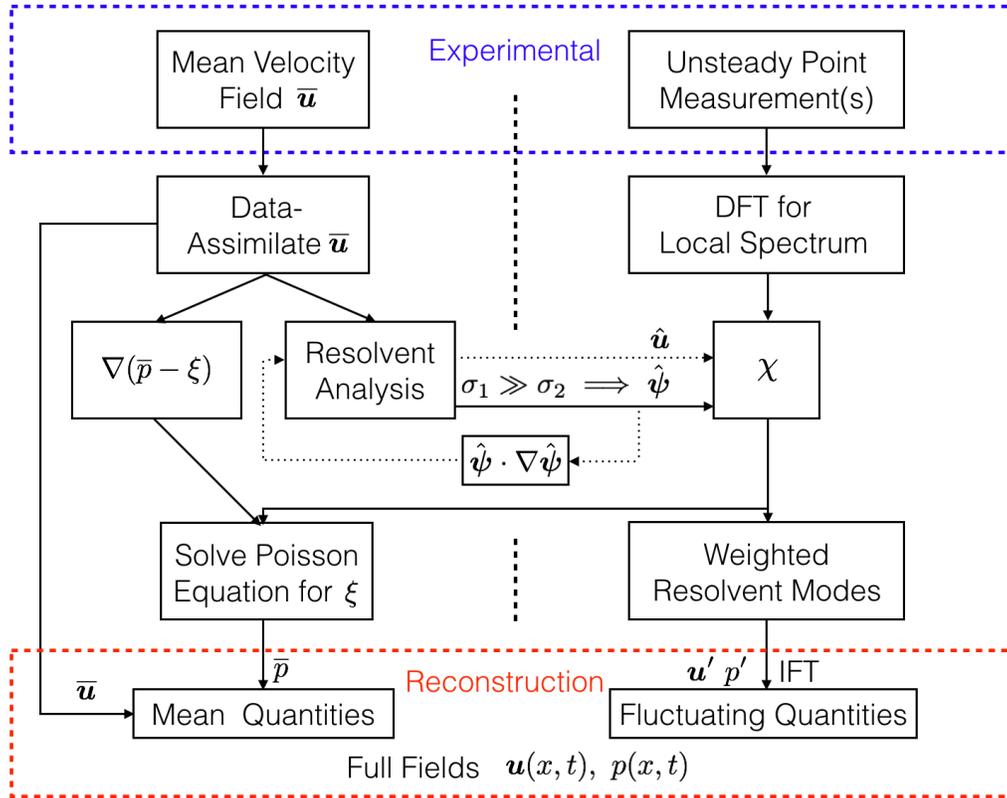


Figure 7.1: Schematic of the flow reconstruction procedure used in this chapter. Experimental quantities are contained in the dashed blue box while the reconstructed variables are contained in the dashed red box.

instead of u since the spectral content is easier to discern. For the general case, the signal can be obtained from any flow quantity.

The quantity $\nabla(\bar{p} - \xi)$ is another output from the data-assimilation. In order to isolate \bar{p} , it is first necessary to obtain information about the fluctuations. Since the unsteady velocity field can be represented in the frequency domain by its Fourier modes, the next step is to evaluate the frequencies at which proportionality can be expected between the dominant resolvent and Fourier modes. If $\sigma_1 \gg \sigma_2$, the dominant resolvent mode $\hat{\psi}_1$ is used in the reconstruction with the caveat from Chapter 4 that it is $\sigma_i \bar{\chi}_i$ which determines the energetic contribution of a given resolvent mode. Since there are temporal frequencies which contribute energy to the flow where this assumption does not hold, e.g., higher frequency harmonics, the nonlinear forcing is approximated by triadic interactions of resolvent modes which sum to the desired frequency (see Chapter 4). Once the correct mode shapes are obtained, their complex amplitude χ is computed using Equation 2.65 (third row).

Once the resolvent modes have been weighted, a Poisson equation for ξ (Equation 2.63) can be solved to obtain \bar{p} . The fluctuating quantities u' and p' can be recovered by performing an inverse Fourier-transform of the weighted resolvent modes which are an approximation of the Fourier modes (fourth row). The final result is an improved approximation of both the mean and fluctuating states using a rudimentary mean profile and a unsteady point measurement (fifth row). All quantities are derived from the NSE and obey the continuity equation.

7.2 Resolvent Analysis Frequency Sweep

Since the first two rows of Figure 7.1 have been completed in Chapter 6, the next step is to perform a resolvent analysis of the data-assimilated mean profile. These results are compared with interpolating the mean profile obtained from PIV onto the same FreeFem mesh. This is accomplished using linear interpolation in *matlab* and manually enforcing the no-slip boundary condition on the airfoil surface since PIV data do not exist on these points. In regions of the mesh where PIV data do not exist such as at the inlet and outlet, the flow is set to $\bar{\mathbf{u}} = [1, 0]$. The number of PIV vectors is approximately 25,000, which is comparable with the 50,000 FreeFem mesh points. However, the FreeFem mesh is significantly more resolved near the airfoil than the PIV grid while the converse is true away from the body.

In Figure 7.2(a-b), the first three singular values of the resolvent operator are plotted for $0 < \omega < 35$. The shape for both the experimental and data-assimilated profiles is quite similar to the cylinder, in that there is a distinct range of frequencies where the first singular value is an order of magnitude higher than the second singular value. The range where all three singular values are the same order of magnitude is also similar. There are some differences between the interpolated versus the assimilated results in Figure 7.2. Panels (c) and (d) are magnified versions of (a) and (b), respectively, which show a discrepancy with respect to the globally most amplified frequency ω_g . In the interpolated case, this maximum occurs around $\omega = 11.9$ while, in the data-assimilated case, $\omega_g = 12.24$. It will be seen, based on the DMD eigenspectrum as well as the power spectra at various points in the wake, that the data-assimilated mean predicts ω_g more accurately.

There are other differences between Figure 7.2(a) and (b). The singular values are generally higher for the data-assimilated case suggesting that the eigenvalues are closer to the imaginary axis. While the separation between the singular values is generally the same, the interpolated case predicts greater separation between the

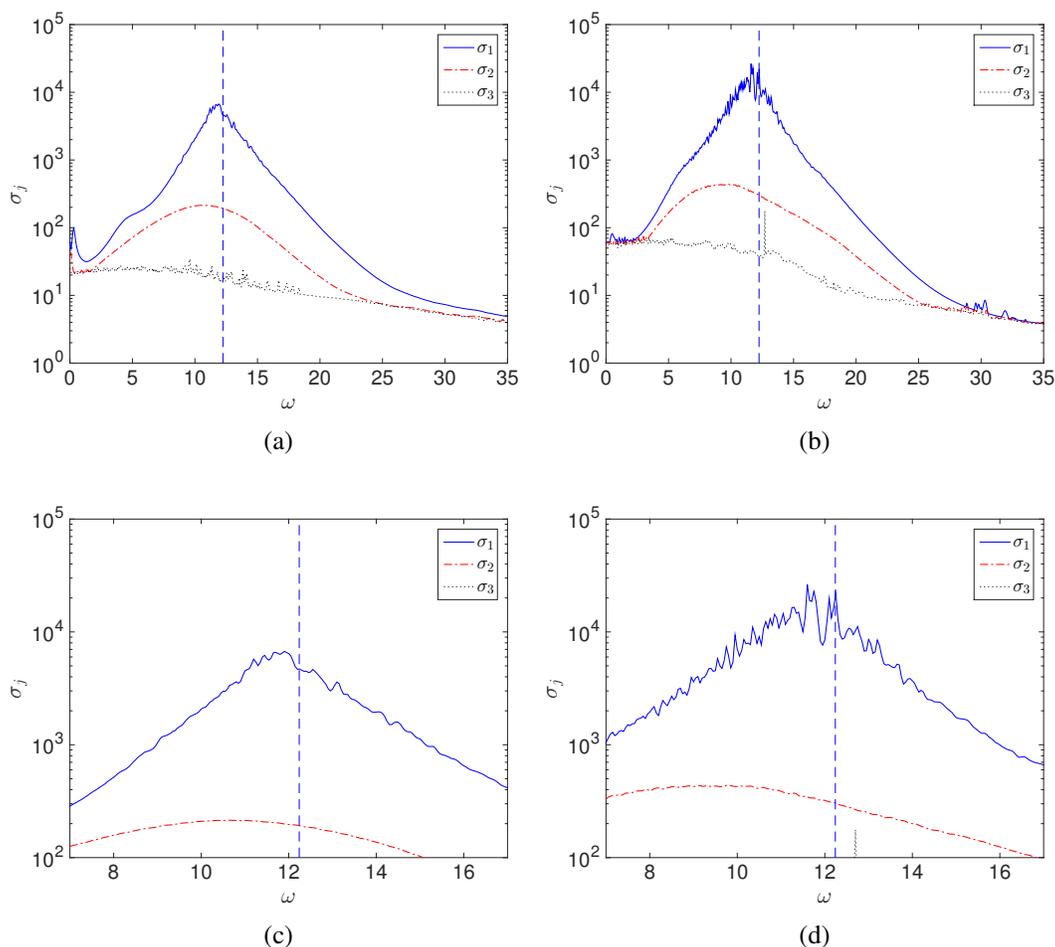


Figure 7.2: First three singular values of the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$ using the (a) interpolated experimental mean and the (b) data-assimilated mean. (c) and (d) are zoomed in versions of (a) and (b), respectively. The vertical, blue dashed line indicates the globally most amplified frequency $\omega_g = 12.24$.

singular values, particularly between σ_1 and σ_2 at very low frequencies ($\omega < 3$) and very high frequencies ($\omega > 30$). The data-assimilated results, as will be seen in Section 7.3, more accurately predict the range of frequencies where the reconstruction is good using a single resolvent mode. In comparison to the interpolation, the resolvent norm is slightly more jagged near the peak. One potential explanation for this behavior is that the shedding frequency is not completely constant in time, and this may be reflected in the broadening of the peak in the frequency range of $10.3 < \omega < 12.3$.

The resolvent forcing modes at $\omega_g = 12.24$ are plotted in Figure 7.3 for the experimental and data-assimilated profiles. Even though the resolvent forcing modes

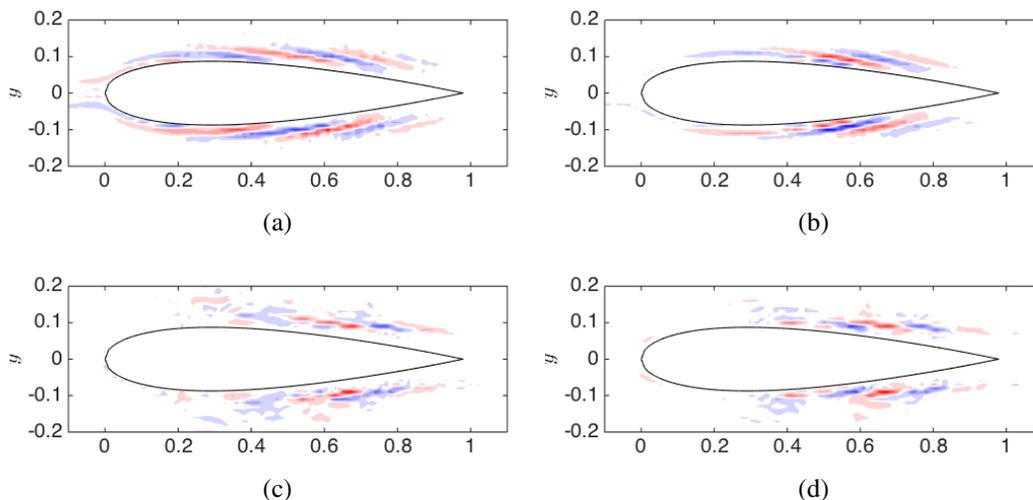


Figure 7.3: Comparison of $\hat{\phi}_1$ for $\omega_g = 12.24$ of the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$. The modes computed from the interpolated experimental profile are shown on the left while those from the data-assimilated profile are on the right. The u -component of the mode is the top row while the v -component of the mode is the bottom row. The modes have been normalized such that the maximum contours are unity.

are similar, the structure is captured slightly better by the data-assimilated profile. The resolvent response modes (which have not been plotted to avoid redundancy), on the other hand, are nearly identical suggesting that the wake dictates the shape of those modes. The PIV data are able to resolve well the wake behind the airfoil hence a good prediction of the most amplified response mode. Due to the convective non-normality of this flow, the forcing modes are upstream and quite close to the airfoil surface since the flow is attached. Because there are only two PIV vectors in the cross stream direction, the interpolation onto the FreeFem mesh, which has a resolution five times larger, is not likely to be very accurate. Furthermore, the no-slip boundary condition has to be manually enforced since the interpolated values are not zero. The data-assimilated profile, on the other hand, does not need to interpolate between mesh points and yields a slightly improved prediction of the forcing mode.

Even if there is a lack of experimental data near the airfoil surface, it is still important for the mesh, onto which the experimental data is interpolated or the data-assimilation equations are solved, to be highly resolved. The spatial support of the forcing mode is confined to a very narrow strip just above the airfoil. Consequently, if the mesh resolution is too sparse, it affects the resolution of the forcing mode and

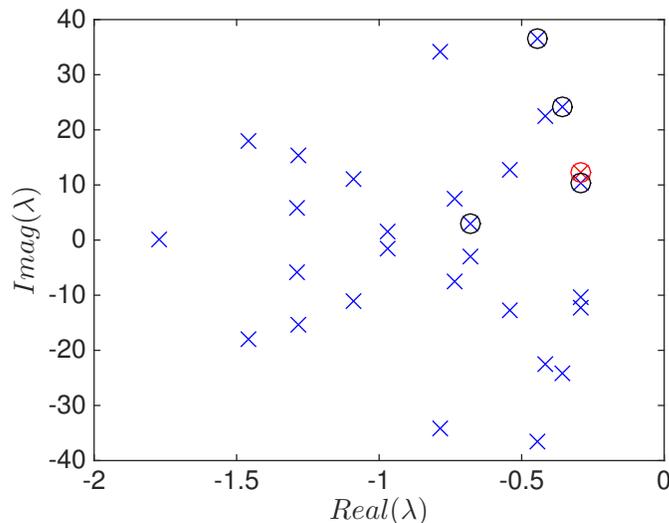


Figure 7.4: DMD eigenvalues of the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$. Only 29 have been plotted as crosses. The eigenvectors of the circled eigenvalues are compared to their resolvent mode counterparts in subsequent figures. The red and circled eigenvalue corresponds to $\omega_g = 12.24$.

eventually the resolvent response mode.

7.3 Comparison Between Resolvent and DMD Modes

Similar to the cylinder case in Chapter 4, the resolvent modes can be compared to their DMD counterparts. As explained by Towne et al. (2018), the two should match when the resolvent operator is low-rank and the forcing is unstructured. The DMD eigenvalues are plotted in Figure 7.4 and the imaginary part of the closest eigenvalue to the globally most amplified frequency is $\omega_g = 12.24$. Notably, this DMD eigenvalue is closest to the imaginary axis. Both velocity components of the resolvent and DMD mode are plotted in Figure 7.5. The agreement is relatively good in terms of the spatial structure although the amplitude of the DMD mode is highest nearest the trailing edge of the body, especially for the streamwise velocity component. These results are reminiscent of those for the cylinder in Chapter 4.

In the interest of brevity, only the v -component of the resolvent and DMD modes is plotted in Figure 7.6 for three other temporal frequencies: one in the vicinity of the globally most amplified frequency ($\omega = 10.39$), one which is (approximately) the first harmonic of the globally most amplified frequency ($\omega = 24.49$), and one at a low temporal frequency ($\omega = 2.90$) where the resolvent operator is not low-rank. There is some separation between σ_1 and σ_2 for the first harmonic mode

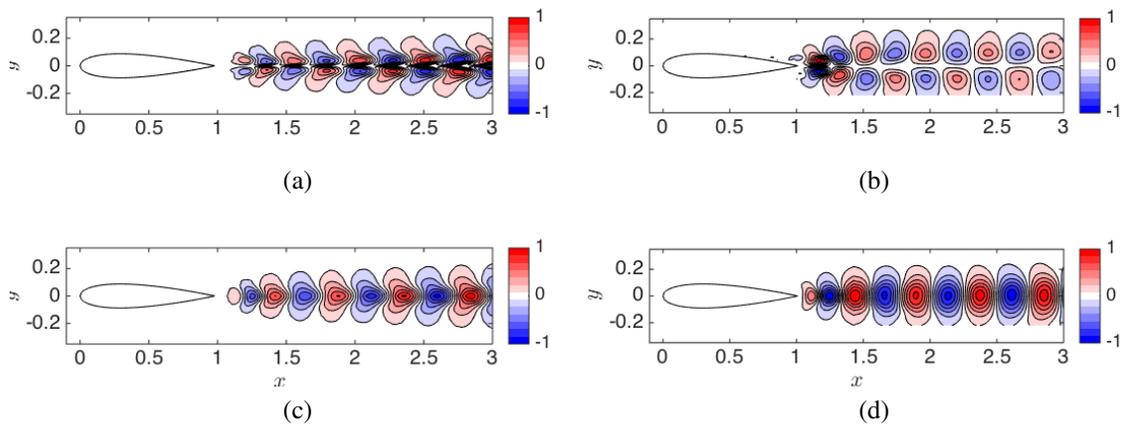


Figure 7.5: Comparison of $\hat{\psi}_1$ (left) and its DMD counterpart (right) for the globally most amplified frequency of the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$. The u -component of the mode is the top row while the v -component of the mode is the bottom row. The modes have been normalized such that the maximum contours are unity.

although it is only a factor of two whereas the separation for the globally most amplified frequency is more than one order of magnitude. The $\omega = 10.39$ mode is of particular interest since the DMD spectrum shows that this mode is quite close to the imaginary axis. It is hypothesized that this mode is physically significant and accounts for the jittery behavior of σ_1 observed in Figure 7.2.

There is agreement between the resolvent and DMD modes for $\omega = 10.39$ only; there are major discrepancies for the other frequencies. The mode predicted from resolvent analysis for $\omega = 24.49$ resembles the shedding mode which has been compressed in the streamwise direction. DMD identifies an anti-symmetric mode that mirrors the $2\omega_s$ mode observed in the cylinder wake (see Chapter 4). The DMD mode at $\omega = 2.90$ has a similar structure to the high frequency mode although it is a far more stretched version in the streamwise direction. The resolvent mode, on the other hand, still resembles the shedding mode, albeit a more compressed version.

It may not be possible to compare resolvent modes with DMD modes due to lack of data or insufficient time resolution. In scenarios where DMD modes are not available, the degree of separation between σ_1 and σ_2 indicates whether it is appropriate to use resolvent modes. If $\sigma_1/\sigma_2 > 10$, it suggests the influence of a strong linear mechanism. The structure of the response, therefore, is insensitive to the structure of the nonlinear forcing. This has been observed not only for the flows in this thesis, but also for wall-bounded turbulence (McKeon and Sharma, 2010), backward

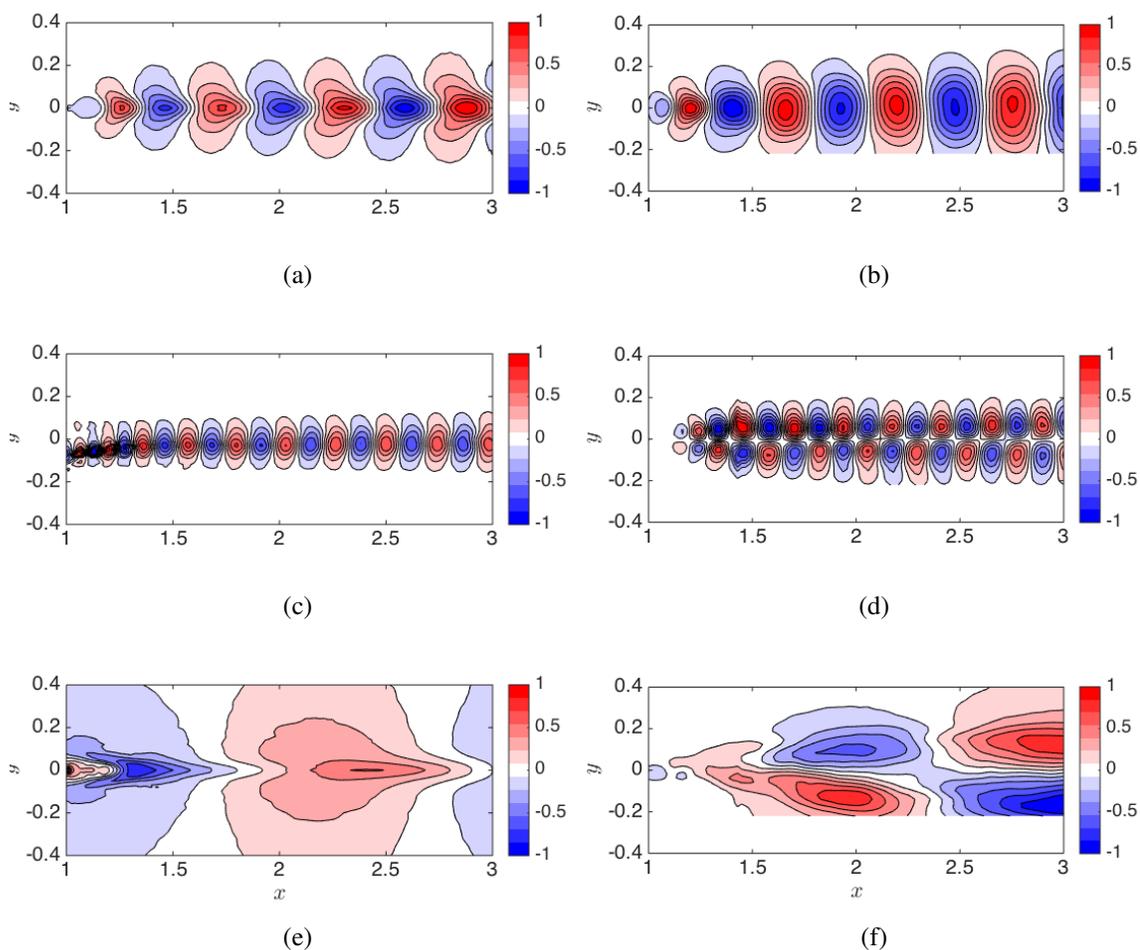


Figure 7.6: Resolvent (left) and DMD (right) mode comparison at $\omega = 10.39$ (a-b), $\omega = 24.49$ (c-d), and $\omega = 2.90$ (e-f) for the NACA 0018 airfoil where $\alpha = 0^\circ$ and $Re = 10250$. Only the v -component is plotted and the modes are normalized such that the maximum contours are unity.

facing steps (Beneddine et al., 2016), and turbulent jets (Schmidt et al., 2017b). In cases where this separation does not exist, the nonlinear forcing can be approximated. Even if the ‘true’ mode shape in the flow looks like $\hat{\psi}_1$, the forced resolvent mode will match $\hat{\psi}_1$. This was observed for ω_s in Chapter 4 when the resolvent was forced by $\hat{\psi}_1(-\omega_s) \cdot \nabla \hat{\mathbf{u}}(2\omega_s) + \hat{\mathbf{u}}(2\omega_s) \cdot \nabla \hat{\psi}_1(-\omega_s)$. There is nothing to lose, therefore, by being conservative with the choice of the σ_1/σ_2 threshold.

7.4 Nonlinear Interactions for Low and High Frequency Modes

There is a transition in the v -component of the DMD modes from symmetric to anti-symmetric which the resolvent modes are failing to capture for modes outside

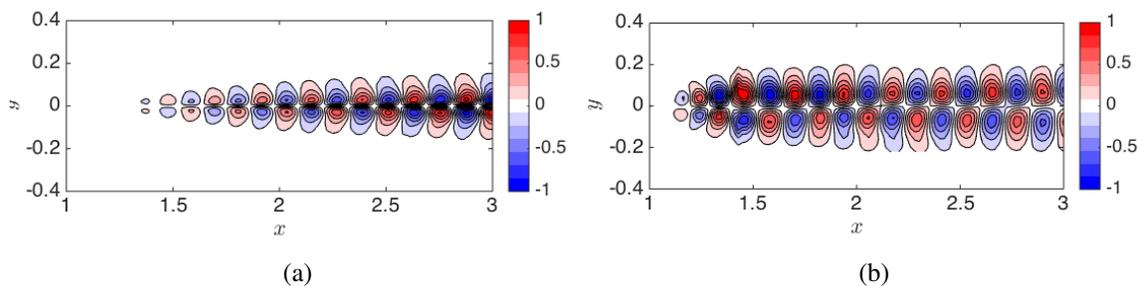


Figure 7.7: Resolvent mode (left) obtained from forcing the resolvent operator with the nonlinear interaction of two triadically consistent response modes at $\omega = 24.49$. The DMD counterpart is in (b). Only the v -component is plotted and the modes are normalized such that the maximum contours are unity.

the rank-1 range. The structure of the first harmonic DMD mode is quite similar to the cylinder case, so it is reasonable to suggest that the nonlinear forcing to this harmonic can be approximated by the self-interaction of the globally most amplified frequency. The nonlinear forcing is computed from

$$\hat{f}(\omega = 24.49) \approx \hat{\psi}_1(\omega = 12.24) \cdot \nabla \hat{\psi}_1(\omega = 12.24), \quad (7.1)$$

where other triads which sum to the same ω have been neglected. The resulting structure is compared to the DMD mode in Figure 7.7. There is much better agreement between this resolvent mode compared to the one obtained from a singular value decomposition of the resolvent operator. Notably the y antisymmetry is now captured. There are some differences in Figure 7.7 such as the thickness in the y -direction and the streamwise location of the mode. The resolvent mode is thinner, concentrated closer to the centerline, and has non-negligible amplitude further downstream than its DMD counterpart. These differences are likely symptoms of the differences observed between the resolvent and DMD modes at the globally most amplified frequency since they were used to compute the nonlinear forcing which generated these modes. The results in Figure 4.18 substantiate this theory since it was observed that the nonlinear forcing computed from the DMD modes was almost an exact match to the nonlinear forcing obtained directly from the DNS.

Up until this point, sums of triads, or modes with $\omega_1 + \omega_2 = \omega_3$, have been considered to obtain the structure at very high frequencies relative to the dominant instability mechanism. The same principle can be applied for very low frequency modes by considering the differences of triads, or $\omega_1 - \omega_2 = \omega_3$. To obtain a closer prediction of the structure at $\omega = 2.90$, the nonlinear forcing $\hat{f}(\omega = 2.90)$ is

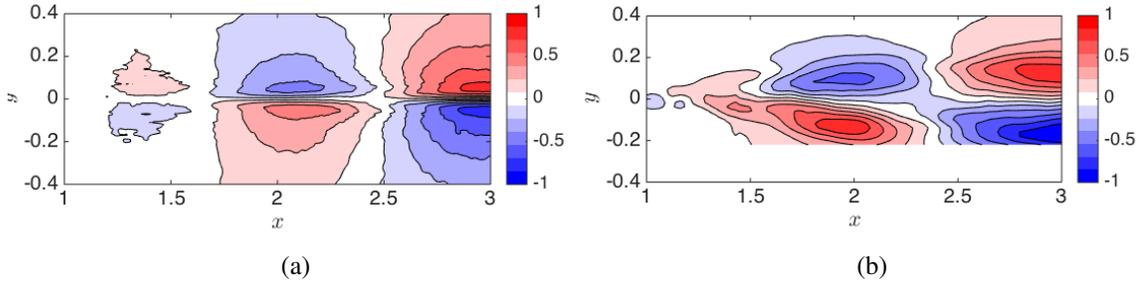


Figure 7.8: Resolvent mode (left) obtained from forcing the resolvent operator with the nonlinear interaction of two triadically consistent response modes at $\omega = 2.90$. The DMD counterpart is in (b). Only the v -component is plotted and the modes are normalized such that the maximum contours are unity.

proposed:

$$\hat{\mathbf{f}}(\omega = 2.90) \approx \hat{\boldsymbol{\psi}}_1(\omega = 12.24) \cdot \nabla \hat{\boldsymbol{\psi}}_1(\omega = -9.06) + \hat{\boldsymbol{\psi}}_1(\omega = -9.06) \cdot \nabla \hat{\boldsymbol{\psi}}_1(\omega = 12.24). \quad (7.2)$$

The velocity response from the resolvent is compared with the DMD mode in Figure 7.8. Once again, there is qualitative agreement between the general spatial variation of the structures obtained and the differences observed are similar to those for the $\omega = 24.49$ modes. More triads would need to be considered to recover the complete nonlinear forcing and therefore the most accurate resolvent mode.

The question of how many triads are necessary to recover an acceptable approximation of the nonlinear forcing is beyond the scope of this thesis, as is the consideration of recovering suboptimal resolvent response modes. The preliminary results which are presented here rely on the rank-1 approximation for frequencies at which σ_1 is sufficiently larger than σ_2 . For frequencies where this assumption does not hold, the nonlinear forcing is approximated using two response modes whose frequencies lie in the range where rank-1 is acceptable. The addition of frequencies recovers modes at the higher end while the differences cover the lower end. It can be hypothesized that the structure of the modes at very high frequencies can be approximated by considering a nonlinear forcing computed from

$$\hat{\mathbf{f}}(\omega_3) \approx \hat{\boldsymbol{\psi}}_1(\omega_1) \cdot \nabla \hat{\mathbf{u}}(\omega_2) + \hat{\mathbf{u}}(\omega_2) \cdot \nabla \hat{\boldsymbol{\psi}}_1(\omega_1), \quad (7.3)$$

where $\omega_1 + \omega_2 = \omega_3$ and ω_3 is the frequency of interest. In the above equation, ω_1 is in the rank-1 range and ω_2 lies outside this range hence the use of the mode shape $\hat{\mathbf{u}}$ instead of $\hat{\boldsymbol{\psi}}_1$.

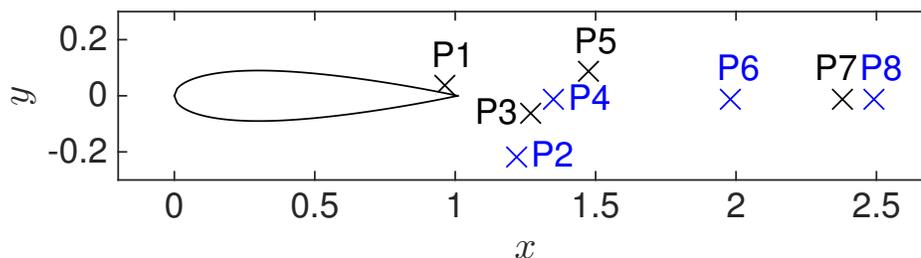


Figure 7.9: Probe points for NACA 0018 at $\alpha = 0^\circ$ and $Re = 10250$ flow reconstruction. The points are numbered from left to right (i.e. the first point is the furthest upstream) and the colors are used solely to avoid confusion for points which happen to be close together such as P3 and P4.

	Position	Description
P1	(0.97, 0.04)	Above Airfoil
P2	(1.22, -0.22)	Below Near Wake
P3	(1.27, -0.06)	Near Wake Below Centerline
P4	(1.35, 0)	Near Wake on Centerline
P5	(1.47, 0.09)	Near Wake Above Centerline
P6	(1.98, 0)	Far Wake on Centerline
P7	(2.38, 0)	Far Wake on Centerline
P8	(2.49, 0)	Far Wake on Centerline

Table 7.1: Points used in the NACA0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$ flow reconstruction.

7.5 Probe Measurements for Velocity Fluctuations

In order to calibrate the complex amplitudes of the mode shapes computed using either the singular value decomposition or forcing the resolvent operator, a time-resolved probe point needs to be selected. Multiple locations may be chosen but this thesis will consider knowledge at one point only. The performance using several points, which are plotted in Figure 7.9 and tabulated in Table 7.1, is compared in Section 7.6. The ordering of the points is determined by their streamwise location. P1 corresponds to the point furthest upstream while P8 corresponds to the point furthest downstream. Only the v -component of the velocity is used to inform the amplitude and phase of the modes at each temporal frequency as the spectral signature is clearer than the u -component.

The power spectrum for two of the candidate points in Table 7.1 is plotted in Figure

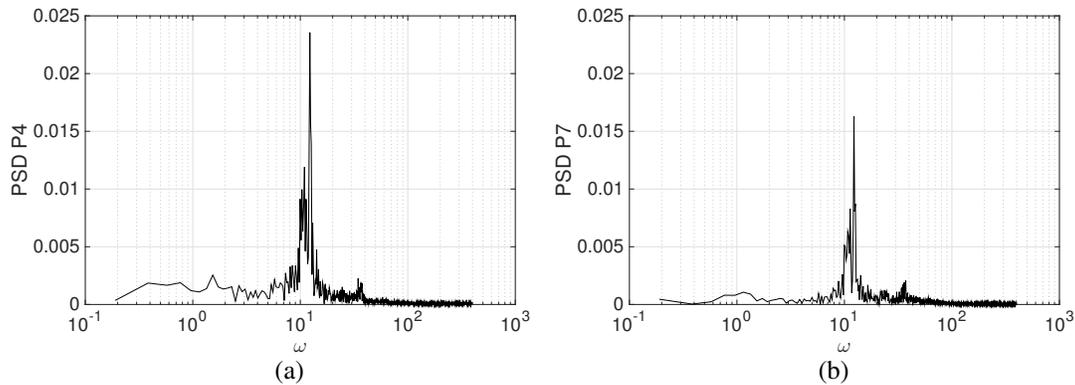


Figure 7.10: Power spectrum at points P4 and P7 in the airfoil wake

7.10. Even though the streamwise separation between P4 and P7 exceeds one chord length, the spectra are quite similar. It can be observed, perhaps more easily for P4, that there exist two main peaks which occur at $\omega = 10.35$ and $\omega = 12.27$. The latter corresponds to ω_g identified by the resolvent of the data-assimilated mean while the former is within the range where σ_1 is not smooth. It does not coincide with the peak of σ_2 which occurs at $\omega = 9.15$. The presence of two distinct frequencies in the flow leads to a beating effect which is explored in the next section. The main dissimilarity to remark upon in Figure 7.10 is the amplitude of the two peaks which decrease as one moves downstream from P4 to P7. There is a small peak for both P4 and P7 at $\omega = 36$, so the resolvent modes are computed for frequencies up to $\omega = 39$.

7.6 Unsteady Flow Reconstruction

In this section, the fluctuating velocity is reconstructed by computing the complex weights of the resolvent mode at each temporal frequency. Doing so yields predictions of the spectra at other points in the flow as shown by Beneddine et al. (2016) and Thomareis (2017). The inverse Fourier-transform of the weighted modes yields the unsteady velocity and pressure fields (Gómez et al., 2016a; Beneddine et al., 2017). The results that are presented here focus on the flow reconstruction aspect although it is possible to make predictions about the spectra at other points in the flow. The quality of the reconstructed fluctuations is evaluated by comparing the unsteady velocity fields with the PIV snapshots. Note that these PIV velocity fields are not constrained to obey continuity, so exact agreement is not necessarily expected. Since pressure data are not available for comparison, the unsteady pressure field results are omitted although the method is also capable of reconstructing the

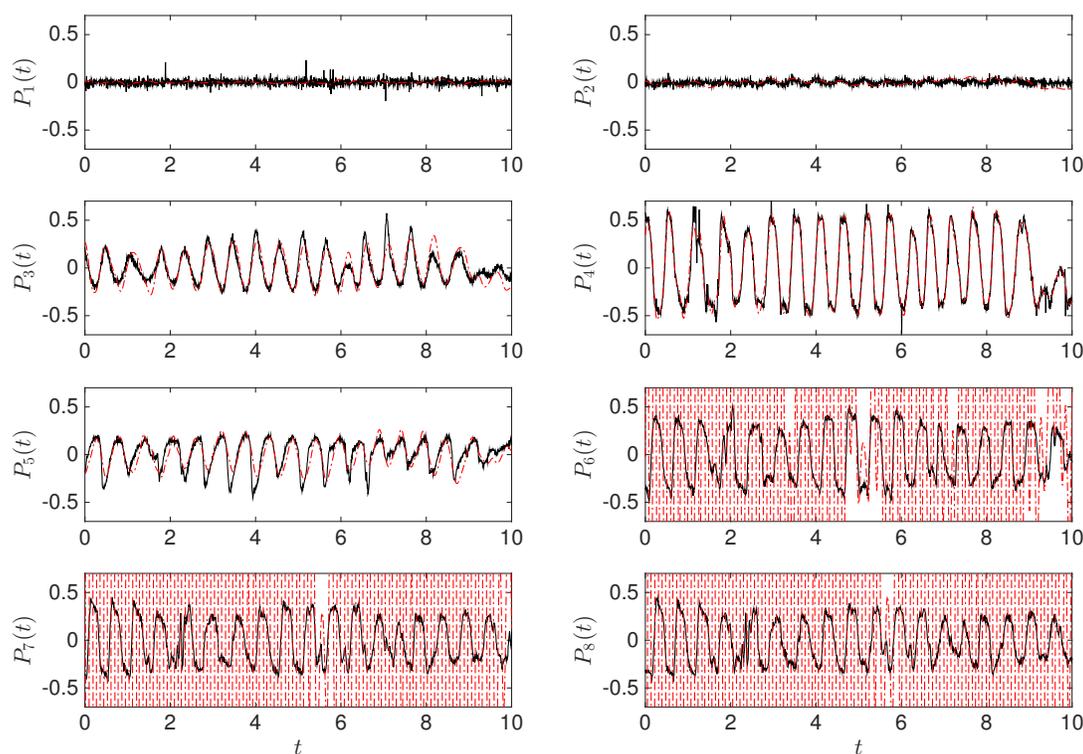


Figure 7.11: Time traces of the reconstructed v -component when the probe point is located at P4. The solid black lines are the raw PIV data while the dashed-dotted red lines are the reconstruction.

pressure fluctuations.

7.6.1 Selection of Probe Point and Point Reconstructions

The choice of the best probe point is examined here before comparing the flow reconstruction to the original PIV snapshots. As will be seen below, it is not necessary to introduce an error metric to make this judgment. The flow reconstruction uses frequencies in the range $0.22 < \omega < 38.83$. The results for probe point 4 are summarized in Figure 7.11 and only the first ten seconds of the data are plotted for greater clarity. Even though the calibrated amplitudes are based on the DFT of Point P4, it can be concluded that Points P1 and P2 are not suitable choices for the probe point since the fluctuation energy is very low in these regions (see Figures 7.5 and 7.6). P1 is too far upstream while P2 lies outside the wake region. The probe point must lie inside the region where it can discern the velocity fluctuations that arise from the vortex shedding. The agreement between the reconstruction and

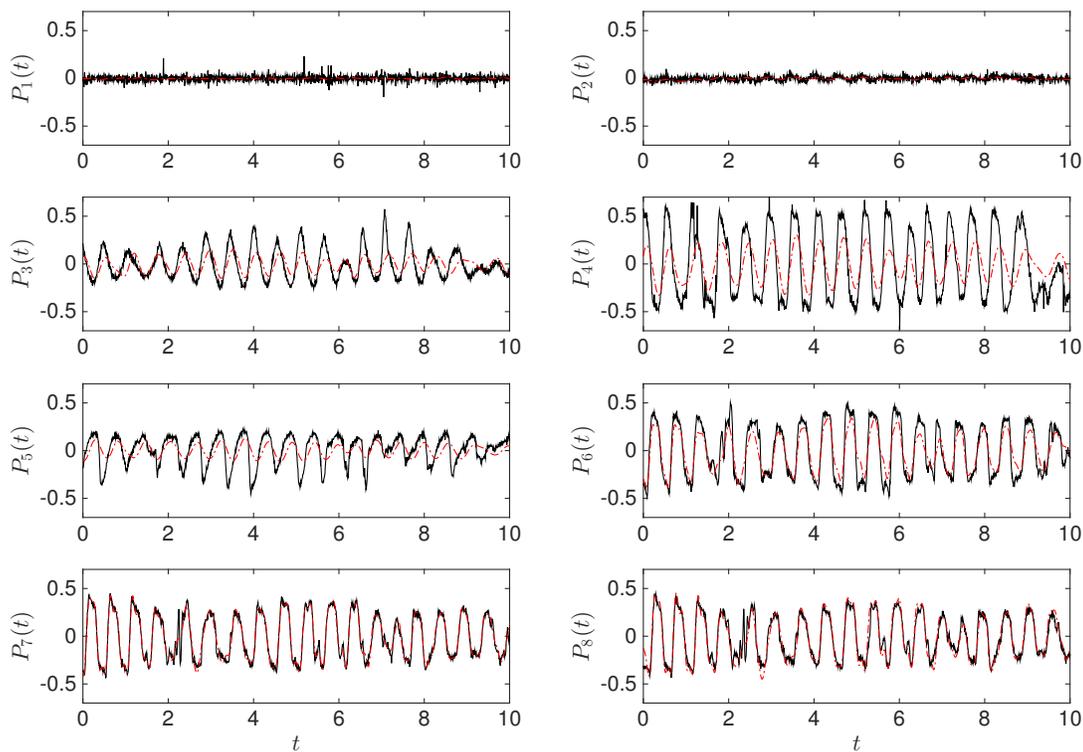


Figure 7.12: Time traces of the reconstructed v -component when the probe point is located at P7. The solid black lines are the raw PIV data while the dashed-dotted red lines are the reconstruction.

the data is good for Points P3-P5, particularly P4 since this is the reference point. The reconstruction for Points P6-P8 is not good as the amplitudes are significantly overestimated. This suggests that the choice of Points P3-P5, which all yield similar results, does not result in very successful flow reconstruction.

The results for probe point 7 are summarized in Figure 7.12. Once again, Points P1 and P2 are in regions of the flow where the fluctuations are weak, so the flow reconstruction ‘matches’ in the sense that it predicts no activity at these points. The agreement for P6-P8 is very good in terms of both the amplitude and phase of the reconstruction. On the other hand, the agreement at points P3-P5, which are further upstream, suffers when calibrating from a probe point downstream. While the phase of the reconstruction matches reasonably well with the original PIV data, the amplitudes are underestimated by a factor of two in some cases. The overall agreement, nevertheless, is far better for P7 than it is for P4 and so P7 will be

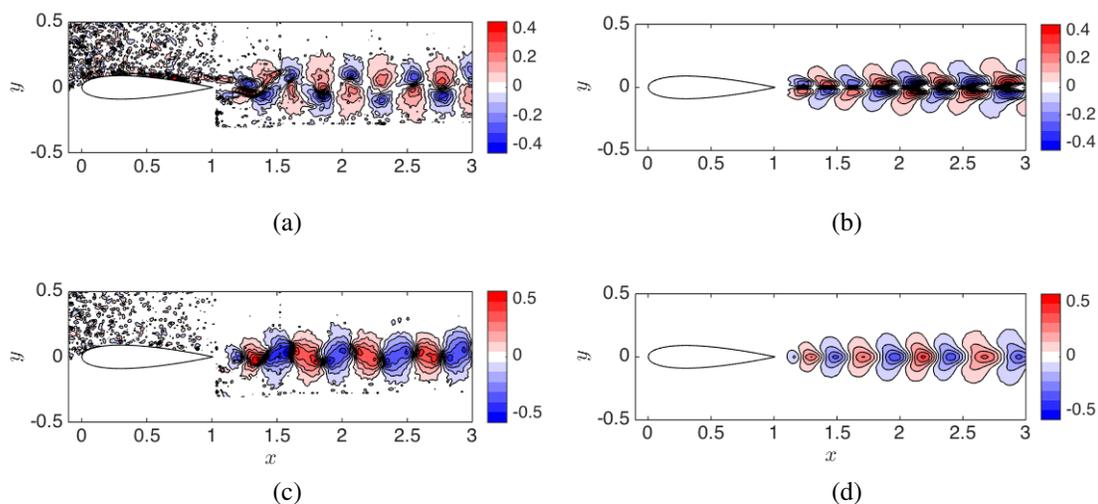


Figure 7.13: Reconstructed flow (right) at $t = 0$ using rank-1 modes for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$. The PIV snapshots without frequency filtering are on the left for comparison. The u -component is the top row while the v -component is the bottom row.

selected for the remainder of the chapter. Taking multiple points into account to calibrate the amplitudes of the modes is not considered in this thesis although it could lead to improved results (see Beneddine et al., 2017).

7.6.2 Full Field Approximation Using Rank-1 Frequencies Only

In this subsection, the flow is reconstructed using frequencies in the rank-1 range only ($3.1425 < \omega < 20.2017$). Recall that rank-1 validity as used here implies that the first resolvent response mode matches the structure of its DMD counterpart. Due to the dominance of the vortex shedding mode, the resolvent modes at less amplified frequencies in the rank-1 range tend to be a stretched or shrunken version of the globally most amplified mode in the streamwise direction. Both components of the reconstructed velocity fluctuations are plotted in Figure 7.13 and are compared to PIV snapshots at the same instant in time. The particular snapshot in Figure 7.13 is the very first, or time $t = 0$. Note that no filtering has been applied to the PIV data, i.e., all frequencies are included whereas the cutoff for the reconstruction is $\omega = 39$.

One of the most striking differences between the PIV snapshots and the reconstructed flow is the amount of noise. No effort has been made to mask it in order to emphasize one of the primary benefits of the flow reconstruction procedure: significantly less measurement noise. There is very good agreement between the

amplitudes of the modes for both velocity components. Both the height and width of the structures is also in good agreement. There are two main areas of disagreement, nonetheless. First is that the intensity of the fluctuations is under predicted immediately downstream of the airfoil. Second is that the shapes of the contours are not exactly the same. The PIV contours tend to be more square in shape and the maximum contour tilts in either the upstream or downstream direction. The shape of the resolvent contours, on the other hand, is symmetric with respect to the centerline. Adding more frequencies using the ‘forced’ resolvent modes results in closer agreement in the shape as will be seen in the next subsection. Overall, however, the bandwidth of frequencies included in the flow reconstruction, which obeys the NSE, yields an acceptable and clean representation of the main features of the instantaneous flow.

7.6.3 Adding ‘Forced’ Resolvent Modes

The flow reconstruction is now performed using all frequencies at which a resolvent mode is computed (i.e. $\omega < 38.73$). The motivation for choosing this cutoff stems from the power spectrum of Point P7 as seen in Figure 7.10, which has small peaks at harmonics of the globally most amplified frequency ω_g . The objective, therefore, is to compute modes just beyond $3\omega_g$ to recover as much of the fluctuating energy as possible. The flow can be reconstructed at every instant in time (only the first snapshot was considered in Figure 7.13 but only three will be presented in the interest of brevity).

The three snapshots correspond to $T_1 = 0$, $T_2 = 10$, and $T_3 = 20.4$ which are explicitly shown in Figure 7.14. T_1 is chosen to compare the effect of adding all frequencies versus just the rank-1 modes. T_2 corresponds to an extreme event where the flow changes considerably for a short amount of time. The amplitude of the velocity signal in Figure 7.14 decreases dramatically around T_2 . This type of behavior is reminiscent of beating between frequencies. Whether the reconstruction is capable of capturing this behavior will be answered shortly. Finally, T_3 is chosen as an instant where the phase of the velocity signal at P7 is opposite of the phase at T_1 .

The reconstructed flows are plotted and compared to the PIV snapshots in Figure 7.15. Only the v -component is considered as there is a greater discussion of the u -component results in Section 7.7. The T_1 snapshot in Figure 7.15(a,b) can be compared to the result in Figure 7.13(c,d). The qualitative agreement is notice-

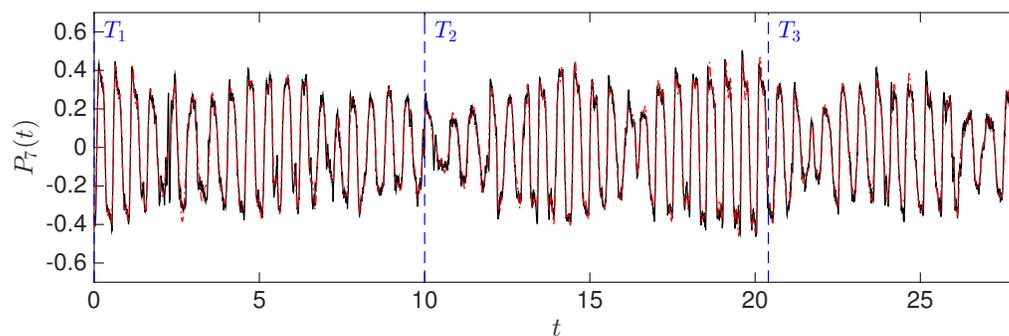


Figure 7.14: Fluctuating v -component at Point P7. The black solid line is the PIV measurement while the dashed-dotted red line is the reconstruction using modes for all frequencies considered. The dashed blue lines indicate the three instants in time where the PIV snapshots are compared to the reconstructed flow.

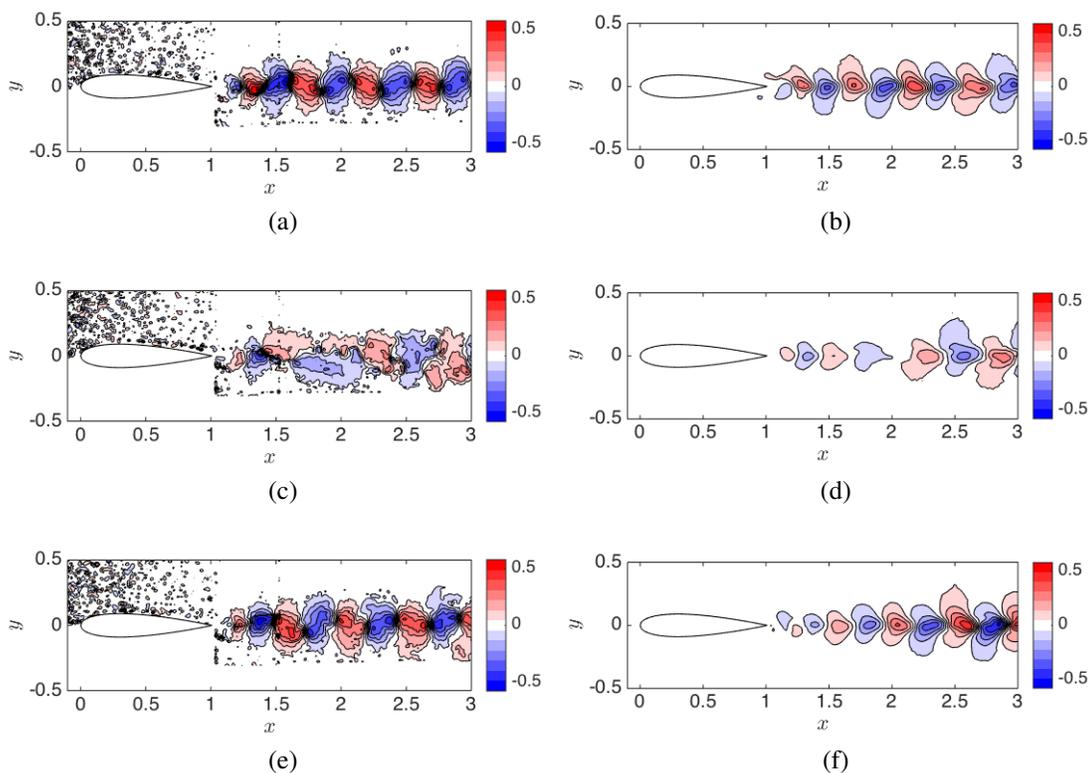


Figure 7.15: Reconstructed flow (right) for various times using all modes for the NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$ compared to the PIV snapshots (left). The snapshots corresponds to T_1 in (a,b), T_2 in panels (c,d), and T_3 in (e,f). Only the v -component has been plotted.

ably better when all frequencies are used as opposed to just the rank-1 frequencies. The shape of the contours in Figure 7.15(b) resemble their PIV counterparts more closely as the negative contours tend to tilt downstream while the positive contours tend to tilt upstream. Furthermore, there is less symmetry with respect to the centerline. The agreement of the mode amplitude in the region $1 < x < 1.5$ has not improved after adding more frequencies as the reconstruction still underestimates the strength of the fluctuations immediately behind the airfoil. The reconstruction also does well for the T_3 snapshot although it suffers from similar defects as the T_1 snapshot.

The T_2 snapshot was chosen intentionally to capture an extreme event in the flow rather than its ‘average’ behavior as represented by T_1 and T_3 . Remarkably, the flow reconstruction is able to approximate the flow’s behavior at this point despite its complexity. While the PIV data look somewhat disorganized and incoherent, the position of each contour’s centroid matches approximately with the flow reconstruction. There are fewer structures at T_2 compared to the other snapshots (seven versus nine) and this is captured by the flow reconstruction.

7.7 Discussion

Despite promising initial results, there are some drawbacks to this method which are discussed in this section. To begin with, a global error metric is defined to compare the effect of the probe location, the velocity component measured, and the number of frequencies accounted for in the model. The second topic of discussion is the influence of mode position and amplitude. It has been remarked on several occasions in this thesis that the resolvent response modes, forced or unforced, tend to have higher amplitudes downstream of the airfoil and closer to the centerline when compared to their DMD counterparts. These differences make it slightly more difficult for the model to converge to the PIV snapshots. The third topic of discussion is the reconstruction of the u -component using v measurements only. It is more difficult to reconstruct u for reasons which will be explored in this section.

7.7.1 Error Metric

Quantitative comparisons between the raw PIV snapshots and the reconstructed flow are based on the same error metrics as those outlined in Beneddine et al.

Point	Component	Frequencies	E_u	E_v
P7	u	rank-1	0.012	0.028
P7	u	all	4.556	0.050
P4	v	rank-1	0.034	0.018
P7	v	rank-1	0.011	0.012
P7	v	most	0.014	0.012
P7	v	all	0.190	0.013

Table 7.2: Comparison of the global error for different input points, velocity components, and range of frequencies.

(2017). The instantaneous error $I_v(t)$ is defined as

$$I_v(t) = \sqrt{\frac{1}{N_p} \sum_{i=1}^{N_p} (v^i(t) - v_{ref}^i(t))^2}, \quad (7.4)$$

where N_p is the number of discrete points where the velocity is known, v^i is the i th discrete transverse velocity value of the reconstructed field, and v_{ref}^i is the i th discrete transverse velocity of the reference field. The global error E is obtained by integrating $I(t)$ in time:

$$E_v = (1/T_{max}) \int_0^{T_{max}} I_v(t) dt. \quad (7.5)$$

If u had been measured instead of v , then it is possible to define $I_u(t)$ and E_u by making the necessary substitutions in Equations 7.4 and 7.5. In this study, $T_{max} = 3499$, or the number of snapshots obtained for a single PIV run as mentioned in Chapter 2.

The global error is computed for various cases, some of which have been presented thus far, in Table 7.2. The domain where E_u and E_v are computed is $x \in [1, 3]$ and $y \in [-0.3, 0.3]$. Choosing a larger domain in the y -direction would artificially increase the quality's reconstruction since there is no fluctuating activity outside the wake region. The length of the domain in the streamwise direction is based on where PIV measurements are available. A more sophisticated approach to minimize the influence of noise might involve limiting the PIV domain to regions where the velocity magnitude is greater than some threshold value. The model's capability with regard to predicting the downstream development of the flow outside the PIV domain is discussed in the next subsection.

For now, the discussion will be limited to analyzing trends associated with E_v as the u -component reconstruction is discussed in Section 7.7.3. The best performers from

Table 7.2 are those corresponding to measuring v at P7. Measuring v at P4 does 50% worse while measuring u does at least a factor of two worse. The inclusion of more frequencies has little impact on E_v . Using only rank-1 frequencies results in $E_v = 0.012$. If most frequencies are used ($2.0202 < \omega < 38.8322$), then E_v remains constant even though the shape improves while adding frequencies below $\omega = 2.0202$ (the motivation for this cutoff is discussed in Section 7.7.3) results in a slight uptick of E_v .

7.7.2 Influence of Mode Position and Amplitude

The reconstruction method is not able to match the PIV snapshots due to inaccuracies associated with the streamwise position and amplitude of the resolvent modes. As seen in Figure 7.7, the spatial support of the resolvent mode is situated further downstream than the DMD mode. The amplitude of the resolvent mode, moreover, is small when it first appears in the airfoil and gradually increases until it reaches its highest amplitude just before the edge of the domain. The DMD mode, on the other hand, indicates that the fluctuations for $\omega = 24.2$ begin further upstream immediately behind the airfoil and decrease in amplitude as a function of streamwise distance from the airfoil. From these observations, it becomes clear why a point further downstream where the fluctuations are weaker yields better results than a point further upstream. P7 compensates for the amplitude mismatch between the resolvent and DMD modes by calibrating the modes such that the upstream fluctuations are weak. The downstream fluctuations will match the probe point. Calibrating based on P4 would result in very large fluctuations downstream.

It is for this reason that the model does not predict with good fidelity the activity further downstream in the wake where PIV data are not available. The data-assimilated mean is computed on a much larger domain where $x = 15$, so it is possible to compute resolvent modes which extend this far downstream as well. Since they do not correctly predict the attenuation of fluctuations as one moves downstream, the model in its present form overestimates the intensity of fluctuations which occur further downstream.

7.7.3 Influence of Velocity Component

Finally, the reconstruction of the u -component is discussed using various ranges of frequencies. In Figure 7.16, the results from the P7 v -component calibration are illustrated for $t = 0$. Similar to its v counterpart, the use of rank-1 modes only in Figure 7.16 fares quite well as the agreement between the PIV and the

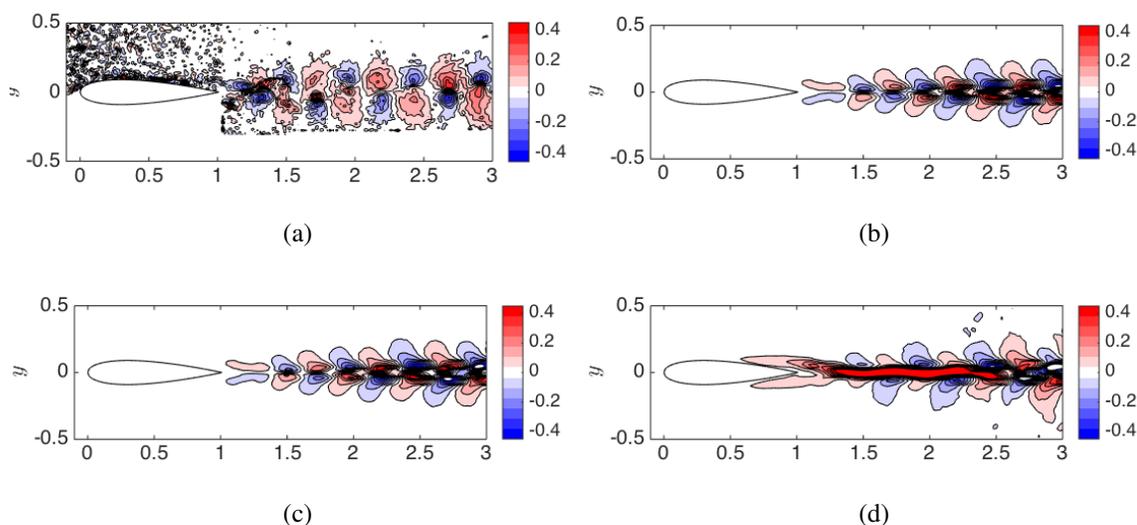


Figure 7.16: u -component of the reconstructed flow (b-d) compared to the PIV snapshot (a) using various amounts of frequencies. Only the rank-1 frequencies are used in (b) while (d) uses all frequencies. (c) includes all $\omega > 2.0202$.

reconstruction is quite good. There are discrepancies in the region immediately behind the airfoil, however, as the model fails to reproduce the first set of structures directly behind the trailing edge. The inclusion of more frequencies in panel (c) does not visibly alter the quality of the reconstruction. According to Table 7.2, there is a slight increase in E_u when including more modes, particularly at the lower frequencies. The motivation for cutting off very low frequencies becomes clear in panel (d) where the quality of the reconstructed flow has deteriorated significantly. There is a small layer of high speed velocity along the centerline which does not appear in the flow. This problem is not ameliorated when the resolvent modes are calibrated based on a u measurement as seen in Table 7.2. In fact, both E_u and E_v increase when measuring u and including either rank-1 frequencies or all frequencies in the model.

7.8 Mean Pressure Reconstruction

One can now return to the left-hand side of the schematic in Figure 7.1 to obtain an estimate for the mean pressure. The method relies on using the correctly weighted resolvent response modes to obtain $\nabla \cdot \mathbf{f}$. Due to the difficulties associated with reconstructing the flow immediately behind the airfoil, the amplitude of $\nabla \cdot \mathbf{f}$ is underestimated. Since the PIV data are time-resolved, the Reynolds stresses can be computed and these are used to obtain an estimate of the mean pressure instead

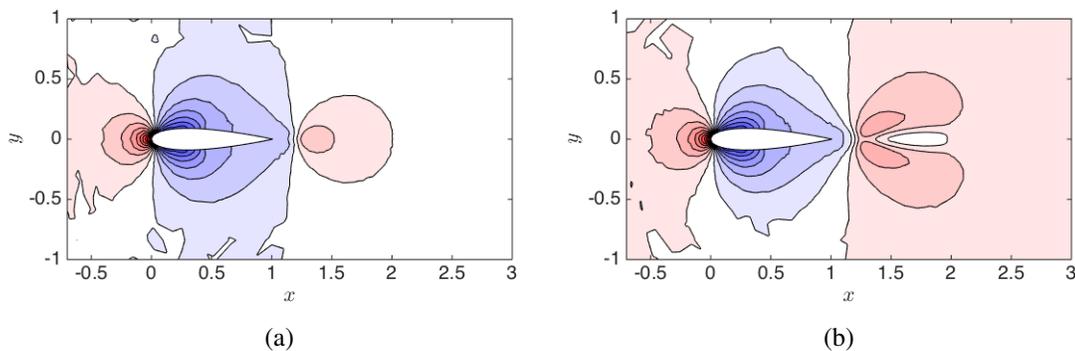


Figure 7.17: The mean pressure $\bar{p} + \xi$ of the data-assimilated mean (a) compared to the corrected mean pressure \bar{p} using Reynolds stress information (b).

of the approximation from the resolvent reconstruction along with the results from the data-assimilation which computes $\bar{p} + \xi$. The objective here is to compute ξ by solving

$$\nabla^2 \xi = -\nabla \cdot \mathbf{f}, \quad (7.6)$$

where $\nabla \cdot \mathbf{f}$ is estimated from the Reynolds stresses. This differs from the reconstruction procedure in Figure 7.1, which assumes knowledge of the mean velocity profile only. If the Reynolds stresses are assumed to be unknown, then Equation 2.63 could be utilized to correct the mean pressure obtained from data-assimilation. Since the quality of the PIV data is sufficiently good to compute gradients of the Reynolds stresses, the mean pressure correction can be obtained for this flow.

The data-assimilated mean pressure, or $\bar{p} + \xi$, is plotted in Figure 7.17(a) alongside the mean pressure obtained from solving Equation 7.6 in panel (b). The prediction from the data-assimilation is already quite accurate since the fluctuations are strong in the wake region only. The only differences, therefore, are concentrated immediately behind the trailing edge where the gradients of the Reynolds stresses are quite strong.

7.9 Convergence of the Method

As mentioned before, the singular values of the resolvent operator are a good indication of whether or not the predicted resolvent modes match DMD modes which are computed directly from the data. The convergence of the method can still be assessed without validating the resolvent modes against the DMD modes. One avenue, for example, is to compute the Reynolds stresses via the triadic interactions of resolvent modes and their complex conjugate counterparts. The result can be

compared with the Reynolds stress gradients obtained from the data-assimilation although it is necessary to isolate the solenoidal component before doing so. The benefit of this method is its simplicity although it has not been applied in this thesis since the PIV snapshots are available and the Reynolds stresses are underestimated.

A second, more rigorous validation approach would be to solve the fully nonlinear system using an optimization procedure analogous to the one developed by Rosenberg (2018). The procedure is formulated on determining the correct complex weights of all resolvent response modes and verifying that the nonlinear forcing generated by their triadic interactions both supports the mean profile and drives the velocity fluctuations at all frequencies. Due to the high degrees of freedom in this system, this may not be practical compared to the first method. The weights approximated using a single probe point, however, could dramatically speed up the optimization process as it helps eliminate temporal frequencies where the fluctuations are weak. The mode shapes computed by approximating the nonlinear forcing, furthermore, are closer to the truth than the most highly amplified resolvent modes. They suggest at which frequencies suboptimal modes should be computed as well as which suboptimal modes are important based on how the estimated mode shape, from the approximated nonlinear forcing, projects onto them.

7.10 Summary

Preliminary results have been presented for reconstructing the flow around a NACA 0018 airfoil at $\alpha = 0^\circ$ and $Re = 10250$. The methodology outlined in Figure 7.1 is followed quite closely with the exception of the mean pressure computation which was performed using knowledge of the Reynolds stresses. Ideally, the reconstructed fluctuations would be able to predict the Reynolds stresses but the amplitude disagreement behind the airfoil when using point P7 meant that these predictions were less accurate than using the Reynolds stresses computed from the time-resolved PIV directly. It would require multiple points in the wake to obtain the correct amplitudes for longer streamwise distances.

The flow reconstruction procedure relies on data-assimilating the PIV mean profile to obtain $\bar{\mathbf{u}}$ and $\bar{p} + \xi$. As mentioned in the previous paragraph, the Reynolds stresses were used to solve a Poisson equation for ξ , yielding a good approximation of the mean pressure field even though no pressure measurements were made. A resolvent analysis of the mean profile revealed low-rank behavior for a bandwidth of frequencies where the first resolvent response mode matched the DMD mode. In

the general case, access to DMD modes may not be possible so the rank-1 approximation should be used when there is one order of magnitude separating σ_1 and σ_2 . At frequencies where this is not the case, the nonlinear forcing is approximated from the interaction of two modes in the rank-1 region. The mode shape is then obtained by forcing the resolvent operator with this nonlinear forcing. At very high frequencies, it may be necessary to approximate the nonlinear forcing from a rank-1 mode and a 'forced' mode.

The unsteady velocity and pressure (not presented here) fluctuations are obtained by calibrating the complex amplitude of the resolvent modes from the discrete Fourier transform of a single point in the flow. The use of a single mode and a single point in the wake is a significant simplification yet it is appropriate for this flow. With just rank-1 modes, the reconstruction agrees reasonably well with the PIV snapshots at various instances in time and is able to identify extreme events in the flow which result from the beating of two distinct frequencies. The addition of modes outside the rank-1 region results in qualitative improvement of v as the reconstruction is able to capture some of the finer details of the structure such as the tilting of the contours. The quality of u , however, declines as a high amplitude region near the centerline appears. Even if the probe point measures u instead of v , the agreement does not improve- in fact, it gets considerably worse. Two possible sources of this problem are suggested. First, the mode shape at very low frequencies deviates from the true structure. This is particularly a problem for u which has a more complicated structure than v . The mode shapes at low frequencies tend to be symmetric and the amplitude of the component along the centerline increases downstream when it ought to decrease. Thus, when it is included in the reconstruction, this artifact manifests itself as a high amplitude waviness along the centerline.

Chapter 8

CONCLUSIONS AND FUTURE WORK

The objective of this thesis was to gain a better understanding of how resolvent analysis selects amplification mechanisms based on the mean profile. It could then be exploited to educe coherent structures in experimental flows and approximate the fluctuating velocity fields at any instant in time using very limited measurements. The final result is an estimate of the flow which obeys the NSE and reproduces the main features observed in the experiment. The complete reconstruction methodology in Figure 7.1 can thus be interpreted as a tool to perfect incomplete and/or noisy data from experiments.

The mean cylinder wake, computed from simulation data at low Reynolds numbers, was an ideal flow for developing the method due to the sparsity of temporal frequencies active in the flow. There were challenges when dealing with experimental data, namely a lack of resolution near walls or, in the case of flows around airfoils, missing data near the airfoil surface. The data-assimilation algorithm first developed by Foures et al. (2014) was extended to experimental data, yielding improved predictions of the mean velocity near a body. The flow around a NACA 0018 airfoil was then reconstructed using a data-assimilated mean velocity profile and resolvent analysis. A more detailed summary and the conclusions are presented in the next section. The thesis draws to a close with a discussion of future work and implications for estimating turbulent flows.

8.1 Summary and Conclusions

8.1.1 Fundamentals of Resolvent Analysis and Application to Flows Around Bodies

In Chapter 3 stability and resolvent analyses were juxtaposed directly to understand the linear amplification mechanisms identified. These mechanisms were categorized as resonant if they could be attributed to a single eigenvalue of the mean eigenspectrum or pseudoresonant if the response is composed of several nonorthogonal eigenfunctions. A scalar measure of the non-normality of a mechanism, initially proposed by Chomaz (2005) as a local measure, was computed globally from $|\hat{\phi}_1^* \hat{\psi}|^{-1}$. When it agreed with $\sigma_1 |i\omega - \lambda|$, then mean stability analysis was capable of identifying the linear mechanism; otherwise, it was pseudoresonant and resol-

vent analysis was required to understand the origin of the structure. In the example flows studied in Chapter 3, cylinder flow exhibited a strong resonant mechanism, hence the equivalence of stability and resolvent modes, which exhibits convective-type non-normality (Marquet et al., 2009) since the resolvent forcing modes are upstream of the resolvent response modes. It was also possible to quantify non-normality in two ways since $|\hat{\phi}_1^* \hat{\psi}|^{-1} = \sigma_1 |i\omega - \lambda|$.

Turbulent channel flow was also studied using wavenumber combinations representative of known turbulent structures. The strength of mean shear often resulted in a response dominated by the streamwise velocity component and highly non-normal mechanisms as quantified by $|\hat{\phi}_1^* \hat{\psi}|^{-1}$. Perhaps one of the most important contributions from this chapter was observing when the resolvent operator was likely to be low-rank. For resonant mechanisms, the resolvent operator was low-rank at frequencies where the eigenvalue was sufficiently separated from the rest of the spectrum as seen in cylinder flow. Pseudo-resonant mechanisms, on the other hand, tend to be low-rank since the first resolvent mode projects onto many nonorthogonal eigenvectors. Suboptimal modes are required to be orthogonal to the first mode. As such, they are much less likely to project onto eigenvectors prone to amplification.

The mean cylinder wake was explored in much greater depth in Chapter 4. A relationship between the scaling of the dominant resolvent modes and the mean profile was established through the shedding frequency. It was also shown that the length of the recirculation bubble, which scales with the size of the wavemaker, can be related to the shedding frequency. It can be concluded that highly amplified linear mechanisms leave a significant footprint on the mean profile. Resolvent analysis is therefore able to identify the coherent structures and their frequency by inspecting the geometry of the mean profile. Only one linear mechanism, the vortex shedding, is active in the low Reynolds number cylinder wake and manifests itself as a marginally stable eigenvalue in the mean LNS operator. Its influence spans over a range of temporal frequencies which coincide with where the resolvent operator is low-rank. Outside of this range, there is no separation of singular values and the first singular modes do not correspond to structures which appear in the flow. It was surmised that the self-interaction of the resolvent mode at the shedding frequency was the necessary nonlinear forcing of $\mathcal{H}(2\omega_s)$. After performing the computation, the ‘forced’ resolvent mode was a good match with the DMD mode at $2\omega_s$, reinforcing the notion that the structure of the nonlinear forcing must be taken into account to back out the correct mode shape. A crude approximation of the nonlin-

ear forcing for higher frequency modes, which can be considered parasitic in that they are driven by the vortex shedding mode, yielded successful predictions of the mode shapes when compared to the DMD modes.

8.1.2 Data-Assimilation of Time-Averaged Simulation and Experimental Data

Resolvent analysis was used to inform the measurement domain for data-assimilation of the mean cylinder wake at low Reynolds numbers in Chapter 5. The most vital part of the flow was the edge of the wavemaker region where the nonlinear forcing $\hat{\mathbf{f}}(\omega_s)$ overlaps with the first resolvent forcing mode $\hat{\phi}_1(\omega_s)$. It was also determined that the input measurements must coincide with the spatial support of the Reynolds stress gradients in the wake ($x \in [1, 4]$) to successfully capture $\nabla \cdot \mathbf{R}$. The mean pressure from data-assimilation of velocity-only measurements can be corrected by solving a Poisson equation which is forced by the divergence of the correctly weighted resolvent modes interacting nonlinearly with their complex conjugate counterparts. The advantage of this method is that it avoids having to compute gradients of Reynolds stresses which amplify underlying noise in the experiment. The weighting of measurement points to emphasize important flow physics was also considered. Although it was concluded that doing so leads to a slightly faster rate of convergence, there was non meaningful improvement in the final result.

The data-assimilation algorithm was adapted for experimental data in Chapter 6 using the smoothing procedure described in Chapter 2. The flow around an idealized airfoil at $Re_{chord} = 13500$ experiences flow separation near the leading edge so the PIV field of view included the area above the airfoil where the Reynolds stresses were likely to be strong as well as the wake. It was observed for the idealized airfoil that the PIV resolution necessary to obtain a good reconstruction was 30 vectors per chord length and the domain could be truncated to a region sufficiently small that only one camera would be necessary for data acquisition. The largest residual discrepancies were near the leading edge of the airfoil where the velocity gradients are highest and the PIV has difficulties pinpointing the separation point.

The mean flows around a NACA 0018 airfoil at $\alpha = 0^\circ$ and $\alpha = 10^\circ$ were also data-assimilated. The $\alpha = 0^\circ$ case was more difficult to converge, particularly at higher Reynolds numbers where the boundary layer and wake region are significantly thinner in the y -direction. As such, the reconstructed Reynolds stress gradients were less smooth than their experimental counterparts for $Re = 20700$ although the assimilation procedure worked well for $Re = 10250$. The reconstruction results for

$\alpha = 10^\circ$ where the flow is separated were a significant improvement over the experimental data since the algorithm was able to fill in missing data near the separation point and obtain smoother Reynolds stress gradients.

8.1.3 Flow Reconstruction from Very Limited Measurements

The data-assimilated profile for $\alpha = 0^\circ$ and $Re = 10250$ was used as an input to the flow reconstruction procedure in Chapter 7. The $\alpha = 10^\circ$ case is more complex and is discussed in Appendix B. The first three singular values were computed across ω to determine where $\sigma_1 \gg \sigma_2$. The sole linear mechanism active in the flow is a vortex shedding mode which was most amplified at $\omega = 12.24$. The frequency range where there was separation in the singular values ranged from $3.14 < \omega < 20.2$. The DFT of a single point in the wake where the fluctuations are energetic was used to calibrate the amplitude and phase of the resolvent modes. Using only rank-1 modes resulted in a satisfactory reconstruction of the u and v fluctuations when compared to the PIV snapshots. They also contained significantly less noise. Modes outside the rank-1 frequency range were calculated by approximating the nonlinear forcing from the triadically consistent interactions of rank-1 modes and running this through the resolvent operator at the frequency of interest. The addition of ‘forced’ resolvent modes to the reconstruction procedure led to discernible qualitative improvement of the v -fluctuations at the expense of the quality of the u reconstruction. While these preliminary results are promising, they suggest several directions for future work which are discussed in the next section.

8.2 Future Work

Some of the ways for improving the performance of the flow reconstruction were mentioned briefly in Chapter 7 and are mentioned in greater detail here. The reason why a probe point further downstream performed better than points further upstream was to compensate for the growing amplitude of the resolvent modes with downstream distance. The DMD modes, on the other hand, tend to decay more rapidly downstream. Such a correction may be provided by the suboptimal modes, if the expansion coefficients are uncorrelated (Towne et al., 2018), which were not considered in this thesis. On the other hand, it may be necessary to reconsider the inner product which is currently defined as the kinetic energy in the entire domain in order to obtain a more accurate mode shape. Other extensions to the reconstruction procedure, which were not considered here, include incorporating multiple probe points, using suboptimal modes in the reconstruction, and determining the optimal

placement of flow sensors. The work of Beneddine et al. (2017) looked into the first option yet obtained fairly minor improvement with the addition of probe points.

The flow reconstruction procedure was not performed for the $\alpha = 10^\circ$ case. The results of the resolvent analysis are summarized in Appendix B where it is observed that there are two peaks in the resolvent norm. The lower frequency peak corresponds to a wake mode that resembles vortex shedding while the higher frequency peak is associated with a shear layer instability near the leading edge. As proof that this is not an artifact from the data-assimilation procedure, the power spectrum is compared from a point in the shear layer to a point in the wake. The peak of the former point is close to the higher frequency peak in the resolvent norm while the peak of the latter point is closer to the lower frequency peak. This phenomenon could be similar to the Kelvin-Helmholtz mechanism and vortex pairing observed by Beneddine et al. (2017) for a round jet. Since the resolvent norm was not plotted therein, it is not known if it displayed two distinct peaks. It is also not known whether the spatial support of these modes were in close proximity and hence could be calibrated with a single probe measurement. Future work would be to investigate the minimum number of measurements necessary to recover an acceptable reconstruction of the flow for the $\alpha = 10^\circ$ case.

In its current form, the flow reconstruction method considered two-dimensional flows only. This is similar to Beneddine et al. (2016) and Beneddine et al. (2017), where the spanwise and azimuthal wavenumbers were set to zero, respectively. In Gómez et al. (2016a), the flow was Fourier-transformed in the spanwise direction giving insight into which spatial wavenumbers were important. The reconstructed flow, therefore, contained resolvent modes computed for energetic frequencies and spanwise wavenumbers. Since the flow around the airfoils is homogeneous in the spanwise direction, the methodology could be extended to model the three-dimensionality in the wake as the resolvent operator can easily accommodate a spanwise wavenumber. A single point measurement would need to collect data at various spanwise stations, however, to identify the spanwise wavenumber content in the flow. At higher Reynolds numbers, however, the wake may be turbulent, at which point there may be spatially uncorrelated fluctuations the resolvent fails to identify (see Beneddine et al., 2017). The influence of the stall cells, furthermore, may disrupt the homogeneity in the spanwise direction in a manner which is difficult to model with the current procedure.

One of the future considerations in terms of the data-assimilation framework is to

account for the out-of-plane velocities which are not captured in the measurements or the processed data. While the 2D incompressible RANS equations seem justified for the flows in this thesis, an algorithm in the future could account for out-of-plane velocity gradients by introducing an additional compliance parameter $d = -\partial_z w = \partial_x u + \partial_y v$ and replacing the divergence-related term in Equation 2.53 by $\langle p^\dagger, \nabla \cdot \bar{\mathbf{u}} - d \rangle$. For a 2D flow ($d = 0$) this expression reverts back to the previous augmented Lagrangian given by Equation 2.53. A second would be to compare the predictions of a RANS simulation with a turbulence model to the experiment and assimilated results. The degree of success is dependent on the model chosen so there is not a unique comparison between traditional RANS and the approach outlined in this study. It would be worthwhile to apply data-assimilation to three-dimensional mean profiles obtained from tomographic PIV (Elsinga et al., 2006), but the ambiguity of the correct boundary conditions would make it difficult to implement.

Finally, it would be interesting to consider the potential of resolvent-based modal expansions for purposes of control, particularly in relation to the placement of sensors and actuators. The resolvent response modes are outputs which are highly amplified by the linear dynamics of the NSE, and so sensors could be placed where these are likely to be strong. The resolvent forcing modes are the ‘trigger’ or input which leads to high amplification, and so the actuators could be placed to manipulate the flow in such a way that suppresses these disturbances. In cases where there is large spatial separation between resolvent forcing and response modes, it is possible that improved performance could be attained by sensing and actuating within a wavemaker region (e.g. Giannetti and Luchini, 2007; Chen and Rowley, 2011), which, as discussed in Chapter 3, may also be obtained from resolvent analysis. While the resolvent decomposition shows potential for control applications (Luhar et al., 2014), further refinements could seek to balance the observability and controllability of the reduced-order model (Moore, 1981; Rowley, 2005; Barbagallo et al., 2009), subject to known sensor and actuator locations, and information about the nature of the nonlinear forcing (Gómez and Blackburn, 2017).

8.3 Implications for Estimating Turbulent Flows

While it is outside the scope of this thesis, one of the eventual goals is to design effective estimators of turbulent flows for closed-loop control. Hoepffner et al. (2005) and Jones et al. (2011) have applied linear estimation to wall-bounded flows at laminar Reynolds numbers while Chevalier et al. (2006) have done so for very low tur-

bulent Reynolds numbers. The first work to apply similar tools to fully developed turbulence at a relatively high Reynolds number ($Re_\tau = 1000$) was Illingworth et al. (2018), who used time-resolved velocity measurements at a single wall-normal location in turbulent channel flow to estimate the velocity field at other wall-normal locations. They demonstrated that it was possible to estimate the largest scales with reasonable accuracy using linear models with eddy viscosity. The ranges of spatial wavenumbers and wall heights over which the estimator will perform well could also be predicted by the linear models. The eddy viscosity model accounted for the influence of the small-scale turbulent fluctuations on the large scales. Without it, the linear estimator overpredicted the amplitude of velocity fluctuations, particularly for high aspect ratio structures (low k_x , high k_z).

The work in this thesis suggests that the resolvent operator in wall-bounded shear flows is predisposed to selecting mechanisms which result in a strong streamwise response. The high mean shear in the LNS operator couples many nonorthogonal eigenfunctions leading to non-normal mechanisms exhibiting component-type non-normality. The singular value decomposition of the resolvent operator attempts to maximize non-normality of the linear mechanism since the resolvent norm is inversely proportional to the inner product of the forcing and response modes. This is particularly evident for streamwise constant disturbances (i.e. $k_x = 0$) in which the OS and SQ operators are normal but the full OS-SQ system is highly non-normal due to the coupling operator $ik_z\bar{u}'$. Progress has been made by Rosenberg and McKeon (2018) to improve the use of resolvent modes as a basis by isolating the contributions from the OS and SQ modes. While it remains an open question for whether an analogous procedure in the two-dimensional case is necessary, this thesis has demonstrated two main points for the spatially-developing case. First, a singular value decomposition of the resolvent operator identifies strong linear mechanisms which leave a footprint on the mean profile. Second, the resolvent operator is able to predict structure even when a linear mechanism is not important by forcing it with the interaction of coherent structures. It is hoped that the closed-loop nature of the resolvent framework (McKeon et al., 2013) can be exploited further to identify the coherent structures in the flow.

BIBLIOGRAPHY

- Abdessemed, N., Sharma, A. S., Sherwin, S. J., and Theofilis, V. (2009). “Transient growth analysis of the flow past a circular cylinder”. In: *Phys. Fluids* 21.044103.
- Adrian, R J (1991). “Particle-imaging techniques for experimental fluid mechanics”. In: *Annu. Rev. Fluid Mech.* 23, pp. 261–304.
- Amestoy, P. R., Duff, I. S., L’Excellent, J.-Y., and Koster, J. (2001). “A fully asynchronous multifrontal solver using distributed dynamic shedding”. In: *SIAM J. Matrix Anal. Appl.* 23.1, pp. 15–41.
- Balachandar, S., Mittal, R., and Najjar, F. M. (1997). “Properties of the mean recirculation region in the wakes of two-dimensional bluff bodies”. In: *J. Fluid Mech.* 351, pp. 167–199.
- Barbagallo, A., Sipp, D., and Schmid, P. J. (2009). “Closed-loop control of an open cavity flow using reduced-order models”. In: *J. Fluid Mech.* 641, pp. 1–50.
- Barkley, D. (2006). “Linear analysis of the cylinder wake mean flow”. In: *Europhys. Lett.* 75, pp. 750–756.
- Barkley, D. and Henderson, R. D. (1996). “Three-dimensional Floquet stability analysis of the wake of a circular cylinder”. In: *J. Fluid Mech.* 322, pp. 215–241.
- Beneddine, S., Sipp, D., Arnault, A., Dandois, J., and Lesshafft, L. (2016). “Conditions for validity of mean flow stability analysis”. In: *J. Fluid Mech.* 798, pp. 485–504.
- Beneddine, S., Yegavian, R., Sipp, D., and Leclaire, B. (2017). “Unsteady flow dynamics reconstruction from mean flow and point sensors: an experimental study”. In: *J. Fluid Mech.* 824, pp. 174–201.
- Bewley, T R, Moin, P, and Temam, R (2001). “DNS-based predictive control of turbulence: an optimal benchmark for feedback algorithms”. In: *J. Fluid Mech.* 447, pp. 179–225.
- Blackburn, H. M. and Sherwin, S. J. (2004). “Formulation of a Galerkin spectral element-Fourier method for three-dimensional incompressible flows in cylindrical geometries”. In: *J. Comp. Phys.* 197.2, pp. 759–778.
- Bourguignon, J.-L. (2012). “Models of turbulent pipe flow”. PhD thesis. California Institute of Technology.
- Brandt, L., Sipp, D., Pralits, J. O., and Marquet, O. (2011). “Effect of base-flow variation in noise amplifiers: the flat-plate boundary layer”. In: *J. Fluid Mech.* 687, pp. 503–528.
- Butler, K. M. and Farrell, B. F. (1992). “Three-dimensional optimal perturbations in viscous shear flow”. In: *Phys. Fluids A* 4, pp. 1637–1650.

- Chandler, G. J., Juniper, M. P., Nichols, J. W., and Schmid, P. J. (2012). “Adjoint algorithms for the Navier-Stokes equations in the low Mach number limit”. In: *J. Comp. Phys.* 231.4, pp. 1900–1916.
- Chen, K. K. and Rowley, C. W. (2011). “ H_2 optimal actuator and sensor placement in the linearised Ginzburg-Landau system”. In: *J. Fluid Mech.* 681, pp. 241–260.
- Chevalier, M., Hoepffner, J., Bewley, T. R., and Henningson, D. S. (2006). “State estimation in wall-bounded flow systems. Part 2. Turbulent flows”. In: *J. Fluid Mech.* 552, pp. 167–187.
- Chomaz, J. M. (2005). “Global instabilities in spatially developing flows: non-normality and nonlinearity”. In: *Annu. Rev. Fluid Mech.* 37, pp. 357–392.
- Chomaz, J. - M., Huerre, P., and Redekopp, L. G. (1988). “Bifurcations to local and global modes in spatially developing flows”. In: *Phys. Rev. Lett.* 60.25.
- Cimbala, J. M., Nagib, H. M., and Roshko, A. (1988). “Large structure in the far wakes of two-dimensional bluff bodies”. In: *J. Fluid Mech.* 190, pp. 265–298.
- Colburn, C. H., Cessna, J. B., and Bewley, T. R. (2011). “State estimation in wall-bounded flow systems. Part 3. The ensemble Kalman filter”. In: *J. Fluid Mech.* 682, pp. 289–303.
- Dergham, G., Sipp, D., and Robinet, J.-Ch. (2013). “Stochastic dynamics and model reduction of amplifier flows: the backward facing step flow”. In: *J. Fluid Mech.* 719, pp. 406–430.
- Dunne, R. (2016). “Dynamic stall on vertical axis wind turbine blades”. PhD thesis. California Institute of Technology.
- Elsinga, G. E., Scarano, F., Wieneke, B., and Oudheusden, B. W. van (2006). “Tomographic particle image velocimetry”. In: *Exp. in Fluids* 41, pp. 933–947.
- Evensen, G. (1994). “Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics”. In: *J. Geophys. Res.* 99, pp. 10143–10162.
- Farrell, B. F. (1987). “Developing disturbances in shear”. In: *J. Atmos. Sci.* 44.16, pp. 2191–2199.
- Farrell, B. F. and Ioannou, P. J. (1993). “Stochastic forcing of the linearized Navier-Stokes equations”. In: *Phys. Fluids* 5, pp. 2600–2609.
- Foures, D. P. G., Dovetta, N., Sipp, D., and Schmid, P. J. (2014). “A data-assimilation method for Reynolds-averaged Navier-Stokes-driven mean flow reconstruction”. In: *J. Fluid Mech.* 759, pp. 404–431.
- Gebhardt, T. and Grossmann, S. (1994). “Chaos transition despite linear stability”. In: *Phys. Rev. E* 50.5, p. 3705.
- Gelfand, I. (1941). “Zur Theorie der Charaktere der Abelschen topologischen Gruppen”. In: *Rec. Math. N. S.* 9, pp. 49–50.

- Gerakopoulos, G, Boutilier, M. S. H., and Yarusevych, S. (2010). “Aerodynamic characterization of a NACA 0018 airfoil at low Reynolds numbers”. In: *40th AIAA Fluid Dynamics Conference, Chicago, Illinois 2010-4629*.
- Gharib, M (1983). “The effect of flow oscillations on cavity drag, and a technique for their control”. PhD thesis. California Institute of Technology.
- Giannetti, F. and Luchini, P. (2007). “Structural sensitivity of the first instability of the cylinder wake”. In: *J. Fluid Mech.* 581, pp. 167–197.
- Gómez, F. and Blackburn, H. M. (2017). “Data-driven approach to design of passive flow control strategies”. In: *Phys. Rev. Fluids* 2, p. 021901.
- Gómez, F., Blackburn, H. M., Rudman, M., McKeon, B. J., Luhar, M., Moarref, R., and Sharma, A. S. (2014). “On the origin of frequency sparsity in direct numerical simulations of turbulent pipe flow”. In: *Phys. Fluids* 26.101703.
- Gómez, F., Blackburn, H. M., Rudman, M., Sharma, A. S., and McKeon, B. J. (2016a). “A reduced-order model of three-dimensional unsteady flow in a cavity based on the resolvent operator”. In: *J. Fluid Mech.* 798.R2.
- Gómez, F., Sharma, A. S., and Blackburn, H. M. (2016b). “Estimation of unsteady aerodynamic forces using pointwise velocity data”. In: *J. Fluid Mech.* 804, R4.
- Gonzalez, C, Bruhat, J-F, Wainfan, B, and McKeon, B J (2010). “Control of laminar separation on an idealized airfoil using periodic dynamic roughness actuation”. In: *Gallery of Fluid Motion, Poster 26*.
- Gronskis, A, Heitz, D, and Mémin, E (2013). “Inflow and initial conditions for direct numerical simulation based on adjoint data assimilation”. In: *J. Comput. Phys.* 242, pp. 480–497.
- Gudmundsson, K. and Colonius, T. (2011). “Instability wave models for the near-field fluctuations of turbulent jets”. In: *J. Fluid Mech.* 689, pp. 97–128.
- Hack, M. J. P. and Moin, P. (2017). “Algebraic disturbance growth by interaction of Orr and lift-up mechanisms”. In: *J. Fluid Mech.* 829, pp. 112–126.
- Hayase, T. (2015). “Numerical simulation of real-world flows”. In: *Fluid Dyn. Res.* 47.051201.
- Hecht, F. (2012). “New development in FreeFem++”. In: *J. Numer. Math.* 20, pp. 251–265.
- Heitz, D., Mémin, E., and Schnorr, C. (2010). “Variational fluid flow measurements from image sequences: synopsis and perspectives”. In: *Exp. Fluids* 48.3, pp. 369–393.
- Hoepffner, J., Chevalier, M., Bewley, T. R., and Henningson, D. S. (2005). “State estimation in wall-bounded flow systems. Part 1. Perturbed laminar flows”. In: *J. Fluid Mech.* 534, pp. 263–294.

- Huerre, P. and Monkewitz, P. A. (1985). “Absolute and convective instabilities in free shear layers”. In: *J. Fluid Mech.* 159, pp. 151–168.
- Hwang, Y. and Cossu, C. (2010). “Linear non-normal energy amplification of harmonic and stochastic forcing in the turbulent channel flow”. In: *J. Fluid Mech.* 664, pp. 51–73.
- Illingworth, S. J., Monty, J. P., and Marusic, I. (2018). “Estimating large-scale structures in wall turbulence using linear models”. In: *J. Fluid Mech.* 842, pp. 146–162.
- Jeun, J., Nichols, J. W., and Jovanović, M. R. (2016). “Input-output analysis of high-speed axisymmetric isothermal jet noise”. In: *Phys. Fluids* 28.047101.
- Jones, B. L., Kerrigan, E. C., Morrison, J. F., and Zaki, T. A. (2011). “Flow estimation of boundary layers using DNS-based wall shear information”. In: *Int. J. Control* 84.8, pp. 1310–1325.
- Jovanović, M. R. and Bamieh, B. (2005). “Componentwise energy amplification in channel flows”. In: *J. Fluid Mech.* 534, pp. 145–183.
- Juniper, M. (2012). “Absolute and convective instability in gas turbine fuel injectors”. In: *Proceedings of ASME Turbo Expo 2012: Power for Land, Sea and Air, Copenhagen*. ASME, New York.
- Kalman, R. E. (1960). “A new approach to linear filtering and prediction problems”. In: *Trans. ASME J. Basic Engng.* 82, pp. 35–45.
- Kang, W. and Xu, L. (2012). “Optimal placement of mobile sensors for data assimilations”. In: *Tellus A* 64.17133.
- Kato, H., Yoshiwaza, A., Ueno, G., and Obayashi, S. (2015). “A data assimilation methodology for reconstructing turbulent flows around aircraft”. In: *J. Comput. Phys.* 283, pp. 559–581.
- Kim, J and Bewley, T R (2007). “A linear systems approach to flow control”. In: *Annu. Rev. Fluid Mech.* 39, pp. 383–417.
- Kompenhans, J, Raffel, M, and Willert, C (2007). *Particle Image Velocimetry - A practical guide*. Springer, Berlin.
- Landahl, M. T. (1980). “A note on an algebraic instability of inviscid parallel shear flows”. In: *J. Fluid Mech.* 98, pp. 243–251.
- Le Dimet, F.-X. and Talagrand, O. (1986). “Variational algorithms for analysis and assimilation of meteorological observations: theoretical aspects”. In: *Tellus A* 38A, pp. 97–110.
- Lehoucq, R. B. and Sorensen, D. C. (1996). “Deflation techniques for an implicitly restarted Arnoldi iteration”. In: *SIAM J. Matrix Anal. Appl.* 17.4, pp. 789–821.
- Leontini, J. S., Thompson, M. C., and Hourigan, K. (2010). “A numerical study of global frequency selection in the time-mean wake of a circular cylinder”. In: *J. Fluid Mech.* 645, pp. 435–446.

- Lewis, J M, Lakshmivarahan, S, and Dhall, S (2006). *Dynamic data assimilation: A least squares approach*. Vol. 13. Encyclopedia of mathematics and its application 104. Cambridge University Press.
- Lin, C. and Hsieh, S.-C. (2003). “Convection velocity of vortex structures in the near wake of a circular cylinder”. In: *J. Eng. Mech.* 129.10, pp. 1108–1118.
- Lu, L. and Papadakis, G. (2014). “An iterative method for the computation of the response of the linearised Navier-Stokes equations to the harmonic forcing and application to forced cylinder wakes”. In: *Int. J. Numer. Methods Fluids* 74, pp. 794–817.
- Luchini, P. and Bottaro, A. (2014). “Adjoint equations in stability analysis”. In: *Annu. Rev. Fluid Mech.* 46, pp. 493–517.
- Luhar, M., Sharma, A. S., and McKeon, B. J. (2014). “Opposition control within the resolvent analysis framework”. In: *J. Fluid Mech.* 749, pp. 597–626.
- Lumley, J. L. (1970). *Stochastic tools in turbulence*. New York: Academic Press.
- Manohar, K., Brunton, B. W., Kutz, J. N., and Brunton, S. L. (2017). “Data-driven sparse sensor placement for reconstruction”. In: *arXiv:1701.07569*.
- Mantič-Lugo, V., Arratia, C., and Gallaire, F. (2014). “Self-consistent mean flow description of the nonlinear saturation of the vortex shedding in the cylinder wake”. In: *Phys. Rev. Lett.* 113.084501.
- Marquet, O., Lombardi, M., Chomaz, J. - M., Sipp, D., and Jacquin, L. (2009). “Direct and adjoint global modes of a recirculation bubble: lift-up and convective non-normalities”. In: *J. Fluid Mech.* 622, pp. 1–21.
- Marquet, O, Sipp, D, and Jacquin, L (2008). “Sensitivity analysis and passive control of cylinder flow”. In: *J. Fluid Mech.* 516, pp. 221–252.
- McKeon, B. J. (2017). “The engine behind (wall) turbulence: perspectives on scale interactions”. In: *J. Fluid Mech.* 817, P1.
- McKeon, B. J. and Sharma, A. S. (2010). “A critical-layer framework for turbulent pipe flow”. In: *J. Fluid Mech.* 658, pp. 336–382.
- McKeon, B. J., Sharma, A. S., and Jacobi, I. (2013). “Experimental manipulation of wall turbulence: A systems approach”. In: *Phys. Fluids* 25.031301.
- Meliga, P., Boujo, E., and Gallaire, F. (2016). “A self-consistent formulation for the sensitivity analysis of finite-amplitude vortex shedding in the cylinder wake”. In: *J. Fluid Mech.* 800, pp. 327–357.
- Meliga, P, Oujals, G, and Serre, E (2012). “Sensitivity of 2-D turbulent flow past a D-shaped cylinder using global stability”. In: *Phys. Fluids* 24.61701.
- Mettot, C, Sipp, D, and Bézard, H (2014). “Quasi-laminar stability and sensitivity analyses for turbulent flows: Prediction of low-frequency unsteadiness and passive control”. In: *Phys. Fluids* 26.45112.

- Moarref, R., Jovanović, M. R., Tropp, J. A., Sharma, A. S., and McKeon, B. J. (2014). “A low-order decomposition of turbulent channel flow via resolvent analysis and convex optimization”. In: *Phys. Fluids* 26.051701.
- Moarref, R., Sharma, A. S., Tropp, J. A., and McKeon, B. J. (2013). “Model-based scaling of the streamwise energy density in high-Reynolds-number turbulent channels”. In: *J. Fluid Mech.* 734, pp. 275–316.
- Monkewitz, P. A. (1988). “The absolute and convective nature of instability in two-dimensional wakes at low Reynolds numbers”. In: *Phys. Fluids* 31.999.
- Mons, V., Chassaing, J. C., Gomez, T., and Sagaut, P. (2016). “Reconstruction of unsteady viscous flows using data assimilation schemes”. In: *J. Comput. Phys.* 316, pp. 255–280.
- Mons, V., Chassaing, J.-C., and Sagaut, P. (2017). “Optimal sensor placement for variational data assimilation of unsteady flows past a rotationally oscillating cylinder”. In: *J. Fluid Mech.* 823, pp. 230–277.
- Moore, B. (1981). “Principal component analysis in linear systems: Controllability, observability, and model reduction”. In: *IEEE Trans. Autom. Control* 26, pp. 17–32.
- Nayar, M. and Ortega, U. (1993). “Computation of selected eigenvalues of generalized eigenvalue problems”. In: *J. Comput. Phys.* 108, pp. 8–14.
- Nisugi, K., Hayase, T., and Shirai, A. (2004). “Fundamental study of hybrid wind tunnel integrating numerical simulation and experiment in analysis of flow field”. In: *JSME Int. J.* 47, pp. 593–604.
- Noack, B. R. and Eckelmann, H. (1994). “A global stability analysis of the steady and periodic cylinder wake”. In: *J. Fluid Mech.* 270, pp. 297–330.
- Oberleithner, K., Rukes, L., and Soria, J. (2014). “Mean flow stability analysis of oscillating jet experiments”. In: *J. Fluid Mech.* 757, pp. 1–32.
- Orr, W. Mc. F. (1907). “The stability or instability of the steady motions of a perfect liquid and of a viscous liquid”. In: *Proc. R. Irish Acad.* 27, pp. 9–138.
- Oudheusden, B. W. van (2013). “PIV-based pressure measurement”. In: *Meas. Sci. Technol.* 24.032001.
- Papadakis, N. and Mémin, E. (2008). “Variational assimilation of fluid motion from image sequence”. In: *SIAM J. Imaging Sci.* 1.343-363.
- Pier, B. (2002). “On the frequency selection of finite-amplitude vortex shedding in the cylinder wake”. In: *J. Fluid Mech.* 458, pp. 407–417.
- Polak, E. and Ribiere, G. (1969). “Note sur la convergence de méthodes de directions conjuguées”. In: *Revue française d’informatique et de recherche opérationnelle, série rouge* 3.1, pp. 35–43.

- Prasad, A and Williamson, C H K (1997). “The instability of the shear layer separating from a bluff body”. In: *J. Fluid Mech.* 333, pp. 375–402.
- Provansal, M., Mathis, C., and Boyer, L. (1987). “Bénard-von Kármán instability: transient and forced regimes”. In: *J. Fluid Mech.* 182, pp. 1–22.
- Qadri, U. A. and Schmid, P. J. (2017). “Frequency selection mechanisms in the flow of a laminar boundary layer over a shallow cavity”. In: *Phys. Rev. Fluids* 2.013902.
- Raffel, M., Willert, C. E., Scarano, F., Kähler, C. J., Wereley, S. T., and Kompenhans, J. (2018). *Particle Image Velocimetry - A practical guide*. Springer.
- Reddy, S. C., Schmid, P. J., and Henningson, D. S. (1993). “Pseudospectra of the Orr-Sommerfeld operator”. In: *SIAM J. Appl. Math.* 53, pp. 15–47.
- Reynolds, W. C. and Tiederman, W. G. (1967). “Stability of turbulent channel flow, with application to Malkus’s theory”. In: *J. Fluid Mech.* 27.2, pp. 253–272.
- Rosenberg, K. (2018). “Resolvent-based modeling of flows in a channel”. PhD thesis. California Institute of Technology.
- Rosenberg, K. and McKeon, B. J. (2018). “Efficient representation of exact coherent states of the Navier-Stokes equations using resolvent analysis”. In: *Fluid Dyn. Res.*
- Rowley, C. W. (2005). “Model reduction for fluids using balanced proper orthogonal decomposition”. In: *Int. J. Bifurc. Chaos* 15, pp. 997–1013.
- Rowley, C. W., Colonius, T., and Basu, A. J. (2002). “On self-sustained oscillations in two-dimensional compressible flow over rectangular cavities”. In: *J. Fluid Mech.* 455, pp. 315–346.
- Rowley, C. W. and Dawson, S. T. M. (2017). “Model reduction for flow analysis and control”. In: *Ann. Rev. Fluid Mech.* 49, pp. 387–417.
- Rowley, C. W., Mezić, I., Bagheri, S., Schlatter, P., and Henningson, D. S. (2009). “Spectral analysis of nonlinear flows”. In: *J. Fluid Mech.* 641, pp. 115–127.
- Schmid, P. J. (2007). “Nonmodal stability theory”. In: *Annu. Rev. Fluid Mech.* 39, pp. 129–162.
- Schmid, P. J. (2010). “Dynamic mode decomposition of numerical and experimental data”. In: *J. Fluid Mech.* 656, pp. 5–28.
- Schmid, P. J. (2011). “Application of the dynamic mode decomposition to experimental data”. In: *Exp. Fluids* 50, pp. 1123–1130.
- Schmid, P. J. and Brandt, L. (2014). “Analysis of fluid systems: stability, receptivity, sensitivity”. In: *Appl. Mech. Rev.* 66.024803.
- Schmid, P. J. and Henningson, D. S. (2001). *Stability and transition in shear flows*. Springer.

- Schmidt, O. T., Towne, A., Colonius, T., Cavalieri, A. V. G., Jordan, P., and Brès, G. A. (2017a). “Wavepackets and trapped acoustic modes in a turbulent jet: coherent structure education and global stability”. In: *J. Fluid Mech.* 825, pp. 1153–1181.
- Schmidt, O. T., Towne, A., Rigas, G., Colonius, T., and Brès, G. A. (2017b). “Spectral analysis of jet turbulence”. In: *arXiv preprint arXiv:1711.06296*.
- Sharma, A. S., Mezić, I., and McKeon, B. J. (2016a). “Correspondence between Koopman mode decomposition, resolvent mode decomposition, and invariant solutions of the Navier-Stokes equations”. In: *Phys. Rev. Fluids* 1.032402(R).
- Sharma, A. S., Moarref, R., McKeon, B. J., Park, J. S., Graham, M. D., and Willis, A. P. (2016b). “Low-dimensional representations of exact coherent states of the Navier-Stokes equations from the resolvent model of wall turbulence”. In: *Phys. Rev. E* 93.2, p. 021102.
- Silva, A. F. C. da and Colonius, T. (2017). “An EnKF-based flow estimator for aerodynamic flows”. In: *8th AIAA Theoretical Fluid Mechanics Conference*, p. 3483.
- Sipp, D. and Lebedev, A. (2007). “Global stability of base and mean flows: a general approach and its applications to cylinder and open cavity flows”. In: *J. Fluid Mech.* 593, pp. 333–358.
- Sipp, D. and Marquet, O. (2013). “Characterization of noise amplifiers with global singular modes: the case of the leading-edge flat-plate boundary layer”. In: *Theor. Comput. Fluid Dyn.* 27, pp. 617–635.
- Sipp, D., Marquet, O., Meliga, P., and Barbagallo, A (2010). “Dynamics and control of global instabilities in open-flows: a linearized approach”. In: *Appl Mech Rev* 63.030801.
- Sreenivasan, K. R., Strykowski, P. J., and Olinger, D. J. (1987). “Hopf bifurcation, Landau equation, and vortex shedding behind circular cylinders”. In: *Proc. Forum on Unsteady Flow Separation*. Ed. by K. Ghia, pp. 1–13.
- Strykowski, P and Sreenivasan, K (1990). “On the formation and suppression of vortex ‘shedding’ at low Reynolds numbers”. In: *J. Fluid Mech.* 218, pp. 71–107.
- Suponitsky, V., Avital, E., and Gaster, M. (2005). “On three-dimensionality and control of incompressible cavity flow”. In: *Phys. Fluids* 17.104103.
- Suzuki, T (2012). “Reduced-order Kalman-filtered hybrid simulation combining particle tracking velocimetry and direct numerical simulation”. In: *J. Fluid Mech.* 709, pp. 249–288.
- Suzuki, T, Ji, H, and Yamamoto, F (2009a). “Unsteady PTV velocity field past an airfoil solved with DNS: Part 1. Algorithm of hybrid simulation and hybrid velocity field at $Re \approx 10^3$ ”. In: *Exp. Fluids* 47, pp. 957–976.

- Suzuki, T, Sanse, A, Mizushima, T, and Yamamoto, F (2009b). “Unsteady PTV velocity field past an airfoil solved with DNS: Part 2. Validation and application at Reynolds numbers up to $Re \leq 10^4$ ”. In: *Exp. Fluids* 47, pp. 977–994.
- Symon, S., Dovetta, N., McKeon, B. J., Sipp, D., and Schmid, P. J. (2017). “Data assimilation of mean velocity from 2D PIV measurements of flow over an idealized airfoil”. In: *Exp. in Fluids* 58.5, p. 61.
- Symon, S., Rosenberg, K., Dawson, S. T. M., and McKeon, B. J. (2018). “Non-normality and classification of amplification mechanisms in stability and resolvent analysis”. In: *Phys. Rev. Fluids* 3.053902.
- Taira, K., Brunton, S. L., Dawson, S. T. M., Rowley, C. W., Colonius, T., McKeon, B. J., Schmidt, O. T., Gordeyev, S., Theofilis, V., and Ukeiley, L. S. (2017). “Modal analysis of fluid flows: an overview”. In: *AIAA J.* 55, pp. 4013–4041.
- Thomareis, N. (2017). “Effect of trailing edge shape on the flow characteristics around an airfoil at low Reynolds number”. PhD thesis. Imperial College London.
- Thompson, M. C., Radi, A., Rao, A., Sheridan, J., and Hourigan, K. (2014). “Low-Reynolds-number wakes of elliptical cylinders: from the circular cylinder to the normal flat plate”. In: *J. Fluid Mech.* 751, pp. 570–600.
- Towne, A. (2016). “Advancements in jet turbulence and noise modeling: accurate one-way solutions and empirical evaluation of the nonlinear forcing of wavepackets”. PhD thesis. California Institute of Technology.
- Towne, A., Colonius, T., Jordan, P., Cavalieri, A. V. G., and Brès, G. A. (2015). “Stochastic and nonlinear forcing of wavepackets in a Mach 0.9 jet”. In: *21st AIAA/CEAS Aeroacoustics Conf.*
- Towne, A., Schmidt, O. T., and Colonius, T. (2018). “Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis”. In: *J. Fluid Mech.* 847, pp. 821–867.
- Trefethen, A. E., Trefethen, L. N., and Schmid, P. J. (1999). “Spectra and pseudospectra for pipe Poiseuille flow”. In: *Comput. Methods in Appl. Mech. Eng.* 175, pp. 413–420.
- Trefethen, L. N., Trefethen, A. E., Reddy, S. C., and Driscoll, T. A. (1993). “Hydrodynamic stability without eigenvalues”. In: *Science* 261, pp. 578–584.
- Tu, J. H., Rowley, C. W., Luchtenburg, D. M., Brunton, S. L., and Kutz, J. N. (2014). “On dynamic mode decomposition: theory and applications”. In: *J. Comput. Dyn.* 1.2, pp. 391–421.
- Turton, S. E., Tuckerman, L. S., and Barkley, D. (2015). “Prediction of frequencies in thermosolutal convection from mean flows”. In: *Phys. Rev. E* 91.043009.
- Wallace, R and McKeon, B. J. (2012). “Laminar separation bubble manipulation with dynamic roughness”. In: *6th AIAA Flow Control Conference, New Orleans, Louisiana* 2012-2680.

- Weideman, J. A. and Reddy, S. C. (2000). “A MATLAB differentiation matrix suite”. In: *ACM Trans. Math. Softw.* 26.4, pp. 465–519.
- Weihs, D. and Katz, J. (1983). “Cellular patterns in poststall flow over unswept wings”. In: *AIAA J.* 21.12, pp. 1757–1759.
- Westerweel, J and Scarano, F (2005). “Universal outlier detection for PIV data”. In: *Exp. Fluids* 39.6, pp. 1096–1100.
- Williamson, C. H. K. (1988). “The existence of two stages in the transition to three-dimensionality of a cylinder wake”. In: *Phys. Fluids* 31.2742.
- Williamson, C. H. K. (1992). “The natural and forced formation of spot-like ‘vortex dislocations’ in the transition of a wake”. In: *J. Fluid Mech.* 243, pp. 393–441.
- Williamson, C. H. K. (1996). “Vortex dynamics in the cylinder wake”. In: *Ann. Rev. Fluid Mech.* 28, pp. 477–539.
- Williamson, C. H. K. and Prasad, A. (1993). “A new mechanism for oblique wave resonance in the ‘natural’ far wake”. In: *J. Fluid Mech.* 256, pp. 269–313.
- Zare, A., Jovanović, M. R., and Georgiou, T. T. (2017). “Colour of turbulence”. In: *J. Fluid Mech.* 812, pp. 636–680.
- Zielinska, B J A, Goujon-Durand, S, Dusek, J, and Wesfried, J E (1997). “Strongly nonlinear effect in unstable wakes”. In: *Phys. Rev. Lett.* 79, pp. 3893–3896.

Appendix A

DERIVATION OF ADJOINT RANS EQUATIONS

To begin the derivation of the adjoint RANS equations and their boundary conditions, the cost functional is restated below:

$$\mathcal{L} = \underbrace{\mathcal{E}(\bar{\mathbf{u}})}_1 - \underbrace{\langle \bar{\mathbf{u}}^\dagger, \bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} + \nabla \bar{p} - Re^{-1} \nabla^2 \bar{\mathbf{u}} - \mathbf{f} \rangle}_2 - \underbrace{\langle \bar{p}^\dagger, \nabla \cdot \bar{\mathbf{u}} \rangle}_3, \quad (\text{A.1})$$

where each term is labeled on the right-hand side so that they can be considered separately. Taking the variation of \mathcal{L} with respect to both $\bar{\mathbf{u}}$ and \bar{p} yields the adjoint RANS equations.

The variation of term 1 with respect to $\bar{\mathbf{u}}$ is

$$\lim_{s \rightarrow 0} \frac{\mathcal{E}(\bar{\mathbf{u}} + s\delta\bar{\mathbf{u}}) - \mathcal{E}(\bar{\mathbf{u}})}{s} = \frac{\delta\mathcal{E}}{\delta\bar{\mathbf{u}}}. \quad (\text{A.2})$$

Term 2 is the most complicated:

$$\lim_{s \rightarrow 0} \frac{\langle \bar{\mathbf{u}}^\dagger, (\bar{\mathbf{u}} + s\delta\bar{\mathbf{u}}) \cdot \nabla(\bar{\mathbf{u}} + s\delta\bar{\mathbf{u}}) + \nabla\bar{p} - Re^{-1}\nabla^2(\bar{\mathbf{u}} + s\delta\bar{\mathbf{u}}) - \mathbf{f} \rangle}{s} - \frac{\langle \bar{\mathbf{u}}^\dagger, \bar{\mathbf{u}} \cdot \nabla\bar{\mathbf{u}} + \nabla\bar{p} - Re^{-1}\nabla^2\bar{\mathbf{u}} - \mathbf{f} \rangle}{s}. \quad (\text{A.3})$$

Several terms will disappear and the nonlinear term can be expanded:

$$\lim_{s \rightarrow 0} \frac{\langle \bar{\mathbf{u}}^\dagger, s\bar{\mathbf{u}} \cdot \nabla(\delta\bar{\mathbf{u}}) + s\delta\bar{\mathbf{u}} \cdot \nabla\bar{\mathbf{u}} + s^2\delta\bar{\mathbf{u}} \cdot \nabla\delta\bar{\mathbf{u}} - sRe^{-1}\nabla^2(\delta\bar{\mathbf{u}}) \rangle}{s}. \quad (\text{A.4})$$

Now take the limit:

$$-\langle \bar{\mathbf{u}}^\dagger, \bar{\mathbf{u}} \cdot \nabla(\delta\bar{\mathbf{u}}) + \delta\bar{\mathbf{u}} \cdot \nabla\bar{\mathbf{u}} - Re^{-1}\nabla^2(\delta\bar{\mathbf{u}}) \rangle. \quad (\text{A.5})$$

Term 3 is straightforward:

$$\lim_{s \rightarrow 0} \frac{\langle \bar{p}^\dagger, \nabla \cdot (\bar{\mathbf{u}} + s\delta\bar{\mathbf{u}}) \rangle - \langle \bar{p}^\dagger, \nabla \cdot \bar{\mathbf{u}} \rangle}{s} = -\langle \bar{p}^\dagger, \nabla(\delta\bar{\mathbf{u}}) \rangle. \quad (\text{A.6})$$

Finally, combine all the terms:

$$\frac{\delta\mathcal{L}}{\delta\bar{\mathbf{u}}} = \underbrace{\frac{\delta\mathcal{E}}{\delta\bar{\mathbf{u}}}}_1 - \underbrace{\left\langle \bar{\mathbf{u}}^\dagger, \underbrace{\bar{\mathbf{u}} \cdot \nabla(\delta\bar{\mathbf{u}})}_{2a} + \underbrace{\delta\bar{\mathbf{u}} \cdot \nabla\bar{\mathbf{u}}}_{2b} - \underbrace{Re^{-1}\nabla^2(\delta\bar{\mathbf{u}})}_{2c} \right\rangle}_2 - \underbrace{\langle \bar{p}^\dagger, \nabla(\delta\bar{\mathbf{u}}) \rangle}_3. \quad (\text{A.7})$$

The next step involves integrating by parts and obtaining the adjoint boundary conditions. Once again, each term is considered separately. First, consider Term 3, which includes the adjoint pressure. At this point, it is easier to switch to index notation particularly when applying Gauss's and Green's Theorem of the first kind:

$$- \langle \bar{p}^\dagger, \nabla(\delta\bar{\mathbf{u}}) \rangle = -\bar{p}^\dagger \delta\bar{u}_{j,i}, \quad (\text{A.8})$$

$$(\bar{p}^\dagger \delta\bar{u}_{j,i})_i = \bar{p}_i^\dagger \delta\bar{u}_j + \bar{p}^\dagger \delta\bar{u}_{j,i}, \quad (\text{A.9})$$

$$-\bar{p}^\dagger \delta\bar{u}_{j,i} = -(\bar{p}^\dagger \delta\bar{u}_{j,i})_i + \bar{p}_i^\dagger \delta\bar{u}_j = \boxed{-\bar{p}^\dagger n_i + \nabla\bar{p}^\dagger}. \quad (\text{A.10})$$

The boxed terms represent the final result. Only terms contributing to the boundary conditions will be left in index notation. The next easiest term to consider is Term 2b since it does not involve integration by parts:

$$-\langle \bar{\mathbf{u}}^\dagger, \delta\bar{\mathbf{u}} \cdot \nabla\bar{\mathbf{u}} \rangle = -\bar{u}_j^\dagger \delta\bar{u}_i \bar{u}_{j,i} = -\bar{u}_j^\dagger \bar{u}_{j,i} = -\bar{u}_j^\dagger (\bar{u}_{j,i})^T = \boxed{-\bar{\mathbf{u}}^\dagger \cdot \nabla\bar{\mathbf{u}}^T}. \quad (\text{A.11})$$

For Term 2a, the Divergence Theorem is applied:

$$- \langle \bar{\mathbf{u}}^\dagger, \bar{\mathbf{u}} \cdot \nabla(\delta\bar{\mathbf{u}}) \rangle = -\bar{u}_j^\dagger \bar{u}_i \delta\bar{u}_{j,i}, \quad (\text{A.12})$$

$$-(\bar{u}_j^\dagger \bar{u}_i \delta\bar{u}_{j,i})_i = -\bar{u}_{j,i}^\dagger \bar{u}_i \delta\bar{u}_j - \bar{u}_j^\dagger \bar{u}_{i,i} \delta\bar{u}_{j,i} - \bar{u}_j^\dagger \bar{u}_i \delta\bar{u}_{j,i}, \quad (\text{A.13})$$

$$-\bar{u}_j^\dagger \bar{u}_i \delta\bar{u}_{j,i} = -(\bar{u}_j^\dagger \bar{u}_i \delta\bar{u}_{j,i})_i + \bar{u}_{j,i}^\dagger \bar{u}_i \delta\bar{u}_j, \quad (\text{A.14})$$

$$\bar{u}_{j,i}^\dagger \bar{u}_i - \bar{u}_j^\dagger \bar{u}_i n_i = \boxed{\bar{\mathbf{u}} \cdot \nabla\bar{\mathbf{u}}^\dagger - \bar{u}_i \bar{u}_j^\dagger n_i}. \quad (\text{A.15})$$

The last term to consider is the Laplacian:

$$-\langle \bar{\mathbf{u}}^\dagger, -Re^{-1} \nabla^2(\delta\bar{\mathbf{u}}) \rangle = Re^{-1} \bar{u}_i^\dagger \delta\bar{u}_{i,j,j}. \quad (\text{A.16})$$

Apply Green's Theorem of the first kind to obtain:

$$Re^{-1} \bar{u}_i^\dagger \delta\bar{u}_{i,j,j} = Re^{-1} \left[-\bar{u}_{i,j}^\dagger \delta\bar{u}_{i,j} + \delta\bar{u}_{i,j} \bar{u}_i^\dagger n_j \right]. \quad (\text{A.17})$$

The $\delta\bar{\mathbf{u}}$ terms need to be gradient free so that they can be eliminated from the equation. Once again, apply the Divergence Theorem:

$$(\bar{u}_{i,j}^\dagger \delta\bar{u}_i)_j = \bar{u}_{i,j,j}^\dagger \delta\bar{u}_i + \bar{u}_{i,j}^\dagger \delta\bar{u}_{i,j}, \quad (\text{A.18})$$

$$-\bar{u}_{i,j}^\dagger \delta\bar{u}_{i,j} = \bar{u}_{i,j,j}^\dagger \delta\bar{u}_i - \bar{u}_{i,j}^\dagger \delta\bar{u}_i n_j. \quad (\text{A.19})$$

Putting all the terms together:

$$-\langle \bar{\mathbf{u}}^\dagger, Re^{-1} \nabla^2(\delta\bar{\mathbf{u}}) \rangle = \boxed{Re^{-1} \left[\delta\bar{u}_{i,j} \bar{u}_i^\dagger n_j + \nabla^2 \bar{\mathbf{u}}^\dagger - \bar{u}_{i,j}^\dagger \delta\bar{u}_i n_j \right]}. \quad (\text{A.20})$$

For now, the surface integral terms, or those which are dotted with a normal vector, are ignored (they will be dealt with subsequently to determine the appropriate boundary conditions). The variation of \mathcal{L} with respect to $\bar{\mathbf{u}}$ becomes:

$$\frac{\delta \mathcal{L}}{\delta \bar{\mathbf{u}}} = \frac{\delta \mathcal{E}}{\delta \bar{\mathbf{u}}} + \bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}}^\dagger - \bar{\mathbf{u}}^\dagger \cdot \nabla \bar{\mathbf{u}}^T + \nabla \bar{p}^\dagger + Re^{-1} \nabla^2 \bar{\mathbf{u}}. \quad (\text{A.21})$$

To minimize the cost functional Equation A.21 is set equal to zero to obtain the adjoint RANS equations:

$$\boxed{-\bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}}^\dagger + \bar{\mathbf{u}}^\dagger \cdot \nabla \bar{\mathbf{u}}^T - \nabla \bar{p}^\dagger - Re^{-1} \nabla^2 \bar{\mathbf{u}} = \frac{\delta \mathcal{E}}{\delta \bar{\mathbf{u}}}}. \quad (\text{A.22})$$

Next the variation of the cost function with respect to \bar{p} is considered. This is far more straightforward since \bar{p} only appears in the second inner product, resulting in

$$-\langle \bar{\mathbf{u}}^\dagger, \nabla(\delta \bar{p}) \rangle = \bar{u}_i^\dagger \delta \bar{p}_{,i}, \quad (\text{A.23})$$

$$(\bar{u}_i^\dagger \delta \bar{p})_{,i} = -\bar{u}_{i,i}^\dagger + \bar{u}_i^\dagger \delta \bar{p}_{,i}, \quad (\text{A.24})$$

$$\bar{u}_i^\dagger \delta \bar{p}_{,i} = -\bar{u}_{i,i}^\dagger + \bar{u}_i^\dagger \delta \bar{p} n_i. \quad (\text{A.25})$$

One obtains the condition that the divergence of $\bar{\mathbf{u}}^\dagger$ must be zero since the boundary condition term will be needed to cancel other terms:

$$\boxed{\nabla \cdot \bar{\mathbf{u}}^\dagger = 0}. \quad (\text{A.26})$$

The boundary conditions are derived from the surface integral terms from both the variation with respect to $\bar{\mathbf{u}}$ and \bar{p} . These are collected below:

$$-Re^{-1} \delta \bar{u}_{i,j} \bar{u}_i^\dagger n_j + \bar{u}_j^\dagger n_j \delta \bar{p} + Re^{-1} \bar{u}_{i,j}^\dagger \delta \bar{u}_i n_j + \bar{p}^\dagger \delta \bar{u}_j n_j = -\bar{u}_j \bar{u}_i^\dagger \delta \bar{u}_i n_j. \quad (\text{A.27})$$

Begin with the inlet boundary:

$$Re^{-1} \frac{\partial(\delta \bar{u}_i)}{\partial x} \bar{u}_i^\dagger - \bar{u}^\dagger \delta \bar{p} - Re^{-1} \frac{\partial \bar{u}_i^\dagger}{\partial x} \delta \bar{u}_i - \bar{p}^\dagger \delta \bar{u} = \bar{u} \bar{u}_i^\dagger \delta \bar{u}_i. \quad (\text{A.28})$$

Setting $\boxed{\bar{u}^\dagger = 0; \bar{v}^\dagger = 0}$ causes all boundary terms to disappear. These boundary conditions are also applied to the surface of the airfoil. (Note: A velocity term without a subscript refers to a specific component of the velocity vector, i.e., the streamwise direction \bar{u}^\dagger or vertical direction \bar{v}^\dagger .) For the symmetry boundaries:

$$-Re^{-1} \frac{\partial(\delta \bar{u}_i)}{\partial y} \bar{u}_i^\dagger + \bar{v}^\dagger \delta \bar{p} + Re^{-1} \frac{\partial \bar{u}_i^\dagger}{\partial y} \delta \bar{u}_i + \bar{p}^\dagger \delta \bar{u} = -\bar{v} \bar{u}_i^\dagger \delta \bar{u}_i. \quad (\text{A.29})$$

Since $\bar{v} = 0$ and $\frac{\partial \bar{u}}{\partial y} = 0$ from the forward boundary conditions, it is necessary that $\boxed{\partial \bar{u}^\dagger / \partial y = 0}$ and $\boxed{\bar{v}^\dagger = 0}$ on the symmetry boundaries. Finally, for the outlet boundary:

$$-Re^{-1} \frac{\partial(\delta \bar{u}_i)}{\partial x} \bar{u}_i^\dagger + \bar{u}^\dagger \delta \bar{p} + Re^{-1} \frac{\partial \bar{u}_i^\dagger}{\partial x} \delta \bar{u}_i + \bar{p}^\dagger \delta \bar{u} = -\bar{u} \bar{u}_i^\dagger \delta \bar{u}_i. \quad (\text{A.30})$$

Setting $i = 1$, one obtains:

$$-Re^{-1} \frac{\partial(\delta \bar{u})}{\partial x} \bar{u}^\dagger + \bar{u}^\dagger \delta \bar{p} + Re^{-1} \frac{\partial \bar{u}^\dagger}{\partial x} \delta \bar{u} + \bar{p}^\dagger \delta \bar{u} = -\bar{u} \bar{u}^\dagger \delta \bar{u}. \quad (\text{A.31})$$

One of the forward boundary conditions is

$$Re^{-1} \frac{\partial \bar{u}}{\partial x} - \bar{p} = 0, \quad (\text{A.32})$$

so the first and second terms in Equation A.31 can be eliminated resulting in the first boundary condition at the outlet:

$$\boxed{Re^{-1} \frac{\partial \bar{u}^\dagger}{\partial x} + \bar{p}^\dagger = -\bar{u} \bar{u}^\dagger}. \quad (\text{A.33})$$

Finally, setting $i = 2$, one obtains:

$$-Re^{-1} \frac{\partial(\delta \bar{v})}{\partial x} \bar{v}^\dagger + Re^{-1} \frac{\partial \bar{v}^\dagger}{\partial x} \delta \bar{v} = -\bar{u} \bar{v}^\dagger \delta \bar{v}. \quad (\text{A.34})$$

The first term drops out due to the forward boundary conditions and so what is left becomes the second boundary condition at the outlet for the adjoint RANS equations:

$$\boxed{Re^{-1} \frac{\partial \bar{v}^\dagger}{\partial x} = -\bar{u} \bar{v}^\dagger}. \quad (\text{A.35})$$

Appendix B

NACA 0018 AIRFOIL WITH $\alpha = 10^\circ$

Preliminary results from resolvent analysis of the data-assimilated mean profile are presented for the flow around a NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$. As discussed in Chapter 8, an improved understanding of their implications is necessary before attempting to apply the flow reconstruction methodology to this flow.

B.1 Resolvent Frequency Sweep

The resolvent norm is computed for various ω over a range similar to that of the $\alpha = 0^\circ$ case. The first two singular values are plotted in Figure B.1. Panel (a) contains the results for the experimental mean interpolated onto the FreeFem mesh while panel (b) is for the data-assimilated mean. Unlike the $\alpha = 0^\circ$ case, the results for the interpolated mean and the data-assimilated mean are significantly different. The interpolated mean has a single peak at $\omega = 5.9$ whereas the data-assimilated mean has two peaks at $\omega = 7.3$ and $\omega = 15.4$. There is a very small peak in the interpolated mean around $\omega = 15$, but there is no separation of singular values. Consequently the frequency range over which the resolvent is low-rank is larger for the data-assimilated mean.

It was also observed by Beneddine et al. (2017) that the power spectrum for the

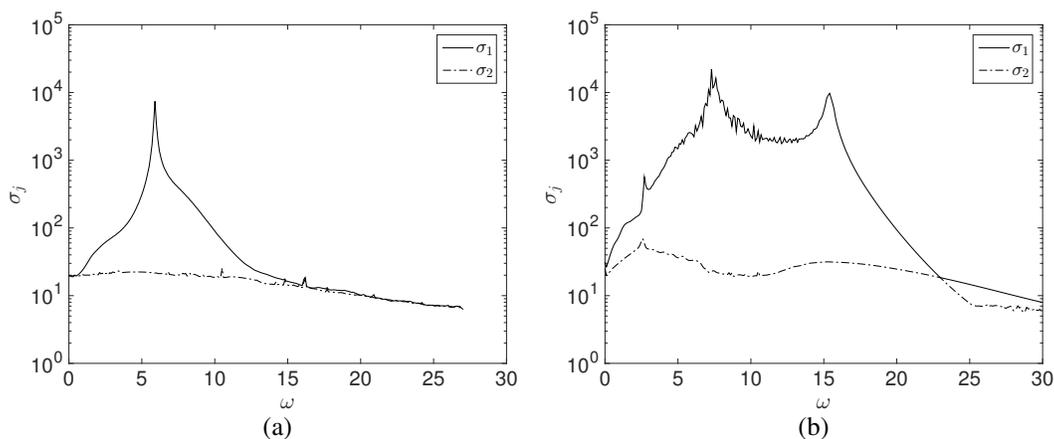


Figure B.1: First two singular values of the NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$. The mean profile was obtained from interpolating the experimental data onto the FreeFem mesh in (a) and data-assimilating the mean flow in (b).

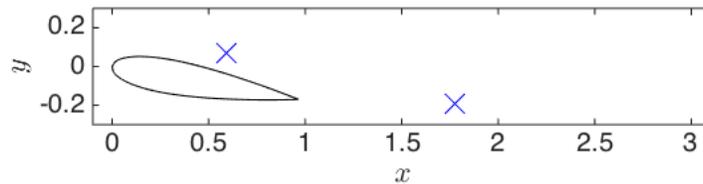


Figure B.2: Points where the power spectrum is computed for the $\alpha = 10^\circ$ flow. The upstream point is located in the shear layer while the downstream point is in the wake region.

round jet displayed peaks at two distinct frequencies. The first was a Kelvin-Helmholtz mechanism which dominated the power spectrum at locations further upstream in the flow. The second peak, which was stronger further downstream, was centered around $St = 0.38$ and corresponded to vortex pairing. It would be interesting to see whether the resolvent norm in Beneddine et al. (2017) also exhibits two peaks. The precise frequency at which the lower peak occurs for the $\alpha = 10^\circ$ flow is difficult to determine from Figure B.1(b), so it is not clear if the upper peak is a harmonic of the lower peak.

B.2 Power Spectrum of Shear Layer and Wake

The power spectrum of two points in the flow are considered to examine how the frequency content in the shear layer differs from that in the wake. The locations are specified in Figure B.2.

The power spectra are plotted in Figure B.3(a) and (b) for the shear layer and wake points, respectively. The peak for the shear layer point occurs at $\omega = 14.2$ while the peak for the wake point occurs at $\omega = 7.1$, which is a subharmonic of the shear layer frequency. These observations are quite similar to those of Beneddine et al. (2017) although it is unclear if vortex pairing occurs in this flow. It is possible that the data-assimilated mean is not sufficiently accurate near the separation point to correctly identify the true frequency in the shear layer. Recall that the interpolated mean, which has practically no reliable data in this region, does not suggest a second peak.

B.3 Resolvent Response Modes

The resolvent response modes for the two peak frequencies $\omega = 7.3$ and $\omega = 15.4$ are plotted in Figure B.4. The lower peak corresponds to a mode whose spatial

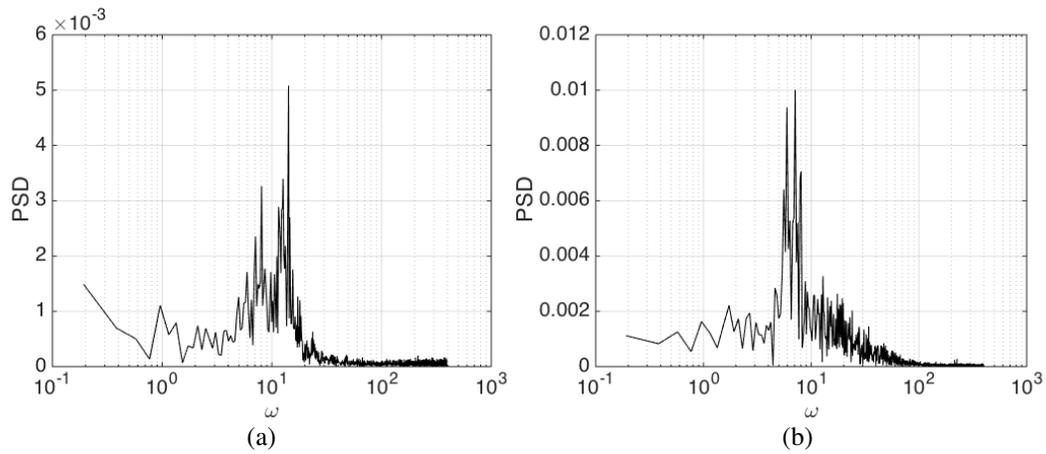


Figure B.3: PSD of a point in the shear layer (a) and wake (b) of the flow around a NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$.

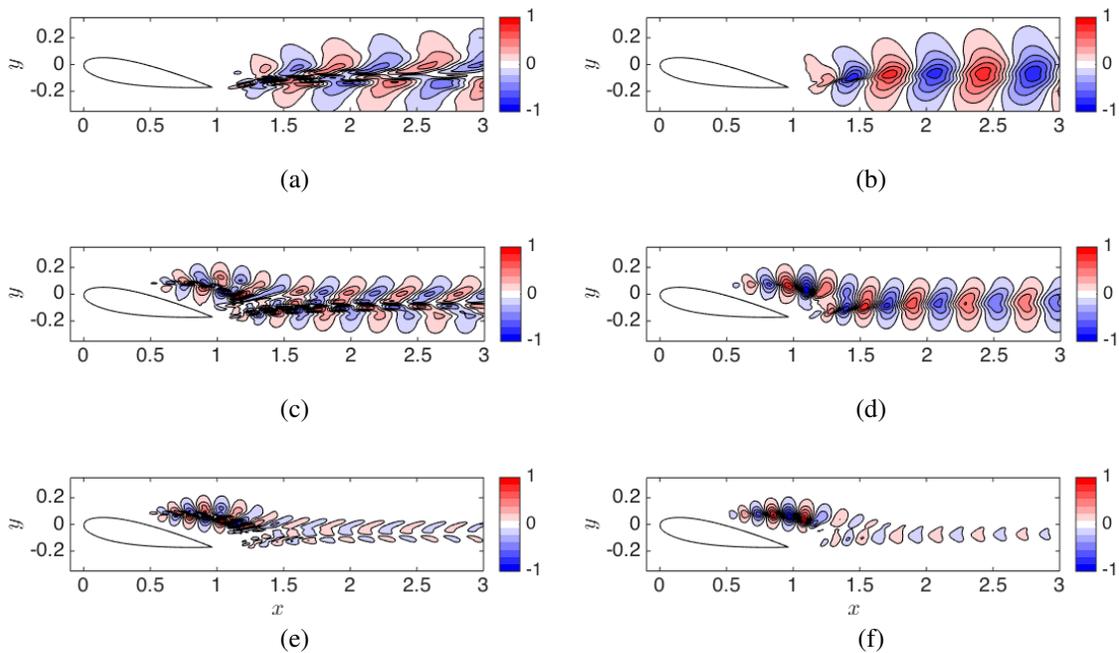


Figure B.4: Resolvent response modes computed for the flow around a NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$. The u -component is plotted in the left-hand column while the v -component is plotted in the right-hand column. The frequencies which have been plotted include $\omega = 7.3$ in (a-b), $\omega = 12.0$ in (c-d), and $\omega = 15.4$ in (e-f). The contours have normalized such that the maximum value is unity.

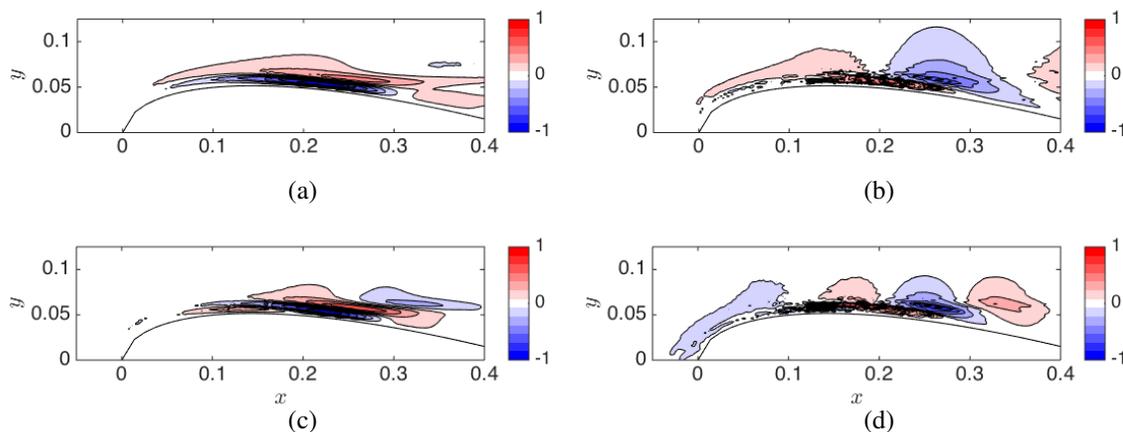


Figure B.5: Resolvent forcing modes computed for the flow around a NACA 0018 airfoil at $\alpha = 10^\circ$ and $Re = 10250$. The u -component is plotted in the left-hand column while the v -component is plotted in the right-hand column. The frequencies which have been plotted include $\omega = 7.3$ in (a-b) and $\omega = 15.4$ in (c-d). The contours have normalized such that the maximum value is unity.

support is entirely behind the airfoil. It also looks remarkably similar to the shedding modes observed for the cylinder wake and $\alpha = 0^\circ$ case. The upper peak corresponds to a mode whose spatial support is mostly confined to the shear layer region. A mode at $\omega = 12$ is also included in Figure B.4 since this frequency lies in between the two peaks. The resolvent mode contains features of both modes discussed earlier as there is energetic activity in both the shear region and the wake.

These results are encouraging since the Reynolds stresses for this flow are concentrated along the shear layer and the region behind the trailing edge of the airfoil as seen in Figure 6.17. The resolvent is identifying linear mechanisms which are responsible for the creation of Reynolds stresses needed to sustain the mean profile. What is not clear at this point is whether the mixing of these mechanisms at intermediate frequencies such as $\omega = 12$ means a rank-1 model can successfully reconstruct the flow. This is a subject of future work.

B.4 Resolving the Forcing Modes

The resolvent forcing modes are plotted for the two peaks identified in Figure B.1(b) are presented in Figure B.5. The modes are concentrated in a very small portion of the domain near the leading edge of the airfoil, similar to those computed by Thomareis (2017), who did not observe such a wide separation of peaks in the resolvent norm for different physical mechanisms. This location coincides with

the beginning of the shear layer. There is almost no overlap between the forcing modes in Figure B.5 and the response modes in Figure B.4 which is characteristic of a convective instability mechanism. The only overlap which exists for $\omega = 15.4$ occurs in the shear layer at $x = 0.5$. The u -component of the forcing is slightly more smooth than its v -counterpart which suffers from a lack of resolution near the airfoil surface. The modes for the $\alpha = 10^\circ$ case are considerably more difficult to resolve when compared to the $\alpha = 0^\circ$ case due to how concentrated they are near the airfoil surface. They do, however, lean against the mean shear which is similar to the those observed in Chapter 7.

It is of interest to note that the structure for both frequencies is remarkably similar. In fact, the forcing mode for $\omega = 15.4$ looks like a compressed version of $\omega = 7.3$ in the streamwise direction. Future work is required to understand how the effect of ω and the mean profile results in stark differences between the spatial support of the modes at each peak in the resolvent norm.